



<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper presented at *Autonome Mobile Systeme, AMS, Karlsruhe, Germany, 4-5 December, 2003*.

Citation for the original published paper:

Lilienthal, A J., Duckett, T. (2003)

Approaches to gas source tracing and declaration by pure chemo-tropotaxis

In: Rüdiger Dillmann, Heinz Wörn, Tilo Gockel (ed.), *Autonome Mobile Systeme 2003* (pp. 161-171). Springer Berlin/Heidelberg

https://doi.org/10.1007/978-3-642-18986-9_17

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:oru:diva-4002>

Approaches to Gas Source Tracing and Declaration by Pure Chemo-Tropotaxis

Achim Lilienthal¹ and Tom Duckett²

¹ WSI, University of Tübingen, Sand 1, 72074 Tübingen, Germany,
lilien@informatik.uni-tuebingen.de

² AASS, Örebro University, Fakultetsgatan 1, S-70182 Örebro, Sweden,
tdt@tech.oru.se

Abstract. This paper addresses the problem of localising a static gas source in an uncontrolled indoor environment by a mobile robot. In contrast to previous works, especially the condition of an environment that is not artificially ventilated to produce a strong unidirectional airflow is considered. Here, the propagation of the analyte molecules is dominated by turbulence and convection flow rather than diffusion, thus creating a patchy distribution of spatially distributed eddies. Positive and negative tropotaxis, based on the spatial concentration gradient measured by a pair of electrochemical gas sensor arrays, were investigated. Both strategies were implemented utilising a direct sensor-motor coupling (a Braitenberg vehicle) and were shown to be useful to accomplish the gas source localisation task. As a possible solution to the problem of gas source declaration (the task of determining with certainty that the gas source has been found), an indirect localisation strategy based on exploration and concentration peak avoidance is suggested. Here, a gas source is located by exploiting the fact that local concentration maxima occur more frequently near the gas source compared to distant regions.

1 Introduction

Unlike visual or auditory stimuli, chemical stimuli are not inherently directional. In order to achieve spatial chemo-orientation an animal has therefore to determine a concentration gradient from comparing either successive stimuli (klinotaxis) or simultaneously sensed intensities from two or more receptors (tropotaxis). As pointed out by Grasso in Holland and McFarland [9] the traditional framework for understanding animal orientation by Fraenkel and Gunn [5] is limited because it assumes a uniform gradient of the distribution. This assumption is usually not fulfilled for gas distributions in natural environments. Due to the low diffusion velocity of gases at room temperature [20], the dispersal of an analyte gas is dominated by turbulence and the prevailing air flow rather than diffusion [29]. A real gas distribution therefore reveals many discontinuous patches of local eddies [8, 23], thus creating local concentration maxima that can mislead a gradient following strategy. Moreover, the absolute maximum of the instantaneous distribution is usually *not* located near the gas source if this source has been active for some time [16].

Localising a gas source is therefore a challenging task. Considering the chaotic nature of turbulent gas transport, it is clear that the same applies also to the sub-problem

of source declaration. Because the source isn't necessarily featured by the highest intensity, the information that the source is in the immediate vicinity with high certainty has to be determined from the sensed concentration pattern. Investigating these problems can lead to a deeper understanding of the physical properties of turbulent motion, as well as the way in which 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 minimising the source-to-sensor distance in *trail following* [22, 26–28] or by assuming a strong unidirectional airstream in the environment. Primarily a strong airstream can be used to get additional information about the local wind speed and direction from an anemometer. Thus strategies become feasible that utilise the instantaneous direction of flow as an estimate of source direction [2] by combining gas searching behaviours with periods of upwind movement [7, 10, 21, 25]. Furthermore, the condition of a constant unidirectional airstream implies a more structured distribution of the odorant. For such situations it is possible to model the time-averaged spread of gas [8] because the complexity of turbulent air movement can be described with a diffusion-like behaviour ruled by an additional diffusion coefficient (see [11]). 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 [11], the bending of an artificial whisker [24] or the influence on the speed of a small rotating paddle [21], reliable readings can be obtained only for wind speeds in the order of at least 10 cm/s.

