Intelligent Systems for Machine Olfaction:
Tools and Methodologies

Evor L. Hines, University of Warwick, UK

Mark S. Leeson, University of Warwick, UK

Medical Information Science
REFERENCE

Senior Editorial Director: Kristin Klinger
Director of Book Publications: Julia Mosemann
Editorial Director: Lindsay Johnston
Acquisitions Editor: Erika Carter
Development Editor: Joel Gamon
Production Coordinator: Jamie Snavely
Typesetters: Keith Glazewski & Natalie Pronio
Cover Design: Nick Newcomer

Published in the United States of America by
Medical Information Science Reference (an imprint of IGI Global)
701 E. Chocolate Avenue
Hershey PA 17033
Tel: 717-533-8845
Fax: 717-533-8661
E-mail: cust@igi-global.com
Web site: http://www.igi-global.com

Library of Congress Cataloging-in-Publication Data

Intelligent systems for machine olfaction : tools and methodologies / Evor L. Hines and Mark S. Leeson, editors.
p. cm.
Includes bibliographical references and index.
ISBN 978-1-61520-915-6 (hardcover) -- ISBN 978-1-61520-916-3 (ebook) 1. Chemical detectors--Automatic control. 2. Gas detectors--Automatic control. 3. Odors. 4. Olfactometry. 5. Artificial intelligence. I. Hines, Evor, 1957- II. Leeson, Mark S., 1963- III. Title.

TP159.C46I48 2011
681'.754--dc22

2011004016

British Cataloguing in Publication Data
A Cataloguing in Publication record for this book is available from the British Library.

Chapter 10
Improved Gas Source Localization with a Mobile Robot by Learning Analytical Gas Dispersal Models from Statistical Gas Distribution Maps Using Evolutionary Algorithms

Achim J. Lilienthal
Örebro University, Sweden

ABSTRACT

The method presented in this chapter computes an estimate of the location of a single gas source from a set of localized gas sensor measurements. The estimation process consists of three steps. First, a statistical model of the time-averaged gas distribution is estimated in the form of a two-dimensional grid map. In order to compute the gas distribution grid map the Kernel DM algorithm is applied, which carries out spatial integration by convolving localized sensor readings and modeling the information content of the point measurements with a Gaussian kernel. The statistical gas distribution grid map averages out the transitory effects of turbulence and converges to a representation of the time-averaged spatial distribution of a target gas. The second step is to learn the parameters of an analytical model of average gas distribution. Learning is achieved by nonlinear least squares fitting of the

DOI: 10.4018/978-1-61520-915-6.ch010

analytical model to the statistical gas distribution map using Evolution Strategies (ES), which are a special type of Evolutionary Algorithm (EA). This step provides an analysis of the statistical gas distribution map regarding the airflow conditions and an alternative estimate of the gas source location, i.e. the location predicted by the analytical model in addition to the location of the maximum in the statistical gas distribution map. In the third step, an improved estimate of the gas source position can then be derived by considering the maximum in the statistical gas distribution map, the best fit, as well as the corresponding fitness value. Different methods to select the most truthful estimate are introduced, and a comparison regarding their accuracy is presented, based on a total of 34 hours of gas distribution mapping experiments with a mobile robot. This chapter is an extended version of the conference paper (Lilienthal et al., 2005).

INTRODUCTION

A major problem for gas source localization in a natural environment is the strong influence of turbulence on the dispersal of gas. Typically, turbulent transport is considerably faster compared to molecular diffusion (Nakamoto et al., 1999; Roberts and Webster, 2002). Apart from very small distances where turbulence is not effective, molecular diffusion can thus be neglected concerning the spread of gas. A second important transport mechanism for gases is advective transport due to prevailing fluid flow. Relatively constant air currents are typically found even in an indoor environment without ventilation (Wandel et al., 2003) as a result of pressure (draught) and temperature inhomogeneities (convection flow).

Turbulent flow comprises at any instant a high degree of vortical motion, which creates packets of gas that follow chaotic trajectories (Shraiman and Siggia, 2000). This results in a concentration field, which consists of fluctuating, intermittent patches of high concentration. The instantaneous concentration field does not exhibit smooth concentration gradients that indicate the direction toward the centre of a gas source (Lilienthal and Duckett, 2004b; Russell, 1999). Figure 1 illustrates actual gas concentration measurements recorded with a mobile robot along a corridor containing a single gas source. It is important to note that the noise is dominated by the large fluctuations of the instantaneous gas distribution and not by the electronic noise of the gas sensors. Turbulence is chaotic in the sense that the instantaneous flow velocity at some instant of time is insufficient to predict the velocity a short time later. Consequently, a snapshot of the distribution of a target gas at a given instant contains little information about the distribution at another time. However, under certain assumptions (e.g. that the air flow is uniform and steady) the *time-averaged* concentration field varies smoothly in space with moderate concentration gradients (Roberts and Webster, 2002).

Figure 1. Normalized raw response readings from an example trial

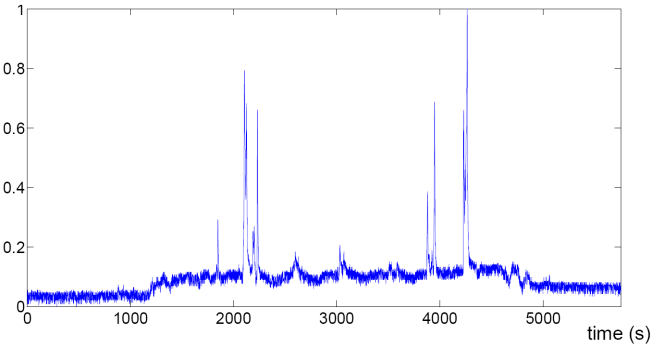

It is often desirable to know the spatial structure of the time-averaged gas distribution. The Kernel Distribution Mapping (Kernel DM) algorithm was introduced by Lilienthal and Duckett (Lilienthal and Duckett, 2003b) to compute a grid map representation of the structure of the time-averaged gas distribution. The input to the Kernel DM algorithm is a set of localized gas sensor readings. In this chapter we consider the case that the readings were collected by a mobile robot. The algorithm is summarized in Section ”The Kernel DM Algorithm”. In itself, gas distribution mapping is useful for any application that requires estimating the average distribution of a certain gas in a particular area of the environment. For example, mobile robots that are able to build such a map can be used for pollution monitoring (DustBot Consortium, 2006), they could indicate contaminated areas in a rescue mission, or could be used in Precision Farming (Blackmore & Griepentrog, 2002) to provide a non-intrusive way of assessing certain soil parameters or the status of plant growth to enable a more efficient usage of fertilizer.

In this chapter, which is an extended version of a paper by the authors (Lilienthal et al., 2005), we describe a method to use statistical gas distribution grid maps in order to locate a gas source. An obvious clue for the gas source position is the maximum in the map. Experiments in an indoor environment indeed demonstrated that the concentration maximum estimate (CME)[1] provides a satisfying approximation of the source location in many cases (Lilienthal & Duckett, 2004a). Under certain assumptions discussed in Sec. “Analytic Gas Distribution Model”, the spread of a gas that evaporates from a stationary source can be approximated as a Fickian diffusion process. Instead of the small diffusion constant that describes molecular diffusion the turbulent diffusion is ruled by a substantially larger turbulent diffusion constant K (eddy diffusivity). In the event of negligible advective transport due to a weak air current, the resulting average gas distribution takes a circular shape. In such cases, it was observed that the distance between the CME and the true source

location was small. By contrast, the localization capability of the CME was found to be considerably degraded when the concentration map showed a stretched out distribution due to a dominant wind direction. According to the equations, which describe the time-averaged stationary gas distribution analytically (see Sec. "Analytic Gas Distribution Model"), the concentration decreases slowly *along* the direction of a constant air current. Thus, even small distortions due to rudiments of turbulent concentration peaks can cause a large displacement of the point of maximum concentration. Also the localization error introduced by the fact that gas sensor measurements are acquired not exactly level with the gas source is more pronounced in the case of stronger air current.

