

Experimental Analysis of Smelling Braitenberg Vehicles

Achim Lilienthal

Tom Duckett

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

Department of Technology, AASS
Örebro University
S-70182 Örebro, Sweden
tdt@tech.oru.se

Abstract

This paper addresses the problem of localisation of a static odour source in an unstructured indoor environment by a mobile robot using electrochemical gas sensors. In particular, reactive localisation strategies based on the instantaneously measured spatial concentration gradient are considered. In contrast to previous works, the environment is not artificially ventilated to produce a strong constant airflow, and thus the distribution of the odour molecules is dominated by turbulence. An experimental set-up is presented that enables different strategies for odour source localisation to be compared directly in a precisely measured experiment. Two alternative strategies that utilise a direct sensor-motor coupling are then investigated and a detailed numerical analysis of the results is presented, including tests of statistical significance. Both strategies proved to be useful to accomplish the localisation task. As a possible solution to the problem of detecting that the odour source - which is usually not corresponding to the global concentration maximum - was found, one of the tested strategies exploits the fact that local concentration maxima occur more frequently near to the odour source compared to distant regions.

1 Introduction

Chemical sensing entered the field of mobile robotics in the beginning of the 1990's. Using electrochemical sensors on a mobile robot is very promising for a broad range of applications. For example, chemical sensing can be useful for an "electronic watchman" to detect, localise and identify odours thus indicating problems such as leaking solvents, hazardous gases or a fire at its initial stage.

The main problem with using gas sensors in real world environments is that the distribution of odourant molecules is dominated usually by turbulence rather than diffusion, which is known to be a considerably slower transport mechanism for gases in general [9]. This point is illustrated by Fig. 1 which shows typical sensor readings in the vicinity of an odour source (ethanol). In this experiment, the robot

passed the source along a straight line at low speed in order to measure the spatial distribution of the analyte accurately. The curve in Fig. 1 indicates that the turbulent gas distribution creates many local concentration maxima. Remarkably the absolute maximum is usually *not* located near an odour source if this source has been active for some time. Additionally the gas distribution varies with time. Due to these physical properties, localising an odour source in an uncontrolled environment is an extremely difficult task. Research in this field is, however, worth the effort not only because it will improve the merits of an electronic watchman but also because this development is likely to be accompanied by a deeper understanding of the physical properties of turbulent motion as well as of the way animals use odours for navigation purposes.

Most work on chemical sensing for mobile robots assumes an experimental setup that minimizes the influence of turbulent transport by either shortening the source-to-sensor distance in *trail following* [11, 13, 12] or by assuming an additional constant airstream in the environment [5, 10].




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



Figure 2: The Örebro Mark III mobile nose. Two sets of 3 gas sensors were mounted inside the two suction tubes mounted at the rear of a Koala mobile robot. The picture also shows the odour source and the coloured “hat” used for determining the absolute position of the robot.

By contrast, the intention of our work is to enable a mobile robot to perform the tasks of detection and localisation of an odour source without being restricted to an environment with a dominant constant airflow [2, 8, 7, 3]. This paper especially addresses the applicability of reactive localisation techniques based on an instantaneously measured spatial gradient. It presents a detailed statistical evaluation of localisation strategies that use a direct sensor-motor coupling. Such systems are known as Braitenberg vehicles due to the famous thought experiments of Valentino Braitenberg [1]. In his book the author mentioned utilising a sense of smell as an example. But so far no evaluation based on real implementation on a mobile robot that navigates guided by airborne chemicals is available to the best of our knowledge.

2 Experimental Setup

2.1 Robot and Gas Sensors

The experiments were performed with a Koala mobile robot (see Fig. 2) equipped with 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 TGS2600, TGS2610 and TGS2620) inside two separate tubes containing a suction fan each. Due to their different selectivities discrimination of different odours is possible. For the investigations presented in this paper, however, the sensor arrays were




Figure 3: Absolute positioning system with 4 cameras. The figure shows a floor plan of the laboratory room and the outline of the region in which the localisation experiments were performed. Also plotted are the fields of view for each camera, shaded according to the number of cameras which can sense a particular region.

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.

2.2 Absolute Positioning System

To record the true position of the robot for the experimental analysis, a vision-based positioning system was developed which tracks a distinctly coloured object mounted on top of the robot. Four Philips PCVC 740K web-cameras (resolution 320×240) were mounted in the corners of the room (see Fig. 3). Each camera first computes an estimate of the angle φ_i to the centre of the coloured object. For each combination of two cameras that can actually “see” the whole coloured object, an estimate of the position \vec{p}_{ij} of that object is then calculated by triangulation. With N cameras up to $N(N-1)/2$ valid position estimates are produced at each time interval, which are then combined to determine a final position estimate \vec{p} . The parameters of the cameras (heading α_i , coordinates X_i, Y_i and angular range $\Delta\alpha_i$) were determined by an initial calibration process that minimizes the average distance \bar{d} between measured and known positions of several locations at which the coloured object is placed ($\bar{d} \approx 1 \text{ cm}$).