The intention of this work is to enable a mobile robot to trace a gas source and declare whether the source has been found without being restricted to an environment with a strong unidirectional airflow (as in [3, 4, 17, 18]). This paper especially addresses the applicability of reactive localisation techniques that use a direct sensor-motor coupling based on an instantaneously measured spatial concentration gradient (pure chemotaxis). Such systems are known as Braitenberg vehicles due to the influential 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 our knowledge.

The rest of this paper is structured as follows: first, the experimental setup and the experiments performed are described in Sections 2 and 3. The corresponding experimental results are then discussed in Section 4 followed by conclusions and suggestions for future work (Section 5).

2 Experimental Setup

2.1 Robot and Gas Sensors

The experiments were performed with a Koala mobile robot (see Fig. 1) equipped with the Mark III mobile nose [14], comprising 6 tin oxide sensors manufactured by Figaro. 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. Namely the low selectivity, the comparatively high power consumption (caused by the heating device) and a weak durability.



Fig. 1. Koala Robot with the Örebro Mark III mobile nose.

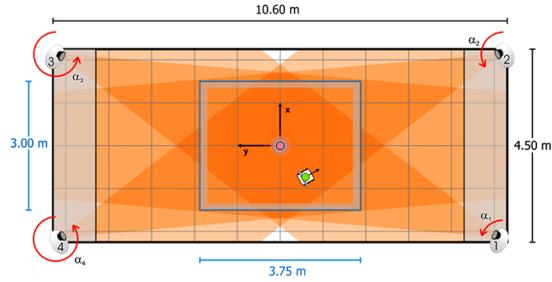


Fig. 2. Floor plan of the used laboratory room together with an outline of the area inspected.

Further on, metal oxide sensors are subject to a long response time and an even longer recovery time [14]. However, this type of gas sensor is most often used for mobile noses because it is inexpensive, highly sensitive and relatively unaffected by changing environmental conditions like room temperature or humidity. 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.

2.2 Absolute Positioning System

To record the position of the robot, the vision-based absolute positioning system W-CAPS [15] was applied, which tracks a distinctly coloured object mounted on top of the robot (the cardboard “hat” shown in Fig. 1). 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. 2 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 by 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. Further details are given in [15].

2.3 Environment and Gas Source

All experiments were performed in a rectangular laboratory room at Örebro University (size $10.6 \times 4.5 \text{ m}^2$). 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. 1). 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 Experiments

3.1 Implementation

Both the positive and the negative tropotactic localisation strategies, were implemented utilising a direct sensor-motor coupling. This kind of steering architecture is frequently referred to as a Braitenberg vehicle. In his famous book [1] 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 assuming a uniform gradient. In this paper inhibitory connections that apply a monotonic transfer function were used. Thus maximum wheel speed results 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 and a monotonic transfer function the wheel on the side that is stimulated more is driven slower and therefore the robot effectively turns to this side. This behaviour was called *permanent love* by Braitenberg because it is supposed to move the vehicle to a source of stimulation and stay near this source in theory. Note that “high concentration” or “stimulation” in this context always means “high sensor values” and that these values do not reflect the actual concentration directly due to the non-zero response and the strong memory effect caused by the long recovery time of the metal oxide sensors (Section 2.1, the response characteristics of the Mark III mobile nose are also discussed in [14]).

With crossed inhibitory connections and a monotonic transfer function the robot is also slowed down by elevated sensor responses but will in contrast turn away from them. Such a vehicle tends to stay at locations nearby a maximum of stimulation, too, but continues to wander if another maximum comes into focus. Accordingly, Braitenberg called this kind of behaviour *exploring love*. Again this description applies to a system with ideal sensors that moves guided by a smooth distribution peaked just at the actual location of existing sources.