A method that formalizes this qualitative argument is presented in Sec. "Step 3 – Selection of Source Location Estimate". The method allows distinguishing situations, where the CME is a reliable approximation of the source location from situations where the CME is unlikely to indicate the gas source position accurately. This is accomplished by comparing how well the statistical gas distribution map can be approximated by the analytical model (detailed in Sec. "Analytic Gas Distribution Model"), which describes the time-averaged gas distribution under certain idealized assumptions.

Apart from providing a measure of the reliability of the CME, the introduced method allows to derive an alternative estimate of the gas source position, which can be used in situations where the CME fails. To determine the analytical model, which approximates the given statistical gas distribution map most closely, the parameter set is optimized by means of nonlinear least squares fitting. Since the model parameters include the position of the gas source, the best fit naturally corresponds to an estimate of the source position. In contrast to the CME, the best fit estimate (BFE) is derived from the whole distribution represented in the statistical gas distribution grid map.

The rest of this chapter is structured as follows. First, the Kernel DM algorithm to compute statistical gas distribution grid maps is described in Sec. "Step 1 - Computation of a Statistical Gas Distribution Model". Second, the adapted physical model that was used to approximate the time-averaged gas concentration is introduced in Sec. "Analytic Gas Distribution Model". Then, the evolutionary strategy method to learn the optimal model parameters is detailed in Sec. "Step 2 – Learning Parameters of the Analytical Model". Next, reasons that cause deviations between the physical model and the statistical gas distribution grid map are discussed in Sec. "Sources of Inaccuracy" and two strategies to select the best estimate of the gas source location are presented in Sec. "Step 3 – Selection of Source Location Estimate". Finally details of the experimental setup are given in Sec. "Experimental Setup" and results are presented in Sec. "Results", followed by conclusions and suggestions for future work in the final Section "Conclusions and Outlook".

STEP 1: COMPUTATION OF A STATISTICAL GAS DISTRIBUTION MODEL

Creating Gas Distribution Grid Maps

By contrast to metric grid maps extracted from sonar or laser range scans, a single measurement from a gas sensor represents the measured quantity (in the case of metal oxide sensors: the rate of redox reactions) only at the comparatively small area of the sensor's surface, typically around 1 cm^2. Nevertheless, the gas sensor readings contain information about the time-averaged gas distribution of a larger area. First, this is due to the smoothness of the time-averaged gas distribution, which allows extrapolating on the averaged gas sensor measurements because the average concentration field does not change drastically in the vicinity of the point of measurement. Second, the metal-oxide gas sensors perform temporal integration of successive readings implicitly due to their slow response and long recovery time. Modeled as a first-order sensor, the time constants of rise and decay for the complete gas sensitive system used here were estimated as $\tau_r \approx 1.8s$ and $\tau_d \approx 11.1s$, respectively (Lilienthal & Duckett, 2003a). Thus the measurements contain information that is spatially integrated along the path driven by the robot.

The Kernel DM Algorithm

Based on the observations mentioned in Sec. "Creating Gas Distribution Grid Maps", the Kernel DM algorithm introduced in (Lilienthal & Duckett, 2004a) uses a Gaussian kernel function to model the decreasing likelihood that a particular reading represents the true quantity (here: the time-averaged relative concentration) with respect to the distance from the point of measurement. For each measurement, two quantities are calculated for grid cells k: an importance weight and a weighted reading. In practice, only those cells in the vicinity of the point of measurement need to be considered, i.e. the cells for which the corresponding centre $\boldsymbol{x}^{(k)}$ lies within a certain radius around the point $\boldsymbol{x}_t$ where the measurement was taken at time t. Cells that are further away from the measurement can be ignored since the effect of the update is negligible. The importance weight is calculated by evaluating the two-dimensional, uni-variate Gaussian function

$$N(\Delta \mathbf{x}_t^{(k)}) = \frac{1}{2\pi\sigma^2} e^{-\frac{(\Delta \mathbf{x}_t^{(k)})^2}{2\sigma^2}} \quad (1)$$

at the displacement $\Delta \boldsymbol{x}_t^{(k)} = \boldsymbol{x}^{(k)} - \boldsymbol{x}_t$ between the centre of grid cell k and the point of measurement $\boldsymbol{x}_t$. This weight models the information content of a particular measurement r_t at the location $\boldsymbol{x}^{(k)}$ of grid cell k. From these weights, two temporary grid maps are computed: $\Omega^{(k)}_t$ by integrating importance weights and $R^{(k)}_t$ by integrating weighted readings up to time t as

$$\begin{aligned} \Omega_t^{(k)} &= \sum_{i=1}^{t} N(|\, \mathbf{x}_i - \mathbf{x}^{(k)} \,|, \sigma), \\ R_t^{(k)} &= \sum_{i=1}^{t} N(|\, \mathbf{x}_i - \mathbf{x}^{(k)} \,|, \sigma) \cdot r_i \end{aligned} \tag{2}$$

The kernel width σ is a parameter of the algorithm. Please note that the gas sensor readings r_i we consider here are first normalized by linear scaling of each sensor to the range of [0, 1]. Please note further that we assume perfect knowledge about the position x_i of a sensor at the time of the measurement. To account for the uncertainty about the sensor position in connection with any gas distribution modeling algorithm, the method in (Lilienthal et al., 2007) can be used.

The integrated weights $\Omega^{(k)}_t$ provide a confidence measure for the estimate at cell k. A high value means that the estimate is based on a large number of readings recorded close to the centre of the respective grid cell. A low value, on the other hand, means that few readings nearby the cell centre are available and that therefore a prediction has to be made using sensor readings taken at a rather large distance. Consequently, if the sum of the weights $\Omega^{(k)}_t$ exceeds a certain threshold value Ω_{min}, the grid cell is set to

$$r_t^{(k)} = R_t^{(k)} / \Omega_t^{(k)} \quad : \quad \Omega_t^{(k)} \geq \Omega_{min} \tag{3}$$

representing an estimate of the value of the distribution in that particular area of the environment. If the sum of weights is below Ω_{min}, the cell is considered unexplored. Since integrated weights $\Omega^{(k)}_t$ are used for normalization of the weighted readings $R^{(k)}_t$ even coverage is not necessary.

$\Omega^{(k)}_t$ models the information content of a series of gas sensor measurements in a way that reflects the sensor characteristics and the trajectory of the robot. An example is shown in Figure 2. First, the certainty about the average gas distribution is modeled as being higher if the gas sensor is moved at a slower speed. A higher information content is assumed by the Kernel DM algorithm especially in cases where a number of successive measurements were performed on the spot. In this case, the estimated value (calculated by averaging over multiple readings) represents

Figure 2. Integrated importance weight Ω_t for an example sensor trajectory. The contributions from each measurement position are shown (in black) together with the integrated value (in red). The trajectory results from a constant velocity movement along a straight path and an immediate stop after the fifth time step (i.e., measurements x_5, x_6 and x_7 were all taken at the same physical location). The Kernel DM algorithm considers a finite discretization of Ω_t to a grid map, which is not shown in this figure for the sake of a better illustration. The course of Ω_t shown here corresponds to a grid map discretization with infinite resolution.