In addition, the current heading ϑ of the robot is estimated when the robot is moving at non-zero speed, by combining the estimates from the robot’s odometry with estimates produced by measuring the tangent to the robot’s path $\Delta\vec{p}_t = \vec{p}_t - \vec{p}_{t-1}$.

2.3 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

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 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 50ml/h. Ethanol was used because it is non-toxic and easily detectable by the tin oxide sensors.

3 Experiments

3.1 Braitenberg Vehicle

The term Braitenberg vehicle is often used to refer to steering architectures with a direct sensor-motor coupling (see Fig. 4). In his book Braitenberg explains which kind of behaviour results for these vehicles (denominated as type 2, 3 and 4) by using different classes of intermediate transfer functions. This paper partic-




Figure 4: Schematic view of Braitenberg vehicles with a direct sensor-motor coupling.

ularly concerns inhibitory connections that apply a monotonous transfer function. In this way maximum wheel speeds result if the sensed concentration is low, which in turn implements a simple sort of exploration behaviour. On the other hand the robot is slowed down by high concentrations of the analyte.

With uncrossed connections the wheel on the side that is stimulated more is driven slower and therefore the robot turns to this side. This behaviour was called *permanent love* by Braitenberg [1] because this sort of vehicle tends to move to a source of stimulation and stay near it in theory. Note that “high concentration” or “stimulation” in this context always means “high sensor values” and that these values do of course not reflect the actual concentration directly, due to the non-zero response and recovery time of the used sensors.

With crossed inhibitory connections and a monotonous transfer function the robot is also slowed

down by increased sensor responses but will in contrast turn away from them. Accordingly, this kind of behaviour was called *exploring love* by Braitenberg [1] because such a vehicle tends to stay at locations that are nearby a maximum response but continues to wander if another maximum comes into focus. Again this statement applies to a system with ideal sensors that moves guided by a smooth distribution peaked just at the actual location of a gas source.

3.2 Sensor Preprocessing

The sensor-motor wiring realises a transfer function $v(x)$ that determines the speed of the connected wheel from the sensed quantity x . But how exactly can this value x be calculated for different gas sensors? This is especially important as metal oxide sensors are known to show seasonal and environmental drift as well as noticeable differences between individual sensors [4]. In this paper a dynamically maintained normalisation of the measured conductance values r_i to the range of [0,1] was chosen. Both the minimum and maximum values were constantly updated and used to calculate the normalised response x_i for each sensor as

$$x_i^{(t)} = \frac{r_i^{(t)} - r_{min,i}^{(t)}}{r_{max,i}^{(t)} - r_{min,i}^{(t)}}. \quad (1)$$

It has to be considered that the normalisation range gets wider and might not cover the actual range of values with time. This causes changes in response to be less pronounced in x . To avoid this problem the normalisation range is dynamically trimmed by means of increasing the minimum and decreasing the maximum value in eqn.1 by a fixed fraction of the normalisation range - Δx_{min}^{trim} , Δx_{max}^{trim} respectively - and constantly repeating this procedure after each Δt^{trim} seconds. During the experiments described in this paper the values $\Delta x_{min}^{trim} = \Delta x_{max}^{trim} = 1\%$ and $\Delta t^{trim} = 30s$ were used.

Finally the normalised response values belonging to one side of the robot were combined by averaging.

3.3 A Testbed for Localisation Strategies

Due to the unpredictable and changing structure of the actual gas distribution, it is apparent that single experiments are not sufficient to derive meaningful conclusions about the performance of a particular localisation strategy. For this reason all localisation methods were tested repeatedly in a testbed given by the following scenario. A 3.75 m \times 3 m field was defined by establishing virtual walls. These boundaries were realised by assigning an artificial potential field [6] that effects a repellent pseudo-force which increases linearly with the penetration depth and starts to be effective at a distance of 20 cm. Both the virtual walls and the area where the repellent pseudo-force is active are shown in Fig. 3.

Now the robot can move freely within this virtual field, and every time it is reflected by one of the walls this is counted as a *wall hit event*. Next, an odour source is placed at a known position inside the field. This might be a real source or just an assumed one for reference tests. Then a series of experiments is performed with this configuration as follows:

- set the robot to a random starting position inside the virtual field (with a clearance of at least 100 cm to the center of the source),
- rotate the robot to a random initial heading,
- start to move the robot controlled by the particular strategy to be tested, and
- count a successful try and restart if the robot enters the obstacle clearance area around the odour source.

These steps are repeated for a fixed amount of time while the actual position and the sensor readings are logged constantly for evaluation purposes.