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 . To calculate the value of x the raw sensor readings r_i were normalised to the range of [0,1]. In order to compensate for the sensitivity mismatch of individual sensors as well as for seasonal and environmental

drifts, a dynamically maintained normalisation 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 each Δt^{trim} seconds by increasing the minimum and decreasing the maximum value in eq. 1 by a fixed fraction of the normalisation range as

$$r_{min,i}^{(t)} = \tilde{r}_{min,i}^{(t)} + \Delta_{min}^{trim} (\tilde{r}_{max,i}^{(t)} - \tilde{r}_{min,i}^{(t)}), \quad (2)$$

$$r_{max,i}^{(t)} = \tilde{r}_{max,i}^{(t)} - \Delta_{max}^{trim} (\tilde{r}_{max,i}^{(t)} - \tilde{r}_{min,i}^{(t)}). \quad (3)$$

Here, $\tilde{r}_{min,i}^{(t)}$ and $\tilde{r}_{max,i}^{(t)}$ refer to the minimum and maximum value at time t before trimming. Finally the normalised response values belonging to one side of the robot were combined by averaging as

$$x_L^{(t)} = \sum_{i=1}^{N_L} x_i^{(t)} / N_L, \quad x_R^{(t)} = \sum_{i=1}^{N_R} x_i^{(t)} / N_R, \quad (4)$$

where N_L and N_R are the number of sensors contained in the sensor array on the left and right side respectively.

For the experiments presented, $N_L = N_R = 3$ sensors and the trimming parameters

$$\Delta_{min}^{trim} = \Delta_{max}^{trim} = 1\% \quad \text{and} \quad \Delta t^{trim} = 30 \text{ s}$$

were used, which were found to be suitable in an initial test sequence.

3.3 A Testbed for Localisation Strategies

Both Braitenberg-type strategies considered were tested repeatedly with the following scenario. A $3.75 \text{ m} \times 3 \text{ m}$ field was defined by establishing virtual walls. These boundaries were realised by assigning an artificial potential field [12] that effects a repellent pseudo-force, which increases linearly with the penetration depth and starts to be effective at a distance of 20 cm. The experiments were performed in the central region of the room where precise and reliable position information is available (see Fig. 2).

Now the robot could move freely within this virtual field while being constantly tracked by the absolute positioning system. Next, a gas source was placed at a known position inside the field. This could be a real source or just an assumed one for reference tests. Then a series of experiments was 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 centre of the source),
- rotate the robot to a random initial heading,

- start to move the robot controlled by the particular strategy to be tested,
- count a successful try and restart if the robot enters the obstacle clearance area around the gas source.

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

4 Results

The following sections present a discussion of the behaviour of gas-sensitive Braitenberg vehicles with uncrossed (“permanent love”) and crossed inhibitory connections (“exploring love”). Typical trajectories are shown in Figures 3 and 4. Here, the path of the robot’s centre is indicated by a hollow circle, while the position of the front corners is plotted using small dots. The starting position and the initial heading of the corresponding trial are marked by an arrow, which originates from the starting position. Also shown are the virtual repellent walls (broken line) that enclose the area where the repellent force increases with the penetration depth of the robot. Finally, the clearance area of the gas source is shown by two circles. A trial was stopped if one of the front corners of the robot entered the inner circle. The outer circle is derived by obstacle growing [13], assuming a circular shaped robot. The radius was chosen to be the minimum distance between the centre of the robot and the source. Note that because the robot actually has a rectangular profile, the outer circle provides just an approximate notion of the obstacle’s boundary with respect to the centre of the robot.

In the according experiments the linear transfer function

$$v(x) = K_v(1 - x) \quad (5)$$

with a velocity gain K_v of either 5 cm/s or 3 cm/s was used and the source was placed in the middle of the virtual field. Frequently the robot could reach the source in a strikingly straightforward way, as in the example of Fig. 3 (a). But quite often the Braitenberg vehicle was also misled by transient concentration maxima and made “decisions” that appear to be exactly the wrong ones to an external observer. This is illustrated in Fig. 3 (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.

