

Gas Source Localisation by Constructing Concentration Gridmaps with a Mobile Robot

Achim Lilienthal¹, Tom Duckett²

¹ W.-Schickard-Inst. of Comp. Science, University of Tübingen, D-72076 Tübingen, Germany
lilien@informatik.uni-tuebingen.de

² Department of Technology, AASS, Örebro University, 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 with a mobile robot equipped with a gas-sensitive system (“mobile nose”). By contrast to metric gridmaps extracted from sonar or laser range scans, a gas sensor measurement 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. The structure of the mapped features is discussed with respect to the parameters of the applied density function, the evolution of the gas distribution over time, and the capability to locate a gas source.

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, provided by another sensor system. Intended applications of this system include detection and localisation of a distant gas source (e.g., as an “electronic watchman” for detecting hazardous substances [1]), 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 [2, 3].

Gridmaps were originally introduced to mobile robotics in the early 1980s as a means of creating maps using wide-angle measurements from sonar sensors [4]. 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 robot’s belief that the corresponding area is occupied by any object. In our approach, the cells in the gridmap correspond to the robot’s 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 were estimated as $\tau_r \approx 1.8$ s and $\tau_d \approx 11.1$ s [5]. Thus, considerable integration of successive measurements is carried out by the sensors themselves. Furthermore, the distribution of gas molecules in an uncontrolled environment tends to be dominated by turbulence rather than diffusion, typically resulting in a jagged pattern of temporally fluctuating eddies [6, 7].

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 descreasing likelihood that a given measurement represents the true concentration with respect to the distance from the point of measurement. By integrating many measurements, the underlying structure of the gas distribution can be separated from the transient variations due to turbulence. The effects of the long response and recovery times are also reduced by incorporating measurements taken as the robot travelled in many different 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 Section 3.2. 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 [9].

2 Building Concentration Gridmaps

Gridmaps extracted from a sequence of measurements are able to represent time-constant features of the measured quantity. Therefore, a properly constructed gas concentration gridmap is expected to show time-averaged structures that might indicate the location of an odour source. In rooms with a constant unidirectional airflow these structures should be plume-like ones [10]. This paper, however, presents investigations performed in an unventilated laboratory room at Örebro University. Recent experiments in unventilated rooms (including this one) showed concentration profiles that are also relatively stable over time [1, 11]. These plume-like structures are likely to be caused by constant air streams that occur as a consequence of spatial temperature differences.

In order to create reasonable gridmaps, the cells have to be updated multiple times. This is not a problem with gridmaps built from laser or sonar scans, where each measurement provides information about an area that overlaps considerably with that of previous scans. By contrast, the gas sensor readings are very local measurements representing the concentration of the analyte at the sensors surface ($\approx 1 \text{ cm}^2$). Nevertheless these readings contain information about a larger area, for two reasons. First, although turbulence tends to create eddies in the gas distribution [6, 7], 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. Thus spatial information is integrated along the path driven by the robot.

As a first approximation, the sensor readings were assumed to represent the real concentration at the current location. The readings r_t were convolved using the two dimensional normalised Gaussian

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

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 within the cutoff radius R_{co} , around the point \mathbf{x}_t where the measurement was taken at time t , the displacement $\delta_t^{(i)}$ to the grid cell's centre $\mathbf{x}^{(i)}$ is calculated as

$$\delta_t^{(i)} = \mathbf{x}^{(i)} - \mathbf{x}_t. \quad (2)$$

- Now the weighting for all the grid cells i is determined by

$$w_t^{(i)} = \begin{cases} f(\delta_t^{(i)}) & \text{if } \delta_t^{(i)} \leq R_{co} \\ 0 & \text{if } \delta_t^{(i)} > R_{co} \end{cases} \quad (3)$$

- Then two temporary values maintained per grid cell are updated with this weighting $w_t^{(i)}$: the total sum of the weights

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

and the total sum of weighted readings

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

For the latter calculation, the normalised readings r_t were used, 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 $W_t^{(i)}$ exceeds the threshold value W_{min} , the value of the grid cell is set to

$$c_t^{(i)} = WR_t^{(i)} / W_t^{(i)} \quad \text{if } W_t^{(i)} \geq W_{min}. \quad (7)$$

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 recorded at positions x_1 - x_5 are spread by the mapping process along the driven path. While the spreading perpendicular to the path is determined mainly by the chosen parameter σ , the spreading along the path is approximately independent of this variable.




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

3 Experimental Setup

3.1 Robot and Gas Sensors

The experiments were performed with a Koala mobile robot (see Fig. 2) equipped with the Mark III mobile nose [5], comprising 6 tin oxide sensors manufactured by Figaro Engineering Inc. This type of chemical sensor shows a decreasing resistance in the presence of combustible 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. Due to their different selectivities, discrimination of different analytes is possible. For the investigations presented in this paper, however, the sensor arrays were used only to increase the robustness of the measured signal. 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.

