Learning to Locate an Odour Source with a Mobile Robot[*]

Tom Duckett, Mikael Axelsson, and Alessandro Saffiotti

Centre for Applied Autonomous Sensor Systems
Dept. of Technology, Örebro University
SE-70182 Örebro, Sweden
http://www.aass.oru.se

Abstract

We address the problem of enabling a mobile robot to locate a stationary odour source using an electronic nose constructed from gas sensors. On the hardware side, we use a stereo nose architecture consisting of two parallel chambers, each containing an identical set of sensors. On the software side, we use a recurrent artificial neural network to learn the direction to a stationary source from a time series of sensor readings. This contrasts with previous approaches, that rely on the existence of a model of the sensor's dynamics. The complete system is able to orient and turn towards the source. An experimental validation was carried out to evaluate the performance of the system.

1 Introduction

The purpose of this work is to enable a mobile robot to locate a stationary odour source in free air. Possible applications of this technology include detection of gas leaks and of dangerous substances, inspection of pipes in factories, and mine sweeping.

The problem of odour source localisation presents a number of challenges connected to the characteristics of solid state gas sensors and of the environment.

1. *Poor selectivity:* the sensors cannot discriminate similar gases, so it is hard to single out the odour trail of the intended substance.

2. *Lack of an established sensor model:* it is very difficult to model the sensor-environment system analytically, due to the chaotic nature of odour propagation and the noise in the sensor readings; a dynamic model is even more difficult to obtain.

3. *Turbulent nature of air currents:* this causes unpredictable variations and ambiguities in the sensor patterns, with different locations of the source producing the same signature. Hence, localisation cannot be performed using only the current sensor readings — instead, we must use historical information.

4. *Latency:* the sensors have long latency and decay periods (see Fig. 2); therefore, unlike more usual sensors such as range-finders, one cannot assume conditional independence of successive sensor readings during real-time operation.

In this paper, we describe an initial set of laboratory experiments conducted using a Koala mobile robot equipped with an electronic "nose" constructed from multiple solid-state gas sensors. We cope with the above challenges as follows. (1) We use a very distinctive, strong-smelling substance (acetone) for the tests; we also use an environment with partially controlled air flows (see Fig. 1). (2) We use an artificial neural
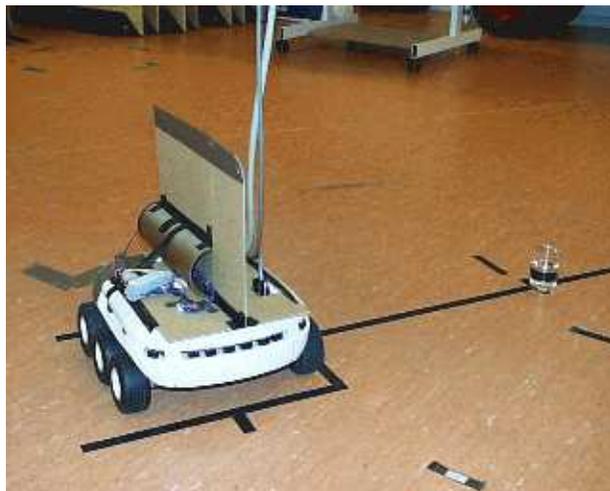

Figure 1: Experimental set-up. The odour sensors are inside the tube mounted on a Koala mobile robot. The odour source is a beaker filled with acetone.
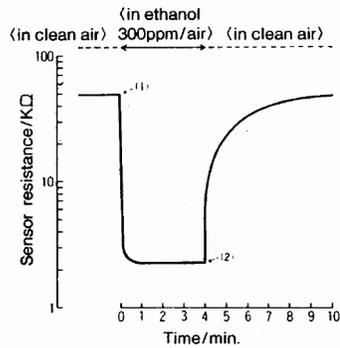
Figure 2: Response of the TGS 822 gas sensor. The sensor takes up to one minute to reach the maximal response to a new odour; on removal of the odour, recovery can take up to 6 minutes.