The statistical evaluation of the experiments, with the gas source placed in the middle of the testbed area, is summarised in Table 1 in terms of the average path length the robot needs to find the gas source, the average distance to the source, the average driving speed and the total path length covered with a particular strategy. All the runs were conducted in the same room (see Fig. 2) and a total of 36.5 hours of localisation experiments was performed where the robot drove almost 5 kilometers.

To validate the results, reference tests were performed without a gas source present. Because the sensor readings were not considered during these tests, the robot moved basically like a ball on a billard table. As in all other experiments, a successful trial was counted when the robot entered the area assigned to be the source.

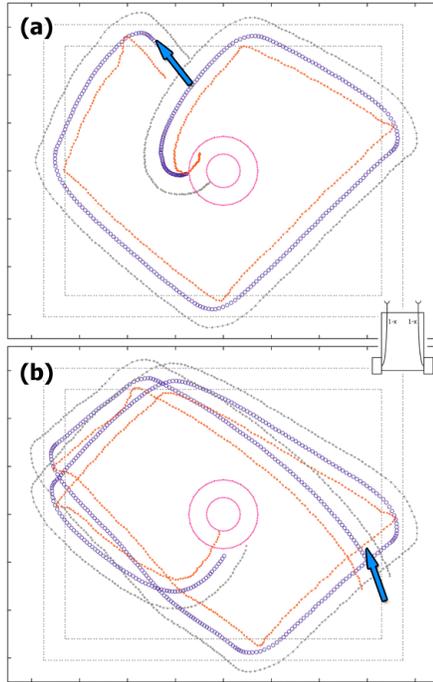


Fig. 3. Examples of the driven path of a Braitenberg-Vehicle with uncrossed (1-x)-connections (*permanent love*).

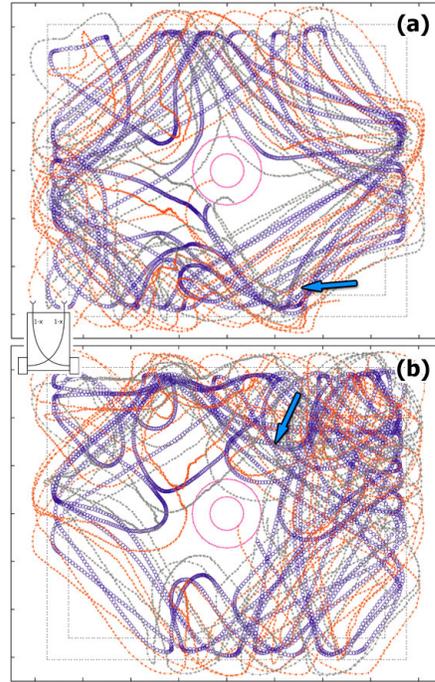


Fig. 4. Examples of the driven path of a Braitenberg-Vehicle with crossed (1-x)-connections (*exploring love*).

4.1 Localisation by Exploration and Hillclimbing

With the source placed in the middle of the inspected area, the average path length of a vehicle with uncrossed (1-x)-connections was 8.49 m compared to 9.67 m for the reference (the median was 6.57 m compared to 8.31 m). Applying the Wilcoxon two sample test³, no significant difference could be determined in this case ($p_{H_0} = 0.2685$)⁴. Thus, a significant improvement of the gas source tracing performance in terms of the average path length cannot be asserted. However, it is important to note that the tested reference strategy, which approximately implements specular reflection at the walls, does not provide uniform coverage over a designated area [6]. Instead, the central area is covered excessively at the expense of the periphery [19]. Hence, the fact that the Braitenberg vehicle does not outperform the reference strategy might be caused by the inherent tendency of the reference strategy to “find” a gas source in the middle of the testbed area.

³ Because the path lengths are not normally distributed, a statistical test was performed that does not assume a specific distribution.

⁴ The Wilcoxon two sample test evaluates the hypothesis H_0 that the populations from which the samples are taken have identical median values.