4 Results

Fig. 5 shows the starting position and the resulting path of two typical runs with uncrossed connections (“permanent love”). In these experiments, the linear transfer function

$$v(x) = K_v(1 - x) \quad (2)$$

with $K_v = 5\text{cm/s}$ was used and the source was placed in the middle of the virtual field. Frequently the robot could localise the source in a strikingly straightforward way, as in the example of Fig. 5(a). But quite often the Braitenberg vehicle was also misled by other local maxima and made “decisions” that appear to be exactly the wrong ones to an external observer. Moreover the results show that changes in the gas distribution can cause a completely different behaviour at the same position. This is illustrated in Fig. 5(b). When the robot first reached the location at which it finally managed to turn towards the source, hardly any reaction was obtained. A few minutes later the source was found directly from almost exactly the same spot.

Again, these examples make clear that a statistical evaluation is needed to judge the performance of different localisation strategies. For the results presented here a total of 36.5 hours of localisation experiments were performed in which the robot drove almost 5 kilometers. All these experiments were conducted in the same room (see Fig. 3) whereas the environmental conditions were varied by partly opening the doors on either side. Several of the derived statistics are listed in Table 1. They are discussed below in terms of the




Figure 5: Examples of the driven path of a Braitenberg-Vehicle with uncrossed (1-x)-connections (permanent love). The plot shows starting position and initial heading (arrow), the measured location of the robot’s center (circle) and both its front corners (small dots), the virtual repellent walls (broken line) and two circles indicating the location of the source.

average path length the robot needs to find the odour source, the average distance to the source, the average driving speed, the average number of wall hit events during the successful trials, and the total path length covered with a particular strategy.

The values obtained in the table need to be validated in order to ensure that the odour source is found using the localisation strategy under investigation and not just by coincidence. In order to do this, the localisation experiments were repeated without an odour source. When the robot moved into the area assigned to be the source, this was counted as a successful trial. Thus the robot moved essentially like a ball on a billiard table. Assuming a virtual source in the middle of the field, such experiments yielded an average path length of 9.67 m. Comparing that value with the corresponding average path length of a vehicle with uncrossed (1-x)-connections in a student-t test reveals no significant difference ($p_{H_0} = 0.4458$). This does not necessarily indicate that the chosen strategy doesn’t improve localisation performance, because it might be also a consequence of the prominent source position which is frequently found by random search.

Therefore an additional set of experiments was conducted where the source was placed at a location near a corner of the field (15 cm away from the beginning of the repellent wall potential - both along the x- and y-axis). For each corner a total of approximately 3 hours of localisation trials were performed both with and without a source. These experiments showed a highly significant improvement in localisation performance in terms of the average path length (Student’s t-test: $p_{H_0} = 0.0002$). This also holds if no normal

Source	Strategy	$K_v [\frac{cm}{s}]$	av. path [m]	av. dist. [cm]	av. wall hits	av. speed/ K_v	tot. path [m]
Middle	Ref (1-x)	5	9.67 ± 7.66	121.9 ± 19.8	3.49 ± 2.75	97.8%	319.0
	PL, 1-x	5	8.49 ± 7.93	136.7 ± 44.9	2.69 ± 2.59	73.4%	1044.0
	EL, 1-x	5/3	51.00 ± 37.36	145.4 ± 13.0	23.17 ± 16.71	75.5%	612.1
Corner	Ref (1-x)	5	20.46 ± 19.38	218.7 ± 33.7	7.24 ± 6.16	97.6%	1554.9
	PL, 1-x	5/3	11.69 ± 11.22	187.6 ± 47.5	5.11 ± 4.49	77.1%	1251.1

Table 1: Statistics of the localisation experiments. The first three columns reference which strategy was tested, while the remaining columns itemise the average path length to find the odour source, the average distance from the source during the search, the average number of wall hit events before the source was found and the average speed of the robot. In the last column the total length of the path driven under control of that particular strategy is given. The applied strategies are referred to as Ref (reference random search), PL (uncrossed connections, “permanent love”) and EL (crossed connections, “endless love”). In cases where different speed gains K_v were tested both of them are given separated by a slash.

distributed observations were assumed by means of performing a distribution-free Wilcoxon two sample test ($p_{H_0} = 0.0005$).




Figure 6: Examples of the driven path of a Braitenberg-Vehicle with crossed (1-x)-connections (permanent love). See Fig. 5 for details of illustration.