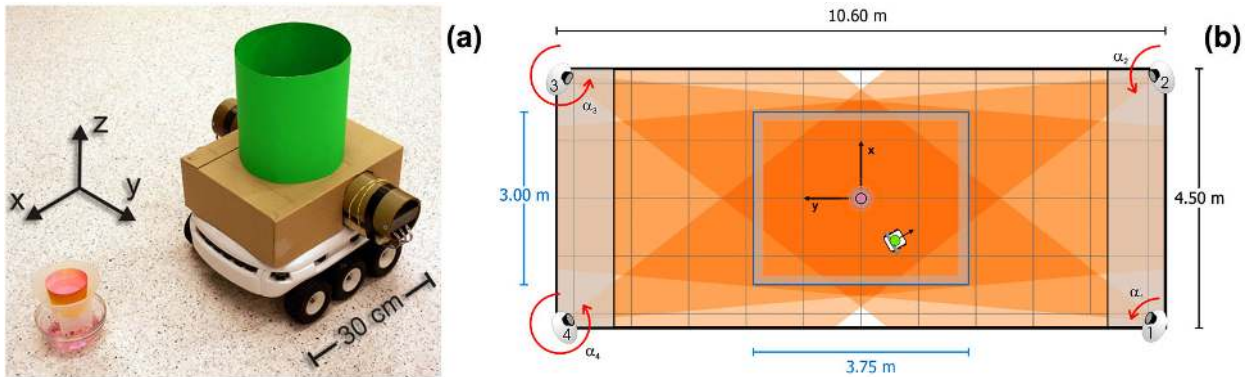


Fig. 2. (a) Koala robot with the Örebro Mark III mobile nose. The picture shows the odour source, the gas sensors inside the two suction tubes mounted at the rear of the robot and the coloured “hat” used for determining the absolute position of the robot. (b) Floor plan of the laboratory room at Örebro University where the experiments were performed.

3.2 Absolute Positioning System

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 positioning system uses four Philips PCVC 740K web-cameras 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. 2(b) 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.

3.3 Environment and Odour Source

All experiments were performed in a rectangular laboratory room at Örebro university (size $10.6 \text{ m} \times 4.5 \text{ m}$). The robot’s movement was restricted such that its centre was always located inside the central region, which is also indicated in Fig. 2(b), where precise and reliable position information is available. The air conditioning system in the room was deactivated to provide an unventilated environment.

To simulate a typical task for an electronic watchman, 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. 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.

3.4 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*, namely 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 [1, 12]. Furthermore the robot was rotated slowly ($10^\circ/s$) at the corners 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 [13]. 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*). While the reactive strategies proved to be useful in moving the robot towards the general direction of the source on average, they are not reliable enough to enable localisation of the gas source with high certainty. Moreover, it is generally not possible to declare that the gas source was found by determining an instantaneous global concentration maximum. This is why the gas concentration mapping technique introduced in this paper should be useful for smelling robots. Further details and especially the gas source localisation performance of this “smelling Braitenberg vehicle” are discussed in [6].

4 Results

4.1 Width of the Likelihood Function

To build a gridmap with the new algorithm (eqn. 1-7), a number of parameters must first be determined: the cell size, the width of the gaussian σ , the cutoff radius R_{co} and the threshold W_{min} . For all experiments presented here, grid cells of size $2.5 \times 2.5 \text{ cm}^2$ were used. The exact magnitude of R_{co} and W_{min} does not have a strong influence, but the parameter σ has a critical effect on the resulting gridmaps. Fig. 3 compares four gridmaps that result from different values of σ , using constant values of $R_{co} = 3\sigma$ and $W_{min} = 1.0 \times (\text{the number of sensors})$. All of the gridmaps were created from the same sensor data collected with all 6 sensors over a period of 100 minutes by a robot that was driven along a sequence of rectangular spirals. (One of these spirals is plotted in Fig. 3 with a thin broken line.) Concentration values are indicated by shadings of grey (dark \rightarrow low, light \rightarrow high) while the values higher than 90% of the maximum are indicated with a second range of dark-to-light shadings (of red). For small values of σ the resulting gridmaps are




Fig. 3. Comparison of concentration gridmaps from the same data set with Gaussians of different width σ (eq. 1).

dominated by local variations (see left part of Fig. 3). Increasing σ causes these maxima to be combined, and thus larger structures of the gas distribution appear as contiguous spots. A value of $\sigma = 15 \text{ cm}$ seems to provide a reasonable tradeoff between these two effects. All of the subsequent gridmaps presented were created using this parameter value.