Source	Strategy	$K_v [\frac{cm}{s}]$	tot. path [m]	av. path [m]	av. dist. [cm]	av. speed/ K_v
Middle	Ref (1-x)	5	319.0	9.67 ± 7.66	136.7 ± 44.9	95.9%
	PL, 1-x	5	1044.0	8.49 ± 7.93	127.4 ± 45.7	73.5%
	EL, 1-x	5/3	731.8	40.66 ± 34.10	143.3 ± 41.0	76.5%
Corner	Ref (1-x)	5	1554.9	20.46 ± 19.38	223.5 ± 78.7	97.6%
	PL, 1-x	5/3	1251.1	11.69 ± 11.22	199.3 ± 88.2	75.5%

Table 1. Statistics of the localisation experiments. The second column references the applied strategy: Ref (reference random search), PL (uncrossed connections, (*permanent love*) or EL (crossed connections, (*exploring love*). In cases where different speed gains K_v were tested both of them are given separated by a slash.

Additional tests were therefore performed with the source placed at a less prominent location near to one corner of the field (15 cm away from the beginning of the repellent wall potential along both the x- and y-axis). For each corner a total of approximately 3 hours of trials was performed both with and without a source. The average path length needed to reach an active gas source was 11.69 m, compared to 20.46 m in the reference experiments (median: 7.25 m compared to 17.09 m). In contrast to the situation where the gas source was placed in the middle of the testbed area, the Wilcoxon two samples test reveals a highly significant improvement in tracing performance in terms of the average path length ($p_{H_0} < 10^{-4}$).

4.2 Localisation by Exploration and Concentration Peak Avoidance

With crossed connections a completely different behaviour results. Although the robot is expected to stay near the source (again assuming a smooth gradient) and thus collisions should be not unlikely, the robot managed to avoid the source most of the time (see Fig. 4). The difference compared to the trials with uncrossed connections is apparent and can be demonstrated with high certainty by a Wilcoxon two sample test ($p_{H_0} < 10^{-7}$).

Though this strategy is obviously not suitable for driving the robot towards a gas source as quickly as possible, it offers an interesting solution to the full gas source localisation problem, including the declaration of the source. This can be seen in Fig. 4. After combining the shown paths, the location of the source is indicated by the part of the picture that remains light. Notice that in contrast to the area that is covered by an obstacle (the inner circle in Figures 3 and 4), the area between the concentric circles should not remain completely light in the case of pure obstacles.

The reason why exploration and concentration peak avoidance might be a solution for gas source localisation can be explained as follows: with crossed inhibitory connections the robot explores the available space and evades each local concentration maximum. Because there exist many of them, it is hard to find a particular maximum that belongs to the actual source by a hillclimbing strategy. On the other hand, concentration maxima occur more frequently near the odour source and thus the source's location remains comparatively light in plots of the driven path such as Fig. 4.

There are several reasons to favour a localisation strategy based on exploration and concentration peak avoidance. First of all it provides a method for declaring a gas source

that does not appear as an obstacle to the robot, a task which can't be accomplished by a hillclimbing strategy because the location of a gas source is usually not featured by a global concentration maximum. It might also be preferable to evade increased gas concentrations in case of an analyte that is corrosive or in other ways offensive to the robot itself.

Note that the experiments point to a feature that can be used to identify a gas source rather than providing a complete solution that is generally applicable. If the source is not detectable as an obstacle, for example, it would not be possible to avoid collisions with the currently used setup. During the experiments the robot moved occasionally into the clearance area around the gas source. A collision with the gas source could corrupt the practicability of the proposed method if the robot spilled the smelling liquid or wetted its tyres, which would cause the robot to become a gas source too.

5 Conclusions and Future Work

This paper is concerned with the problem of tracing and declaring a static gas source by a mobile robot using electrochemical gas sensors. Two Braitenberg-type strategies were investigated in an environment without a controlled airflow and both were shown to be of possible use for gas source localisation.

Using uncrossed inhibitory connections it was found that the average path length the robot needs to move to the source can be decreased. The path length could be reduced by up to a factor of two compared to a strategy that explores the available area ignoring the gas sensor readings. This factor, however, depends on the considered situation and the reference search strategy used.