Figure 3: Top: Figaro gas sensors. Bottom: arrangement of sensors in the tube.

network to acquire the system model, by learning a mapping from the sensor readings to a target output state. (3-4) We use a sensing strategy which involves taking a sequence of sensor readings, and then apply a *recurrent* neural network [4] to combine sensory evidence accumulated over time.

The final system is able to orient and turn towards the source. An experimental validation was carried out to evaluate the performance of this system.

2 Previous Work

Smell is perhaps the least used sense in robotic applications. However, several studies that integrate odour sensors and mobile robots have appeared in the literature. Most of these studies address one of two problems: following an odour trail on the ground (e.g., [15, 14]); or locating a distant odour source. In this paper, we focus on the latter problem.

Early systems for odour source localisation were based on the idea of following the gradient of the gas concentration by looking at the instantaneous response of the sensor [13]. In order to cope with the irregularities in the concentration caused by air turbulence, some systems have used a combination of gas sensors and anemometers [10], or introduced forced ventilation [6]. Several authors took inspiration from biological observations. In return, others have used smelling robots to evaluate hypotheses of animal behaviour: for instance, in [5] a robot is used to study chemo-orientation strategies in lobsters.

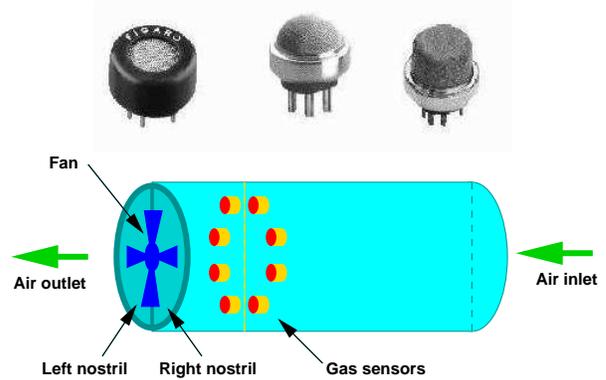Gas sensors respond to changes in gas concentration

with long delays. To compensate for these delays, later systems incorporate a model of the dynamic sensor response. Ishida and colleagues have developed a model-based 'odour compass' which can measure the bearing to an odour source with good speed and accuracy [8]. (See [7] for a three dimensional version.) The parameters of the dynamic model were identified by a recursive least squares method. In another work [9], the authors use a dynamic model of the sensor plus a model of the gas distribution to estimate the *distance* to the odour source.

While the above techniques have produced good results, reliance on a known dynamic model of the sensor response may reduce their applicability. Parameter identification is a difficult and time consuming procedure; moreover, the parameters may change due, e.g., to unpredictable fluctuations in the environmental conditions. The learning approach proposed in this paper may provide an attractive alternative.

3 Hardware Design

The first stage of the project was to develop an appropriate experimental set-up and hardware design. The odour sensor developed is based on a set of Figaro gas sensors (Fig. 3 top). These sensors contain a component whose resistance changes as a function of the concentration of molecules of a given gas.

Fig. 3 shows the schematics of the sensor developed, based on a preliminary design by Berg and Tegebo [2]. Eight gas sensors are included in a tube, and a small fan is used to to pump air through the tube [6]. The sensors are oriented with the wind flow to reduce the cooling effect.

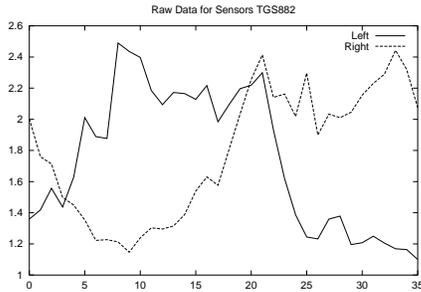The tube is divided into two chambers, or *nostrils*,

Figure 4: Readings for two gas sensors placed in different nostrils during a 360° rotation

and two identical sets of four sensors are placed in the two nostrils. The reason for this construction is as follows. Due to air turbulence and sensor drift effects, the absolute response of the sensors does not provide meaningful information about the odour source location. However, we can find the direction to the source by taking the *relative* strength of two readings from sensors placed in the two parallel chambers. Fig. 4 illustrates this: the figure shows the readings from two TGS 882 sensors placed in the two nostrils as the robot turned anti-clockwise through 360°, starting with the nose facing in the opposite direction to the source, and sampling at 10° intervals.