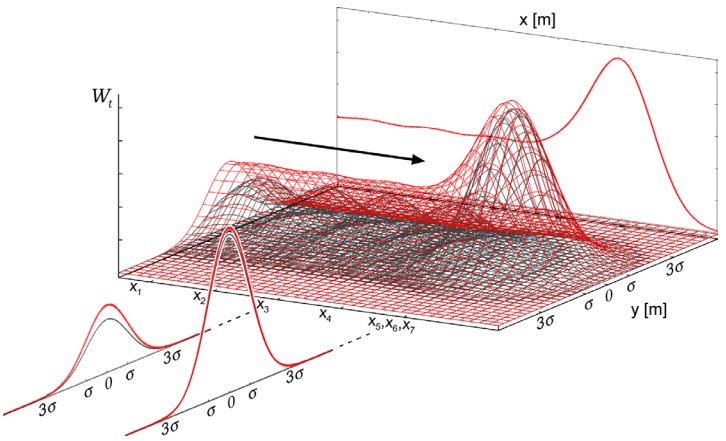

a temporally integrated quantity that naturally contains more information about the gas distribution at this particular location. This value also contains more information about adjacent places because of the higher certainty about the average gas distribution, and because the temporal mean also carries out spatial integration due to the spatial fluctuation of the gas.

Second, the information content is modeled to be particularly high *along* the sensor trajectory, i.e. at places to which the sensors have actually been exposed to. This is particularly important for gas sensors used in this study where, due to the memory effect of the metal oxide sensors, the sensor readings present a low-pass filtered response value integrated along the path driven. Correspondingly the continuously collected measurements contain information about the locations *between* the places where the sensor data were actually logged. By contrast, the importance model decreases quickly orthogonal to the sensor trajectory, corresponding to the fact that the distribution value can only be approximated by extrapolating on the actual measurements assuming smooth transitions in the time-constant gas distribution structure.

Finally, the certainty about the average concentration is modeled as being approximately constant along the path if the robot was driven at a constant, not too

high speed. As long as the time constant of decay is much longer than the time between individual measurements, the information content of the sensor readings about locations along the integration path is in fact approximately independent of the actual points of measurements.

Parameters of the Kernel DM Algorithm

While the actual value of the threshold Ω_{min} was found to have a minor influence on the resulting gas distribution map (Lilienthal & Duckett, 2004a), the width σ of the Gaussian kernel function is a critical parameter. Referring to the exploration path of the robot, σ has to be chosen high enough to satisfy the requirement for sufficient extrapolation on the gas concentration measurements, but low enough to preserve the fine details of the mapped structures. In this work, fixed parameter values of σ = 15 cm, and Ω_{min} = 10.0 × (number of sensors) were chosen based on the considerations in (Lilienthal & Duckett, 2004a).

Impact of the Sensor Dynamics

Due to the response characteristics of metal oxide sensors, a single gas sensor reading represents a temporally and, if the robot is driven at non-zero speed, also a spatially integrated concentration value. The averaging effect is considered implicitly by the model of the information content applied in the Kernel DM algorithm. The function in Equation (1) to compute importance weights contains, on the other hand, no term to model the asymmetry, which is induced by the much longer recovery time compared to the response time of the sensors. It would indeed not be possible to unambiguously determine the actual concentration distribution along the path that caused a given series of gas sensor readings. Therefore, an asymmetric term is not incorporated. Consequently, a certain level of distortion in the mapped gas distribution has to be tolerated. The magnitude of this distortion is estimated below.

As a consequence of the delayed response and the prolonged decay time of the gas sensors, the mapped values show asymmetrically blurred edges and a slightly shifted centre of the area of maximum concentration compared to the real distribution. This effect can be seen in Figure 3, which shows how a rectangular step pulse would be mapped by the Kernel DM algorithm introduced above. In the upper left part (a) the real distribution can be seen, which is a step pulse with an assumed duration of Δt = 10s. In addition, the response of the gas-sensitive system is shown in part (b). This curve was calculated using a first order sensor model with the parameters τ_r and τ_d of the Örebro Mark III mobile nose. In Figure 3 (c), the Gaussian kernel functions multiplied by the corresponding sensor readings are also shown. The samples were assumed to be recorded at a rate of 2 Hz and a width of σ = 1s

Figure 3. Mapping of a rectangular step pulse. The figure shows the step-like concentration course the gas-sensitive system is exposed to (a), the sensor response as calculated for the ¨Orebro Mark III mobile nose (b), the Gaussian weighting functions multiplied by the corresponding sensor readings (c), and the resulting curve of the mapped values (d).

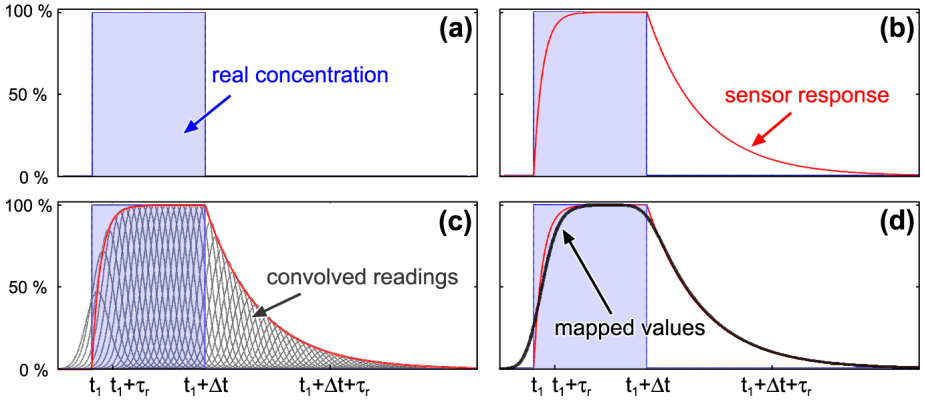

was used for the Gaussian kernel functions in Equation (1). This corresponds to a distance of $\sigma = 5$ cm, if a situation is considered where a robot drives with a constant velocity of 5 cm/s through a 50 cm wide area of constant concentration. Note that the Gaussians vanish in the front part of the graph due to the zero response of the sensor. Finally, the normalized curve of the mapped values is depicted in Figure 3 (d). This curve is calculated according to Equation (3), meaning that the sum of the Gaussians shown in Figure 3 (c) is divided by the sum of the weights.

Comparing the assumed distribution with the course of the mapped values, the asymmetrical shift as well as the blurring effect can be seen in Figure 3 (d). This distortion is, however, not critical. Due to the low speed of the robot, which never exceeded 5 cm/s during the experiments presented here, the expected shift would be in the order of 10 cm at most. This effect is even smaller for smoother distributions, which the metal-oxide sensors can follow more closely than a step-like one. Further on, the directional component of both effects gets averaged out if the robot passed the same point from different directions. For the experiments considered here, this condition is fulfilled because either a predefined exploration path was used, where the path is passed equally often from opposite directions, or the robot was controlled as a gas-sensitive Braitenberg vehicle (Braitenberg, 1984), and particular points were passed equally often from multiple directions on average. Thus, the mapping process results in a representation, which is broadened but not severely shifted compared to the true distribution.