For real world applications this strategy has to be extended by an additional declaration mechanism to determine that the gas source has been found with high certainty. This mechanism could be added by using other sensors, which 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 those that are closest to the source. Due to the fact that maxima occur more frequently near the odour source, the path of the robot covers the whole available area except that near to the actual location of the source.

Applying this strategy therefore offers a solution to the task of gas source declaration without using additional sensors. A further advantage is that the average distance to the gas source is increased (see Table 1) and that direct contact between the robot and the volatile or liquid substance to be detected can be diminished. This might be preferable in the case of an analyte that is corrosive or in other ways offensive to the robot itself. Furthermore, it helps to prevent the robot from wetting its wheels in cases where the gas source is a liquid substance that cannot be sensed as an obstacle. This is generally not desired because the robot would thereby become a gas source too. With the setup of the Mark III mobile nose, however, it was not possible to avoid collisions with the gas source completely. Future work should therefore address possible modifications to achieve a collision-free path even if the gas source cannot be sensed as an obstacle. A suggestion would be to add a third gas sensor array in a tube that sticks out to the front of the robot in order to examine the area the robot is about to drive towards.

If this third tube protrudes over the robot's front sufficiently, it should be possible to decrease the probability to traverse a puddle of the analyte with the robot, even if the robot approaches the source from a down-wind direction.

A possible objection to the suggested method is time consumption. Because the actual location of a gas source is determined by excluding all other possible locations, the time needed to locate the source increases with the size of the area observed. On the other hand the time consumption scales down with the number of robots utilised. And after all, there is as yet no other known method to localise a gas source that does not appear as an obstacle to the robot in an uncontrolled environment.

References

1. Valentino 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, 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.
4. 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.
5. G.S. Fraenkel and D.L. Gunn. *The Orientation of Animals*. Clarendon Press, Oxford, 1940.
6. Douglas 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. 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.
8. J. O. Hinze. *Turbulence*. McGraw-Hill, New York, 1975.
9. Owen Holland and David McFarland. *Artificial Ethology*. Oxford University Press, New York, 2001.
10. Hiroshi Ishida, Yukihiko Kagawa, Takamichi Nakamoto, and Toyosaka Moriizumi. Odour-Source Localization in the Clean Room by an Autonomous Mobile Sensing System. *Sensors and Actuators B*, 33:115–121, 1996.
11. 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.
12. 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.
13. Jean Claude Latombe. *Robot Motion Planning*. Kluwer Academic, 1991.
14. 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.
15. 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.

16. 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.
17. 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.
18. Achim Lilienthal, Michael R. Wandel, Udo Weimar, and Andreas 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, 2001.
19. Michael J. McNish. Effects of uniform target density on random search. Master's thesis, Naval Postgraduate School, Monterey, California, 1987.
20. Takamichi Nakamoto, Hiroshi Ishida, and Toyosaka Moriizumi. A Sensing System for Odor Plumes. *Analytical Chem. News & Features*, 1:531–537, August 1999.
21. 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.
22. R. Andrew Russell. Ant Trails - an Example for Robots to Follow? In *IEEE Int Conf. Robotics and Automation (ICRA 1999)*, pages 2698–2703, 1999.
23. R. Andrew Russell. *Odour Sensing for Mobile Robots*. World Scientific, 1999.
24. R. Andrew Russell and Anies 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. 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.
26. R. Andrew Russell, David Thiel, and Alan Mackay-Sim. Sensing Odour Trails for Mobile Robot Navigation. In *IEEE Int Conf. Robotics and Automation (ICRA 1994)*, pages 2672–2677, 1994.
27. Titus Sharpe and Barbara 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. Distanto. Goal-oriented Mobile Robot Navigation Using an Odour Sensor. In *Proceedings of the Intelligent Vehicles Symposium '95*, pages 147–151, 1995.
29. 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)*, pages 507–512, Coimbra, Portugal, 2003.