Each nostril contains four sensors of a different type, thus endowing the sensor with the ability to discriminate between different substances [11]. We plan to use this ability in future experiments of multi-substance localisation.

The odour sensor is mounted on top of a Koala mobile robot. Both the sensor and the robot are connected by a cable to a PC running the data acquisition, learning, and navigation software.

4   Software Design

4.1   Pre-Processing

To remove some of the outliers and undesired noise from the recorded sensor data, we first applied a median filter to smooth the sequence of sensor readings. Each sensor was treated independently, sorting the window of five successive values around the current reading and then taking the middle value.

The second pre-processing step applied was a simple linear transformation on the sequence of values for each sensor, so that all of the inputs to the neural network had a similar range of values. Again, each sensor

was treated independently, the sequence of readings being transformed to have zero mean and unit standard deviation (see Bishop [3, p. 298]). For each input variable x_i, the set of re-scaled values is given by

$$\tilde{x}_i^m = \frac{x_i^m - \overline{x}_i}{\sigma_i},$$

where $\overline{x}_i$ refers to the mean value for x_i, σ_i is the standard deviation for x_i, and $m = 1, ..., M$ labels the sequence of sensor patterns.

Finally, each individual sensor pattern $\mathbf{x}^m$ was normalized to lie on the unit hyper-sphere, by dividing each sensor value $\tilde{x}_i^m$ by the L2-norm $\|\mathbf{x}^m\|$ as

$$\|\mathbf{x}^m\| = \sqrt{\sum_{i=1}^{N} |\tilde{x}_i^m|^2},$$

where $i = 1, ..., N$ labels the different sensors ($N = 8$). The purpose of this step was to remove the absolute magnitude information from the vector of sensor readings. This was found to be necessary because the magnitude of the readings from both chambers of the nose varies unpredictably with the natural variation in the propagation of odours in the environment. By taking the L2-norm, we constrain the input vector of the neural network in order to focus the learning on the *relative* magnitude of left and right sensor readings.

4.2   Recurrent Neural Network

In early experiments, we found that a non-recurrent neural network was unable to learn the required mapping from the current sensor readings to the direction of the odour source, due to the noisy and ambiguous nature of the sensor patterns (see Results). We therefore applied a recurrent neural network, which is able to take into account a temporal sequence of sensor patterns, and thus to disambiguate similar looking situations. The structure of this network is shown in Fig. 5.

After some initial experiments, the most appropriate network architecture using one hidden layer was found to have 10 hidden units, and therefore 10 recurrent input units in addition to the 8 inputs from the sensors. The recurrent inputs at time t were taken from the outputs of the hidden units at time $t - 1$, with the initial recurrent inputs being set to zero.

The target outputs for training were obtained from the angular displacement of the nose relative to the
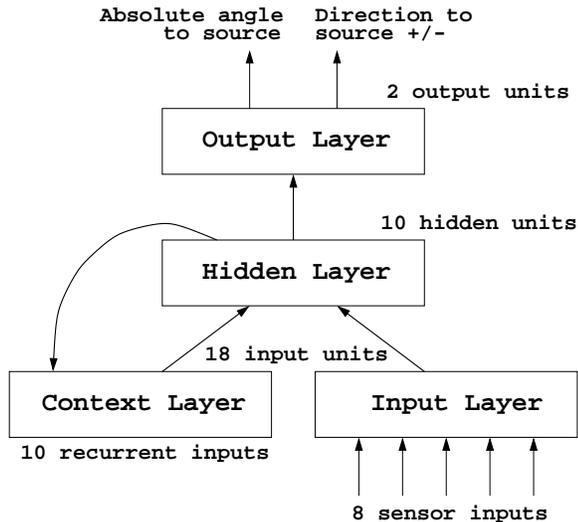
Figure 5: Schematic diagram of the recurrent neural network. A set of 10 extra input units, corresponding to recurrent connections from the 10 units of the hidden layer, were added to a fully connected multi-layer perceptron. The recurrent inputs at time t were taken from the outputs of the hidden units at time $t - 1$.

direction of the odour source. In order to deal with the discontinuity in the angle (between 359 and 0 degrees), two output units were used with the following coding:

- *Absolute angle to the source*, from 0 to 180 degrees, normalized to lie between 0 and 1.