ANALYTIC GAS DISTRIBUTION MODEL

It is currently not feasible to model all aspects of turbulent wind and gas distribution in a realistic environment. A general problem is that many boundary conditions are unknown. Even if sufficiently accurate knowledge about the state of the environment would be available, it would be very time-consuming to achieve the required resolution with a conventional finite element model (Kowadlo & Russell, 2003). For specific situations, however, the time-averaged gas distribution can be described in a computationally inexpensive way. Assuming isotropic and homogenous turbulence and a one-directional, constant wind field, the time-averaged gas distribution of a constantly emitting point source on the floor can be described as (Hinze, 1975; Ishida et al., 1998):

$$C(x,y) = \frac{q}{2\pi K}\frac{1}{r}e^{-\frac{V}{2K}(r-\mathbf{x}_\mathbf{w})}, \tag{4}$$

$$r = \sqrt{(x_S - x)^2 + (y_S - y)^2}\ , \tag{5}$$

$$\mathbf{x}_\mathbf{w} = \frac{\Delta\mathbf{x}\cdot\mathbf{w}}{|\ \mathbf{w}\ |} = (x_S - x)cos\theta + (y_S - y)sin\theta\ . \tag{6}$$

The concentration C at a point (x, y) level with the gas source is determined by the turbulent diffusion coefficient K, the location of the gas source (x_S, y_S), its release rate q, the wind speed V, and the upwind direction θ. Equation (4) comprises a term for symmetric $1/r$ decay and a second term that models asymmetric decay with respect to the wind direction. The variable $\mathbf{x}_\mathbf{w}$ is the projection of the displacement with respect to the source to the upwind direction. Accordingly, the exponential term in Equation (4) is constant along the upwind direction and the asymmetric decay is steepest in downwind direction.

The model introduced in Equation (4) describes a system in the state of equilibrium where the gas source evaporates infinitely long into an infinite space. As a consequence, the model diverges at the source location, which is obviously an unrealistic description of the observed gas concentration. For this reason, the symmetric term is replaced here by the spatial profile of the Green's function of the diffusion equation. For a fixed time (> 0) this function declines with $exp(-const \times r^2)$ and does not diverge at the source location. Consequently, the time-averaged gas distribution was modeled in this work as

$$\tilde{C}(x,y) = C_{00} e^{-C_S r^2} e^{-C_A (r - [(x_S - x)\cos\theta + (y_S - y)\sin\theta])} + C_B \tag{7}$$

This model depends on seven parameters. C_{00} is the maximum concentration at the source location while C_S and C_A describe the magnitude of the symmetric and asymmetric decay, respectively. (x_S, y_S) are the coordinates of the point source and θ is the upwind direction, i.e. the angle between the upwind vector **w** and the x-axis. The parameter C_B is added in order to account for a non-zero base-level.

A comparison of the models specified in Equation (4) and Equation (7) is given in Figure 4. The model of the time-averaged gas distribution as defined in Equation (3.1) is shown in Figure 4 (a). In order to relate the models to each other, the parameters of the adapted model were chosen such that the symmetric terms are equal at a distance of 15cm in case of Figure 4 (b,c), and at a distance of 10cm in case of Figure 4 (d), respectively. The same asymmetric term was used for all four plots. Finally, the parameter C_S was chosen to represent a highly asymmetric distribution with $C_A/C_S = 5$ in Figure 4 (b) and more symmetric distributions with $C_A/C_S = 0.5$ in Figure 4 (c) and Figure 4 (d).

Figure 4. Comparison of models of the time-averaged gas distribution. A plot of the model given by Equation(3.1) (a) is compared with plots of the model given by Equation (3.4) for three different parameter sets: CA/CS = 5, symmetric term equal to the model in Equation (3.1) at a distance of 15 cm from the point source (b); CA/CS = 0.5, symmetric term equal to the model in Equation (3.1) at a distance of 15 cm (c); CA/CS = 0.5, symmetric term equal to the model in Equation (3.1) at a distance of 10 cm (d). The same asymmetric term was used for all four plots.

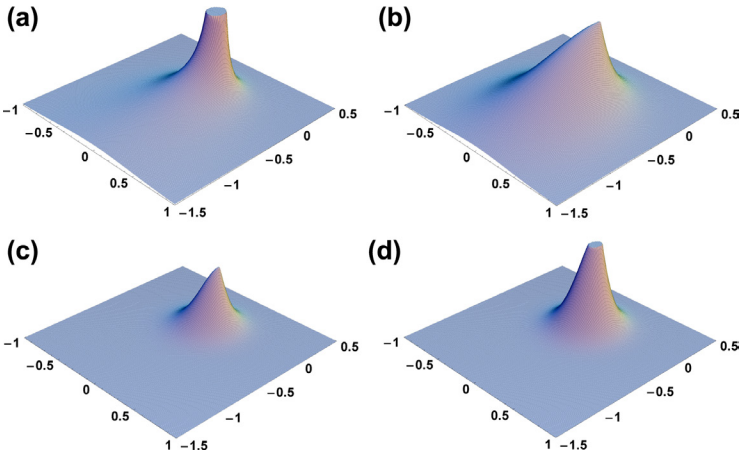

STEP 2: LEARNING PARAMETERS OF THE ANALYTICAL MODEL

The problem is now to find the set of parameters of the analytical model given in Equation (7), which approximates the statistical gas distribution grid map most closely.

In order to calculate the model quality for a given parameter set $\boldsymbol{p}$, the gas distribution predicted by the analytical model C_p is discretized to the grid map and the deviation of $C_p^{(k)}$ from the value $r^{(k)}$ predicted by the statistical gas distribution grid map is determined for all corresponding grid cells $k \in \{1,..., N\}$. The prediction errors are then summed up and normalized over all explored cells, resulting in the average prediction error

$$\Delta_{\mathbf{p}} = \sqrt{\frac{\sum_{k} (C_{\mathbf{p}}^{(k)} - c^{(k)})^2}{N}} \tag{8}$$

with N being the total number of grid cells.

Optimization Using Evolution Strategies

Searching for the best set of parameters $\boldsymbol{p}$ in terms of minimizing the average prediction error is a typical optimization problem. Preliminary experiments showed that the particular optimization problem considered here is multi-modal, i.e. the function defined in Equation (8) comprises multiple possibly deceptive local optima apart from the global optimum.

Therefore, it was decided to use Evolutionary Algorithms (EA), which are known for their capability to perform well in multi-modal search spaces. EA are probabilistic, generational, population based optimization strategies that mimic the natural evolution based on Darwin's principle of the "survival of the fittest" by repeated simulation of a generational life cycle. Starting from a population of randomly initialized solutions (individuals) a generational cycle is started by evaluating all individuals with a target or fitness function, e.g. Equation (8), and stochastically selecting the best individuals to generate offspring for the next generation. Offspring is typically generated by making copies of the selected individuals (parents), which can then be altered through mutation (i.e. random changes) or crossover (i.e. mixing partial copies from multiple parents). While mutation and crossover enables the EA to explore the whole search space, the quality of the solutions represented by the individuals increases gradually over multiple generations simply due to iterated selection.