With crossed connections a completely different behaviour results. Although the robot is expected to stay near the source and thus collisions should not be unlikely in theory, the robot managed to avoid the source most of the time (see Fig. 6). The difference compared to the trials with uncrossed connections is apparent and can be clearly proven by statistical tests (both Student’s t-test and Wilcoxon two sample test revealed that $p_{H_0} < 0.00001$). This strategy provides also a localisation facility that is illustrated in Fig. 7, which shows the robot’s path recorded over 3 hours. Here the location of the source is clearly indicated by the part of the picture that was not covered by the robot. (Notice that the density of path lines between the outer and the inner circle would be considerably higher if the odour source was just an obstacle.) The good performance obtained, in terms of this particular localisation method can be explained as follows. The

robot explores the available space and evades the local concentration maxima. Because there exist many local maxima, it is difficult to find the one that corresponds to the actual location of the source by using a hill-climbing strategy (as with uncrossed connections). However, because the concentration maxima occur more frequently near to the odour source, the density of path lines in the vicinity of the source remains comparatively low using the second strategy (crossed connections), as shown in Fig. 7.




Figure 7: Path of a Braitenberg-Vehicle with crossed (1-x)-connections during 3 hours including 5 single trials (collisions are marked by a star). Note that such a plot - i.e., the area which is not covered by the robot - may provide an alternative method to localise the source.

It is worth mentioning that there are additional reasons to prefer a localisation strategy that is based on exploration and concentration peak avoidance. Avoiding high concentrations might be necessary, for example, because the odourant is harmful or in other ways offensive to the robot itself. Furthermore, applying

such a strategy can prevent the robot from wetting its wheels in cases where the odour source is a dripping liquid, which is generally not desirable because the robot would then effectively become an odour source too.

5 Conclusions and Future Work

This paper is concerned with the task of localisation of a static odour source by a mobile robot using electrochemical gas sensors.

Two Braitenberg-type strategies were investigated and both were shown to be useful for localisation. With uncrossed connections the average path length the robot needs to move to the source is reduced by a factor of two compared to random search. For real world applications this strategy has to be extended by an additional mechanism to detect that the odour source has been found. This mechanism might be provided by adding other sensors that provide clues on possible sources, for example, by recognising a beaker or a puddle by vision.

With crossed connections the robot evades each local concentration maximum including the one that is caused by the source. Due to the fact that they occur more frequently near the odour source the recorded path of the robot covers the whole available area except near the actual location of the source. Although this strategy requires more time it may often be preferable because it provides a possibility to recognize an odour source without using additional sensors. Furthermore direct contact of the robot with the volatile or liquid substance to be detected can be diminished in this way.

Future work will utilise the introduced testbed to analyse the benefits of further localisation strategies. This will include extended Braitenberg-type strategies as well as not purely reactive ones. The presented reactive strategies can be improved in two ways: Either by optimizing the applied set of parameters (Δt^{trim} , Δx_{min}^{trim} , Δx_{max}^{trim} , K_v) or by using nonlinear transfer functions that are likely to be better suited considering the non-linearity of the sensor response.

Acknowledgments This work was sponsored with a Marie Curie grant that was part of the European Commission's 5th framework programme (FP5), which is gratefully acknowledged by the author. All the small problems that arose during the preparation for the described experiments could be fastly fixed due to the kind help of Alexander Skoglund, Per Sporrang, Kevin LeBlanc and Grzegorz Cielniak. Finally the fruitful discussions with Amy Loutfi and Boyko Iliev greatly helped with all the remaining experimental and theoretical problems. Thanks to all these helpful colleagues.

References

- [1] V. Braitenberg. *Vehicles: Experiments in Synthetic Psychology*. MIT Press/Bradford Books, 1984.
- [2] T. Duckett, M. Axelsson, and A. Saffiotti. Learning to Locate an Odour Source with a Mobile Robot. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 2001)*, Seoul, South Korea, May, 21–26 2001.
- [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] J. W. Gardner and P. N. Bartlett. *Electronic Noses - Principles and Applications*. Oxford Science Publications, Oxford, 1999.
- [5] 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.
- [6] 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.
- [7] A. J. Lilienthal, M. R. Wandel, U. Weimar, and A. Zell. Experiences Using Gas Sensors on an Autonomous Mobile Robot. In *Proceedings of EU-ROBOT 2001, 4th European Workshop on Advanced Mobile Robots*, pages 1–8, Lund, Sweden, September, 19–21 2001. IEEE, IEEE Computer Press.
- [8] A. J. Lilienthal, M. R. Wandel, U. Weimar, and A. Zell. Sensing Odour Sources in Indoor Environments Without a Constant Airflow by a Mobile Robot. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 2001)*, pages 4005–4010, Seoul, South Korea, May, 21–26 2001.
- [9] T. Nakamoto, H. Ishida, and T. Moriizumi. A Sensing System for Odor Plumes. *Analytical Chem. News & Features*, 1:531–537, August 1999.
- [10] 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.
- [11] 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.
- [12] 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. Wilson, editors, *Proceedings of the 5th Conference on Simulation of Adaptive Behaviour*, pages 195–204, 1998.
- [13] 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.