- *Direction to the source.* Positive angles were encoded as 1, negative angles as 0, and the intermediate angles 0 and 180 degrees as 0.5.

All of the units were implemented using a sigmoidal activation function, and the network was trained using the standard algorithm for error back-propagation with the sum-of-squares error function [3] and a learning rate of 0.2.

5 Results

For training and testing, the experimental set-up shown in Fig. 1 was used, where the odour source consisted of a beaker containing 80 ml of 99.5% acetone placed at a distance of 50 cm from the robot. Sixty complete "rounds" of sensor data were collected, where a round consists of one complete rotation by the robot at a fixed position. In each round, data collection started with the nose facing in the direction away from the source, turning anti-clockwise to collect 36 sets of sensor readings sampled at 10° intervals and pausing 5 seconds at each sampling angle. The collected data was then treated using the pre-processing steps described in section 4.1. Fig. 6 shows the data collected in ten rounds, and the effect of the median filter: notice that, although some of the outliers are removed by the filtering, the data remain quite noisy and ambiguous. (Performing the experiment in a clean room would probably result in less noisy data, but would be less indicative of the applicability of the technique to real situations.)

The recurrent neural network was trained using data from 50 rounds, and tested on the other 10 rounds. The performance of the recurrent network on the test data was then compared to that of a standard multi-layer perceptron with one hidden layer of 10 units (see Fig. 7). The mean squared error for the recurrent network on the test data was 0.046 compared to 0.101 for the non-recurrent network.

A validation experiment was then conducted on the mobile robot. Here, the robot was programmed to collect two rounds of data under the same conditions as in training and testing, to estimate the direction of the source using the collected data, and then to turn towards the estimated direction. This experiment was repeated over 7 runs, as indicated in Table 1.

To test the ability of the network to generalise to situations on which it had not been trained, we conducted a further three experiments with different distances and starting directions to the source. The results show that the starting direction did not affect the performance of the system (experiment 2). In addition, the performance of the robot improved as the source was moved nearer to the robot (experiment 3), even though the neural network had not been trained with data collected at this distance. However, the robot

Expt.	Dist. to source /cm	Starting direction /deg	Runs	Mean Err. /deg
1	50	180	7	22.9 ± 13.8
2	50	0	7	14.3 ± 14.0
3	25	180	7	18.6 ± 10.7
4	75	180	7	52.9 ± 58.3

Table 1: Summary of results for the validation experiments with the complete system, showing the mean absolute angular error of the robot in finding the odour source and standard deviation.
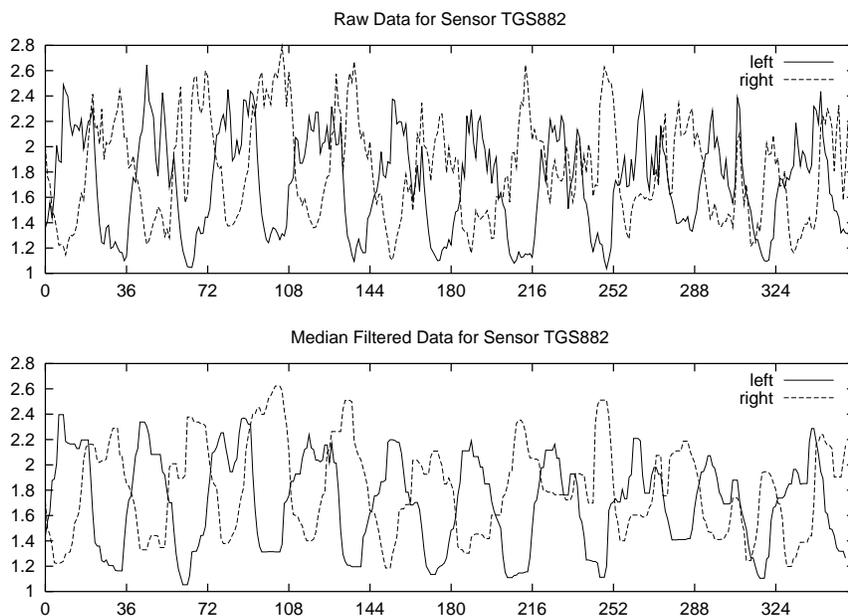
Figure 6: Top: Raw data from left and right nostril. Bottom: Same data, median filtered.