In this work, a special type of EA called Evolution Strategies (ES) is applied to the optimization problem. ES were developed by Rechenberg and Schwefel to solve practical application problems in mechanical engineering (Rechenberg, 1973; Schwefel, 1981). In contrast to Genetic Algorithms (GA) developed by Holland (Holland [1975]), ES abstract the key elements of EA to perform an efficient search especially on real-valued search spaces and typically use specialized mutation operators.

For example, instead of a stochastic selection scheme, ES use typically a $(\mu+\lambda)$- or (μ,λ)-strategy, selecting only the μ best individuals as parents to generate λ offspring. In case of a (μ,λ)-strategy these offspring set up the next generation to be evaluated. A $(\mu+\lambda)$-strategy, on the other hand, combines the μ parents and the λ offspring in the next generation. While a $(\mu+\lambda)$-strategy is guaranteed to improve monotonically due to elitism, the (μ,λ)-strategy is usually better suited for multi-modal search spaces because it enables the ES to escape from local optima.

Self-Adapting Mutation Operators and Crossover

ES often use so called self-adapting mutation operators. A mutation operator that adds random values with a fixed variance to the decision parameters does not perform well in all cases. A fixed mutation step size can be too small in case an individual is in a local optimum or still far from the optimum. In this case the optimization algorithm would either converge prematurely or convergence would be achieved too slowly. On the other hand, a fixed mutation step size can also be too large if the individual is already close to the global optimum. Because oversized perturbations lead the offspring away from the global optimum this might prevent the algorithm to converge to an optimum at all.

Therefore, Rechenberg and Schwefel introduced so called strategy parameters σ^2 that define the mutation step size. Strategy parameters are used either one for all decision parameters (global mutation) or using a vector of strategy parameters one for each decision parameter (local mutation). Local mutation enables independent adaptation of the strategy parameters for each dimension of the problem space. Each individual has a unique set of decision and strategy parameters which are altered with a probability of p_{mut} by calculating the new strategy parameters σ'_j and the new model parameters p'_j as

$$\begin{aligned} \sigma'_j &= \sigma_j \cdot e^{\tau \cdot N(0,1) + \tau' \cdot N_j(0,1)}, \\ p'_j &= p_j + \delta p_j \cdot \sigma'_j \cdot N_j(0,1). \end{aligned} \tag{9}$$

$N(0, 1)$ and $N_j(0, 1)$ are random numbers independently drawn from a normal Gaussian distribution, τ is an overall learning rate and τ' is a coordinate wise learning

rate. The parameters δp_j have to be assigned to each strategy parameter in advance. They specify the range for each decision variable and thus normalize the mutation step size to the size of the search space.

Ultimately, each individual can be altered by crossover of two parents $<p'_1, \sigma'_1>$ and $<p'_2, \sigma'_2>$ with a probability of p_{cross} as

$$\begin{aligned} \sigma''_j &= \frac{\sigma'^1_j + \sigma'^2_j}{2}, \\ \mathbf{p}'' &= \left\langle p_1'^{R_1(1,2)}, p_2'^{R_2(1,2)}, .., p_n'^{R_n(1,2)} \right\rangle. \end{aligned} \tag{10}$$

Here, n gives the number of parameters, and $R_j(1, 2)$ is a random variable that is used to select for each decision parameter p_j with equal probability whether p_j is chosen from parent 1 or parent 2.

A non elitist generation strategy like the (μ,λ)-strategy will in the long run favor individuals with suitable strategy parameters, since their offspring will perform better than offspring from individuals with poorer strategy parameters. This way the ES is able to self-adapt the strategy parameters to the local properties of the search space.

Optimization Strategy Used in this Work

In order to determine suitable parameters for the ES, several test runs were performed where the model function was used with set parameters as the ground truth. These test runs showed that the performance of the ES with local mutation does not depend heavily on the actual value of the parameters used. As a parameter set that produced very good fitting results in the preliminary tests, enables self-adaption, and reduces the chance of premature convergence, the following parameters were used here: $\mu = 10$, $\lambda = 50$, $p_{mut} = 1.0$, $p_{cross} = 0.01$. The initial step width was set to $\sigma_j = \sigma_{init} = 0.1$. According to (Eiben & Smith, 2003), the overall learning rate and the coordinate wise learning rate were chosen to be $\tau = \{\sqrt{(2\cdot\sqrt{n})}\}^{-1} \approx 0.435$ and $\tau' = \{\sqrt{(2n)}\}^{-1} \approx 0.267$, respectively. The parameter ranges p_j were chosen to be [0,5] for C_{00}, [0,20] for C_S and C_A, [-4,4] for x_0 and y_0, [-180°,180°] for θ, and [-1,1] for C_B.

To reduce the chance to get stuck in a local optimum further, four independent ES optimization runs with 25000 fitness calls in each run were performed. Finally, the best individual found was selected and considered to be the best possible fit for the given gas distribution map.

SOURCES OF INACCURACY

A perfect agreement between the statistical gas distribution map and the analytical model given by Equation (7) cannot be expected in general for several reasons. First of all, the assumptions the analytical gas distribution model is based on might not be fulfilled. This mainly applies to the assumption of constant, unidirectional wind. Although no artificial air current was produced to create a dominant constant flow, the gas distribution was strikingly stable in the experiments presented in this paper, most likely due to stable temperature gradients in the room (Wandel et al., 2003). For that reason, the observed gas distribution can often be approximated reasonably well with the analytical model. Assuming only minor variations of the dominant air stream, the model will indicate the average wind direction while poorer fitness values are expected in case of stronger variations.

A further assumption, which is not completely fulfilled, is that the statistical gas distribution map represents the true time-averaged gas distribution as it would appear over infinitely long time. Due to the local character of gas sensor measurements, it takes some time to build the statistical grid map model. In addition to spatial coverage, a certain amount of temporal averaging is also necessary to represent the time-constant profile of the gas distribution. The basic structures in the map were found to stabilize within the first hour of the mapping experiments (Lilienthal & Duckett, 2004a). During this time, transient concentration peaks caused by turbulence might not be sufficiently averaged out and thus can be preserved in the statistical gas distribution map as minor deviations from the smooth course of the distribution. While this is generally more of a problem in regions of low concentration (because the peak to time-average ratio is higher there (Roberts & Webster, 2002), it is especially problematic concerning the experiments where the robot was reactively controlled to avoid low concentrations, causing a low density of measurements in regions where the average concentration is low. The distortions due to rudiments of turbulent peaks tend to influence the fit result because the region of low concentration was typically much larger than the area of high concentration in the experiments considered in this work. In order to compensate for this effect, a modified fitness function was used that compares the square of the prediction of the analytical model $C^{(k)}$ with the square of the corresponding grid cell value $r^{(k)}$.

$$\Delta_{\mathrm{p}} = \sqrt{\frac{\sum_{k} ((C^{(k)})^2 - (r^{(k)})^2)^2}{N}}, \tag{11}$$

Hence, the influence of deviations in regions of low average concentration on the fitness function $\Delta \boldsymbol{p}$ is reduced compared to deviations in regions of high average concentration.

It is important to note that an important result of the fitting procedure, in addition to the obtained value of the fit parameters, is the fitness value itself, i.e. the average prediction error given in Equation (8) and Equation (11). A poor fitness indicates that the applied analytical model cannot describe the observed gas distribution (in terms of the statistical model) faithfully. If so, the source position estimate obtained from the best fit cannot be considered reliable and the best information at hand is the CME. On the other hand, if a good fitness can be achieved, the corresponding parameter set comprises a reasonable estimate of the source location as will be seen in Sec. "Results".