4.2 Evolution of the Gas Distribution over Time

Due to the local character of the concentration measurements, it takes some time to build a concentration map. In addition to spatial coverage, a certain amount of temporal averaging is also necessary until the time-constant structure of the gas distribution is represented in the map. The evolution of the time-constant structures can be seen in Fig. 4, which shows two sequences of gridmaps created from data collected up to the specified time. The gridmaps were




Fig. 4. Evolution of concentration gridmaps over time. The upper row shows a sequence of gridmaps created from the data that were collected along a predefined rectangular spiral up till the specified time. In the same way the lower row depicts gridmaps based on data collected in another experiment with a Braitenberg-vehicle with uncrossed inhibitory connections.

created with the parameters $\sigma = 15$ cm, $R_{co} = 3\sigma$, $W_{min} = 6.0$ using all six sensors. Grid cells that have not yet been inspected (with $W^{(i)} \leq W_{min}$) are indicated with a different colour (green).

In the experiments where the robot moved along a predefined path (section 3.4), it took approximately 30 minutes for the mapped structures to stabilise. This can also be seen in Fig. 5(a), which shows the distance between the centre of the grid cell with the maximum value and the centre of the gas source. After about 1500 s, this distance converges either to a relatively stable value, or it alternates between two otherwise stable values.

Slightly different results were obtained if the robot was controlled reactively as a Braitenberg vehicle, as explained in section 3.4 (see lower row in Fig. 4 and Fig. 5(b)). Due to the fact that areas with high concentration were actively explored, stable structures could be determined much 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.

4.3 Localisation of a Gas Source

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. As mentioned, this assumption is not fulfilled under realistic conditions due to the relatively slow diffusion velocity of gases compared to spreading by turbulence. Nevertheless, the position of the concentration maximum can be used to estimate the approximate location of the source in some cases, as can be seen in Fig. 5. Here, the distance between such an estimate and the centre of the real position of the gas source is plotted against time, both for data collected by driving along the predefined path (a) and with a Braitenberg-type strategy (b). The velocity gain in the case of the reactively steered robot was 5 cm/s and the gas source was placed in the middle of the inspected area. Three experiments were performed with uncrossed sensor-motor connections and another one with crossed connections, indicated by the solid (red) line. An exact agreement with the position of the gas source (which




Fig. 5. Distance between the cell with the highest concentration and the centre of the gas source. The gas sensor readings were collected (a) along a predefined path and (b) applying a Braitenberg-type search strategy.

is shaded area in Fig. 5) was observed only temporarily. On the other hand the error at the end of the experiment was < 75 cm in 6 out of 7 and < 50 cm in 5 out of 7 experiments. It is not guaranteed, however, to get a good estimate of the source location with this method. This can be seen in Fig. 5(a) but was also observed in further experiments with reactive strategies.

5 Outlook

This paper presented a new technique for modelling gas distributions by constructing concentration gridmaps with a mobile robot. Preliminary experimental results for two different exploration strategies were presented. 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 will include development of an actual source-finding strategy based on these self-acquired maps.

References

1. Achim Lilienthal, Michael R. Wandel, Udo Weimar, and Andreas 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.
2. Hiroshi Ishida, K. Suetsugu, Takamichi Nakamoto, and Toyosaka 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.
3. R. Andrew Russell, David Thiel, Reimundo Deveza, and Alan Mackay-Sim. A Robotic System to Locate Hazardous Chemical Leaks. In *IEEE Int Conf. Robotics and Automation (ICRA 1995)*, pages 556–561, 1995.
4. Martin C. Martin and Hans P. Moravec. Robot evidence grids. Technical Report CMU-RI-TR-96-06, The Robotics Institute, Carnegie Mellon University, 1996.
5. Achim Lilienthal and Tom 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.
6. Achim Lilienthal and Tom Duckett. Experimental Analysis of Smelling Braitenberg Vehicles. In *Proceedings of the IEEE International Conference on Advanced Robotics (ICAR 2003)*, Coimbra, Portugal, 2003.
7. R. Andrew Russell. *Odour Sensing for Mobile Robots*. World Scientific, 1999.
8. Achim Lilienthal and Tom Duckett. An Absolute Positioning System for 100 Euros. In *Proceedings of the IEEE International Workshop on Robotic Sensing (ROSE 2003)*, Örebro, Sweden, 2003.
9. 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.
10. J. O. Hinze. *Turbulence*. McGraw-Hill, New York, 1975.
11. Michael R. Wandel, Achim Lilienthal, Tom Duckett, Udo Weimar, and Andreas 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.
12. Ahmed Mohamad Farah and Tom 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.
13. Valentino Braitenberg. *Vehicles: Experiments in Synthetic Psychology*. MIT Press/Bradford Books, 1984.