was not as successful when the source was moved further away from the robot (experiment 4), partly due to one completely failed run when the robot ended up facing 180 degrees away from the source. (Without this run, the mean error in experiment 4 would have been 31.7 ± 17.2.)

To evaluate the significance of the results, a t-test [12] was performed for each experiment. Here, we tested the null hypothesis that these results were generated by chance, comparing the actual results to an assumed population mean of 90 degrees (the average angular error we would expect if the robot turned in random directions). The null hypothesis was rejected at the 99% level of confidence for experiments 1 to 3, but was accepted for experiment 4.

6   Conclusions and Future Work

Using gas sensors on a mobile robot to locate an odour source is a demanding task, that requires the ability to deal with noisy data and historical information. The experiments shown in this paper demonstrate that a recurrent neural network is an adequate tool to solve this task, without having to rely on the existence of a model of the sensors' dynamics.

There are two most notable limitations in the cur-

rent experiments. First, the training data were only collected from a fixed distance, resulting in relatively poor performance from different distances — neural networks trained off-line are very good at interpolating, but not so good at extrapolating. Second, the navigation part of the experiment is trivial, with the robot simply pointing to the estimated direction of the odour source.

Future work will focus on navigation strategies to reach the odour source. It should be noted that simple gradient following may be ineffective, since the odour is typically broken into unpredictable patches due to air turbulence [1]. Therefore, more complex strategies will need to be developed.

References

[1] J. Atema. Eddy chemotaxis and odor landscapes: exploration of nature with animal sensors. *Biol. Bull.*, 191(1):129–138, 1996.

[2] S. Berg and A. Tegebo. Development of an electronic nose. Technical Report Oru.Te-1EL024-A103/99, Örebro University, 1999.
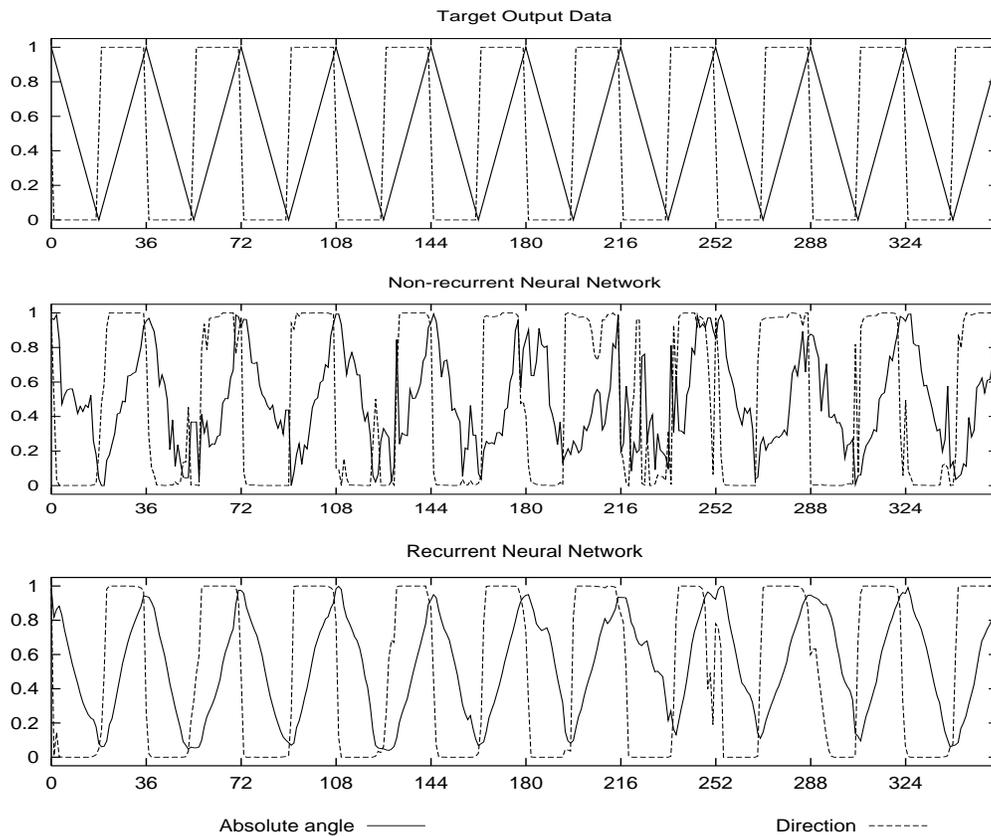
Figure 7: Top: target output for the test data. Middle: actual output for the non-recurrent neural network. Bottom: actual output for the recurrent neural network.