Please note that because of the difficulty to measure absolute concentrations with metal oxide gas sensors in an uncontrolled environment (Lilienthal & Duckett, 2004b), the gas distribution map directly models the sensor readings. Although the calibration function of the gas sensors is nearly linear if a small range of concentrations is considered as in the experiments presented here, it is therefore not possible to determine the absolute value of C_{00} and C_B. The remaining parameters, however, have a meaningful interpretation because they refer to the geometric dimension of the distribution profile. Therefore the fitting process described in Sec. "Step 2 – Learning Parameters of the Analytical Model" can be seen as model-based shape analysis of the statistical gas distribution grid map.

STEP 3: SELECTION OF SOURCE LOCATION ESTIMATE

The results of the second step are the parameters of the best fitting analytical model and the respective fitness value. The parameters of the analytical model include the coordinates of the source position. In the third step, the asymmetry ratio C_A/C_S and the optimal fitness value are used to determine whether the source estimate corresponding to the best analytical model or the estimate given by the maximum in the statistical gas distribution map is more reliable. This step corresponds to an analysis of the observed conditions, particularly the airflow conditions.

SABEC1: SHAPE ANALYSIS-BASED ESTIMATION CHOICE STRATEGY

Based on the observations above we propose the following strategies for determining an improved estimate of the gas source location. The first strategy is referred

Figure 5. Flowchart of the SABEC1 and SABEC2 strategy to select the most truthful estimate of the source location. Based on the fitness and the asymmetry of the best fit model obtained from optimizing the average prediction error with the ES, either the CME or the BFE is chosen, or no prediction is made at all.

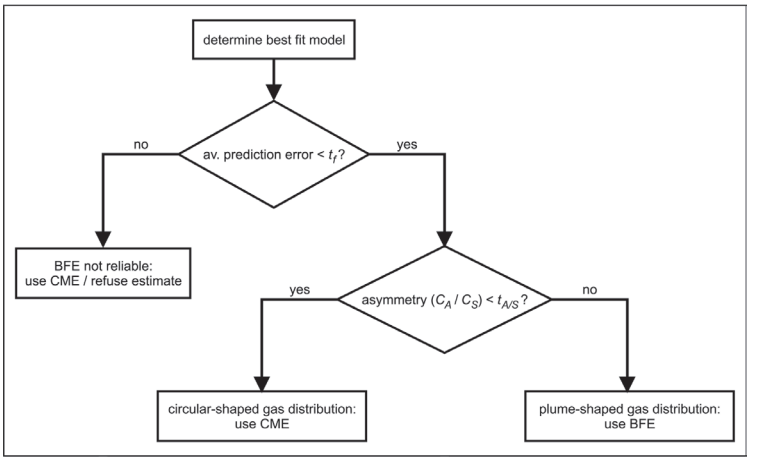

to in the following as SABEC1 strategy (Shape Analysis-Based Estimation Choice strategy). It is sketched in Figure 5, together with the SABEC2 strategy previously discussed. When applying the SABEC1 strategy, the best fitting analytical model is determined with the ES and the average prediction error (fitness) and the asymmetry ratio C_A/C_S are compared with corresponding threshold values t_f and $t_{A/S}$. If the average prediction error is above t_f, the BFE is considered not reliable and the CME is used to estimate the source location. The CME is also chosen in case of a prediction error below t_f and a low value of C_A/C_S, indicating a relatively weak stationary air current. Finally, the BFE is used in case of a small prediction error and a C_A/C_S value above $t_{A/S}$.

SABEC2: EVENTUALLY REFUSING TO PREDICT THE GAS SOURCE LOCATION

A further improvement of the accuracy can be obtained if it is acceptable to refuse a prediction of the source location in case both the BFE and the CME were found to be unreliable. The important observation in this context is that a poor fitness value, corresponding to an unstable wind field, for example, typically indicates an expanded region of high values in the statistical gas distribution map where the maximum can easily vary due to rudiments of turbulent fluctuations. This entails the

SABEC2 strategy also indicated in Figure 5, which selects the BFE or the CME in case of a reliable best fit model (depending on the asymmetry of the model) while it does not make a prediction in case of an unreliable model with a poor fitness.

EXPERIMENTAL SETUP

Robot and Gas Sensors

The experiments were performed with a Koala mobile robot equipped with the Mark III mobile nose (Lilienthal & Duckett, 2003a), comprising 6 tin oxide gas sensors manufactured by Figaro (see Figure 6). 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, including low selectivity, comparatively high power consumption (caused by the heating device) and weak durability. In addition, metal oxide sensors are subject to a long response time and an even longer recovery time (Lilienthal & Duckett, 2003a). However, this type of gas sensor is most often used on mobile robots because it is inexpensive, highly sensitive and relatively unaffected by changing environmental conditions such as room temperature and humidity. The gas 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. Papst Fans (Papst 405F) were used to generate an airflow of $8m^3/h$. Multiple, redundant sensor types were used only to increase the robustness of the system (there was no attempt to discriminate different odors). The distance between the two sets of sensors was 40 cm.

Figure 6. Koala robot with the Örebro Mark III mobile nose and the gas source used in the experiments. The small image in the top left corner shows a Figaro gas sensor used in the Mark III mobile nose.

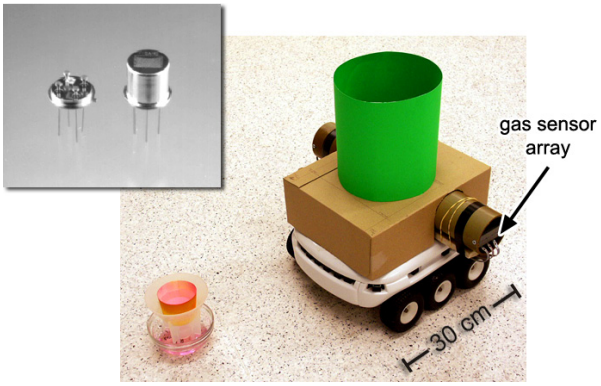

Environment, Gas Source and Absolute Positioning System

All experiments were carried out in a rectangular laboratory room at ¨Orebro University (size 10.6×4.5 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 from the external absolute positioning system WCAPS (Lilienthal & Duckett, 2003c), which was used to track the colored cardboard "hat" on top of the robot (the "hat" can be seen in Figure 6).

To emulate a typical task for an inspection robot, a gas source was chosen to imitate a leaking tank. This was realized by placing a paper cup filled with ethanol on a support in a bowl with a perimeter of 12cm (see Figure 6). 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. The air conditioning system in the room was deactivated.

Results

An illustration of the optimization process described in Sec. "Step 2 – Learning Parameters of the Analytical Model" is given in Figure 7 and Figure 8. At the bottom right of each figure, the statistical gas distribution grid map, which is to be approxi-

Figure 7. Example of the model selection process. The statistical gas distribution grid map that is to be approximated can be seen at the bottom of the figure to the right of a fitness plot, which shows the average prediction error according to Equation (5.1). Four individual solutions corresponding to best model obtained after 500, 2500, 10000 and 25000 fitness calls are depicted in the top row.

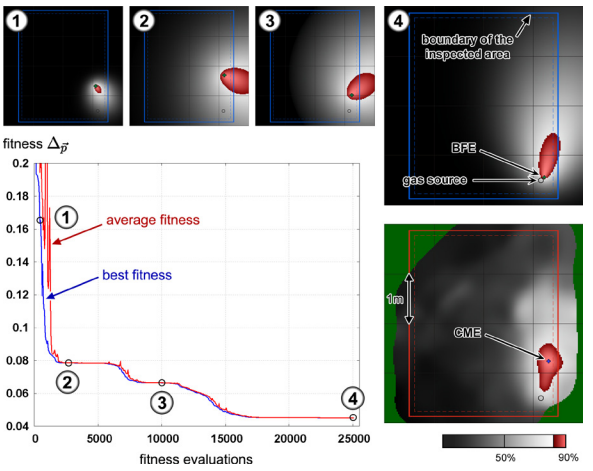

Figure 8. Example of the model selection process. In the same way as in Figure 7, a statistical gas distribution grid map, the fitness plot and best fit models obtained after 500, 2500, 7500 and 25000 fitness calls are shown.

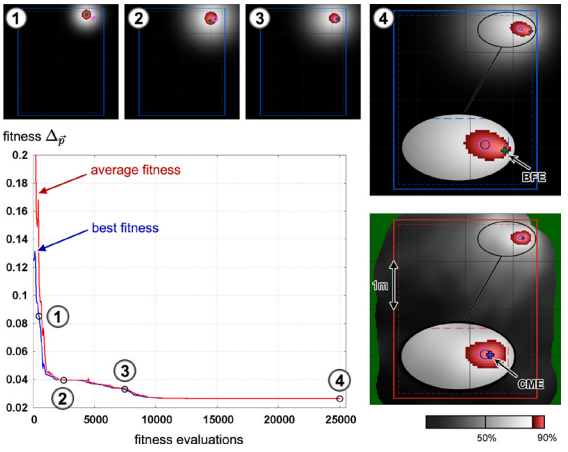

mated, is shown. The gas distribution maps in both figures were created from data collected in 90 minutes of exploration in two different experiments where the robot was reactively controlled as a Braitenberg vehicle which follows the local gradient sensed with a pair of gas sensor arrays (Lilienthal & Duckett, 2004a). To the left of the gas distribution map, the evolution of the average and the best fitness in the population is plotted depending on the number of fitness evaluations. In addition, four individual solutions are depicted in the top row. As indicated in the fitness plot, these solutions correspond to the best model obtained after a certain number of fitness evaluations. Different shadings are used to indicate relative concentration values in the gas distribution grid maps and the visualization of the model output (dark → low, light → high). Values higher than 90% of the maximum are displayed with a second range of dark-to-light shadings (of red), and unexplored cells are shown with a different color (dark green).

A good agreement between the statistical gas distribution map and the best fit model after 25000 fitness evaluations was obtained in both the cases considered in Figure 7 and Figure 8, indicated by a low fitness value. The fitting results, however, suggest that the experiments were carried out under different conditions. While the long stretched out shape of the fitted analytical model displayed in Figure 7 points to a relatively strong air current, a weaker air stream is indicated by the more circular shape of the best fit model shown in Figure 8. Such a difference is expressed by the ratio of the parameters C_A/C_B, which was 15.5 in case of the distribution in Figure 7 but only 2.8 for the distribution in Figure 8. In case of a strong asymmetry, i.e. a high C_A/C_B value, the CME is typically a poor approximation of the true source

location. Due to the small variation of the gas distribution in upwind direction, small inhomogeneities of the wind field can cause a large displacement of the maximum along the direction of the air current. As in the example shown in Figure 7, the BFE is a more reliable estimate of the gas source position in case of a high ratio C_A/C_B.

On the other hand, the distance of the BFE to the gas source was often higher compared to the CME in case of a more symmetric distribution (see the example in Figure 8). A potential reason is that a circular distribution (with a corresponding source estimate in the centre of the region of increased concentration) can easily be interpreted as a plume-like structure (with a source estimate at the boundary of the same region), for example if the distribution was not explored properly around the gas source. Another reason might be due to the distortion and broadening effect due to the slow decay of the gas sensors mentioned when considering the impact of the sensor dynamics above. Regions of high concentration appear expanded in the map and consequently the fit results tend to be displaced with respect to the true gas source location by a certain amount. Here, the CME is typically a more accurate approximation of the gas source that is less sensitive to small deviations from the ideal distribution profile.

SABEC1 Results

The accuracy of different gas source location estimates was compared based on 97 snapshots of concentration maps obtained in 11 mapping experiments, including a total of 34 hours of exploration. In four runs the robot moved along a predefined path (inwards and outwards a rectangular spiral) while it was reactively controlled as a gas-sensitive Braitenberg vehicle in the remaining seven trials. The explored area was approximately $2.4 \times 2.4\text{m}^2$ in the experiments with a predefined exploration path, and it was approximately $3.7 \times 3.7\text{m}^2$ in the Braitenberg vehicle experiments. Snapshots of the gas distribution map were taken in intervals of 15 minutes starting after one hour of exploration to assure that the maps would represent mainly the stationary properties of the gas distribution (see Sec. “Sources of Inaccuracy”).

A comparison of the results can be seen in Table 1. The first column specifies in which cases the BFE was chosen as the source position estimate instead of the

Table 1. Gas source localisation error obtained with the SABEC1 strategy

strategy (when to choose BFE)	*tf , tA/S*	[cm]
never	-, -	26.2 ± 20.9
if av. prediction error < *tf*	0.0470, -	22.2 ± 15.5
if av. prediction error < *tf* and *CA/CS* > *tA/S* (SABEC1)	0.0575, 8.0	17.8 ± 10.4

CME and the second column shows the optimal threshold parameters found. Finally, the average distance between the estimate and the true source location is given in the third column. The obtained accuracy was 15% better compared to using only the CME if the BFE was always selected in case a good fitness, and it was 32% lower if the BFE was selected with the SABEC1 strategy that additionally considers the asymmetry of the fit result. Note that only a weak dependency of the obtained accuracy on the values of the threshold parameters was observed. Thus, the exact choice of these values was found to be not critical.

SABEC2 Results

The accuracy obtained with the SABEC2 strategy and the percentage r_{np} of gas distribution maps, for which no prediction of the gas source location was made, are shown in Figure 9 depending on the chosen fitness threshold t_f and using a constant value of $t_{A/S} = 8.0$. As an example, a fitness threshold of $t_f = 0.055$ corresponded to an average error of (14.8 ± 7.3) cm while a prediction of the gas source location was refused for 31% of the gas distribution map snapshots. The graph also shows the level of accuracy that was obtained using the CME only.

The monotonic evolution of the average error shown in Figure 9 demonstrates that the obtained fitness value is a suitable measure for the confidence about the gas source location estimate obtained from either the CME or the BFE. Thus, it

Figure 9. Dependency of the average error of gas source prediction on the fitness threshold t_f when the SABEC2 strategy is applied to determine whether the CME or the BFE is chosen as the final source location estimate, or a prediction is refused (crosses). The rate of gas distribution maps for which no prediction was made is also indicated (filled circles). The value of $t_{A/S}$ was set to 8.0.