[3] C.M. Bishop. *Neural Networks for Pattern Recognition*. Oxford University Press, 1995.

[4] J.L. Elman. Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning*, 7:195–225, 1991.

[5] F.W. Grasso, J.A. Basil, and J Atema. Toward the convergence: robot and lobster perspectives of tracking odors to their source in the turbulent marine environment. In *Proc. of the IEEE Intl. Symp. on Intelligent Control (ISIC)*, pages 259–264, 1998.

[6] H. Ishida, K. Hayashi, M. Takakusaki, and T. Nakamoto. Odour-source localization system mimicking behaviour of silkworm moth. *Sensors and Actuators A*, 51:225–230, 1996.

[7] H. Ishida, A. Kobayashi, T. Nakamoto, and T. Moriizumi. Three-dimensional odor compass. *IEEE Trans. on Robotics and Automation*, 15(2):251–257, 1999.

[8] H. Ishida, T. Nakamoto, and T. Moriizumi. Study of odor compass. In *Proc. of the IEEE Intl. Conf. on Multisensor Fusion and Integration for Intelligent Systems*, pages 222–226, 1996.

[9] H. Ishida, T. Nakamoto, and T. Moriizumi. Remote sensing and localization of gas/odor source and distribution using mobile sensing system. In *Proc. of the Intl. Conf. on Solid State Sensors and Actuators*, pages 559–562, Chicago, 1997.

[10] H. Ishida, K. Suetsugu, T. Nakamoto, and T. Moriizumi. Study of autonomous mobile sensing system for the localization of odor source using gas sensors and anemometric sensors. *Sensors and Actuators A*, 45:153–157, 1994.

[11] A. Loutfi, S. Coradeschi, T. Duckett, and P. Wide. Odor source identification by grounding linguistic descriptions in an artificial nose. In *SPIE Intl. AeroSense Symp.*, 2000.

[12] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical Recipes in C, 2nd. edition*. Cambridge University Press, 1992.

[13] R. Rozas, J. Morales, and D. Vega. Artificial smell detection for robotic navigation. In *Proc. of the 5th Intl. Conf. on Advanced Robotics*, pages 1730–1733, 1991.

[14] R.A. Russell. Ant trails – an example for robots to follow? In *Proc. of the IEEE Intl. Conf. on Robotics and Automation*, pages 2698–2703, 1999.

[15] E. Stella, F. Musio, L. Vasanelli, and A. Distante. Goal-oriented mobile robot navigation using an odour sensor. In *Proc. of the IEEE Intl. Symp. on Intelligent Vehicles*, pages 147–151, 1995.