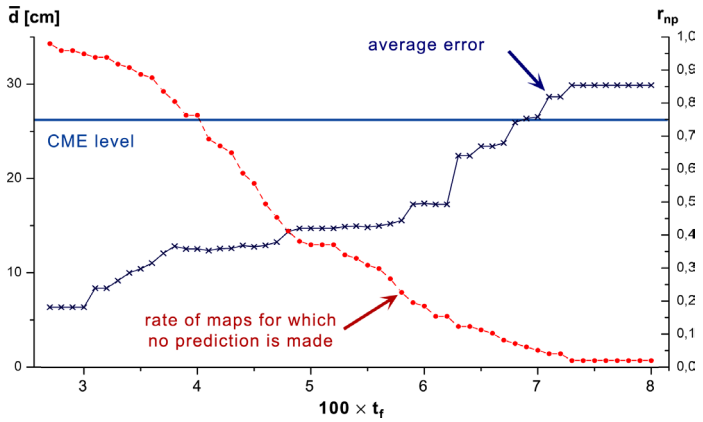

seems possible to compute from the fitness value a radius of an area around the estimate position where the source is expected with high certainty, or to choose a fitness threshold depending on the required accuracy.

CONCLUSION

The method presented in this chapter is visually summarized in Figure 10. It computes an estimate of the location of a single gas source from a set of localized gas sensor measurements and consists of three steps.

In the first step, a statistical model of the gas distribution is computed from a sequence of localized gas sensor measurements using the Kernel DM algorithm. This model represents the spatial structure of the time-averaged gas distribution in a two-dimensional grid map and averages out the transitory effects of turbulent gas dispersal. In this work, gas sensor readings and the corresponding position data were collected with a mobile robot while it was driven according to a particular exploration strategy.

In the second step, the parameters of an analytical model of the time-averaged gas distribution are learned by nonlinear least squares fitting of the analytical model to the statistical gas distribution map using Evolution Strategies. The analytical model

Figure 10. Summary of the proposed method to estimate the location of a single gas source

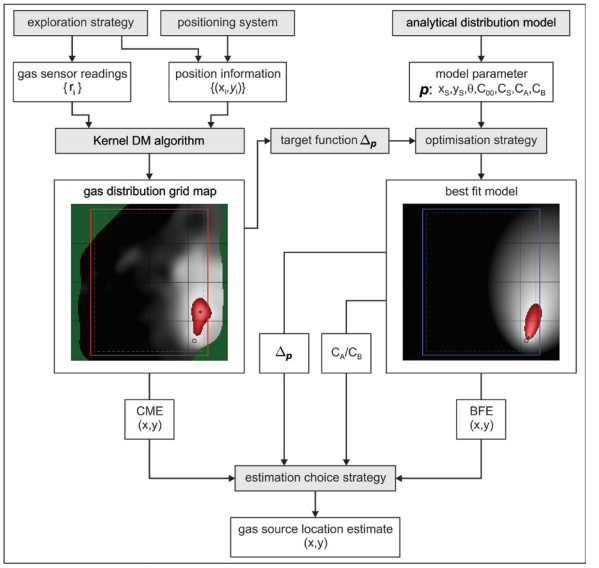

is based on a physical model that describes the time-averaged gas distribution under the assumptions of isotropic and homogenous turbulence, a constant, uniform wind field, and a constantly emitting point source on the floor. Accordingly, the fitting procedure includes an analysis of the statistical gas distribution map, which allows interpreting the shape of the statistically modeled gas distribution by quantitative means. This enables us to distinguish gas distribution maps according to different airflow conditions.

As a result of the second step we obtain the parameters of the best analytical model and the respective fitness value. The parameters of the best fitting analytical model include the coordinates of the source position (best fit estimate, BFE), and the asymmetry ratio C_A/C_S. These values are used in the third step, to determine whether the BFE or the CME is deemed more reliable. The SABEC1 strategy that always selects either the CME or the BFE is described and the SABEC2 strategy that might also refuse to output a source location estimate is also detailed. The third step formalizes the analysis of the observed conditions, particularly the airflow conditions.

The proposed method was evaluated based on a total of 34 hours of mapping experiments. Using the SABEC1 strategy, the obtained estimate of the gas source location was found to be 32% more accurate than the CME. Moreover, it was demonstrated that an even better accuracy can be obtained by using the SABEC2 strategy at the cost of an increasing rate of gas distribution maps for which no prediction about the gas source location is made. The BFE was found to be more reliable than the CME if the mapping experiment was performed under conditions of a relatively strong constant air current, indicated by a high asymmetry ratio C_A/C_S of the best fitting analytical model. Finally, it was found that the quality of the best fitting analytical model (in terms of the achieved fitness value) can be used as a confidence measure for approximating the source position. If the observed airflow is heavily deviating from the assumption of a uniform, constant flow, the best fitness is expected to be still comparatively poor and this indicates that the source location estimate obtained from the parameters of the best analytical model is not very reliable. In fact, it was found that if the best fitting analytical model approximated the gas distribution map rather poorly, both BFE *and* CME were found to be less accurate than in cases where the average deviation between the best fit model and the statistical gas distribution grid map was small.

FUTURE RESEARCH DIRECTIONS

The proposed method can be extended in several ways. First, a more sophisticated algorithm than Kernel DM can be used to build the statistical gas distribution model from a sequence of localized gas sensor measurements. In particular, algorithms that

estimate both the mean field and a predictive variance are of interest, such as Kernel DM+V (Lilienthal et al., 2009) or the method using Gaussian Process mixtures proposed by (Stachniss et al., 2009). The maximum of the predictive variance provides an additional estimate of the gas source location. Another interesting extension is to use wind information collected together with the gas sensor readings for building the statistical gas distribution model as in (Reggente & Lilienthal, 2009a).

Wind information could also be used in the second step of the approach introduced in this chapter, for example as initial guess for the corresponding parameters of the analytical model or as an alternative means to determine the best estimate of the gas source location by distinguishing different airflow conditions. Based on a model of the wind field, more involved analytical models of gas dispersal could be used instead of the model described.

The proposed approach builds a two-dimensional statistical gas distribution model from the gas sensor measurements, which were collected according to a predefined exploration strategy. A further possibility for future research is therefore to investigate adaptive exploration, i.e. strategies that determine where to collect future measurements based on the current gas distribution model. Ultimately, a three-dimensional approach needs to be developed. An initial step is to build three-dimensional statistical gas distribution models based on gas sensor measurements such as in (Reggente & Lilienthal, 2009b). On the long run, building the gas distribution model should consider all available relevant information, including, for example, the temperature distribution and the spatial outline of the environment.

ACKNOWLEDGMENT

This work has partly been supported by the EC under contract number FP7-224318-DIADEM: Distributed Information Acquisition and Decision-Making for Environmental Management.

REFERENCES

Blackmore, B. S., & Griepentrog, H. W. (2002). A future view of precision farming. In *Proceedings of PreAgro Precision Agriculture Conference* (pp. 131-145), Müncheberg, Germany, Center for Agricultural Landscape and Land Use Research (ZALF).

Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. MIT Press/ Bradford Books.

DustBot Consortium. (2006). *DustBot-networked and cooperating robots for urban hygiene*. Retrieved from http://www.dustbot.org

Eiben, A. E., & Smith, J. E. (2003). *Introduction to evolutionary computing*. New York, NY: Springer.

Hinze, J. O. (1975). *Turbulence*. New York, NY: McGraw-Hill.

Holland, J. H. (1975). *Adaptation in natural and artificial systems*. MIT Press.

Ishida, H., Nakamoto, T., & Moriizumi, T. (1998). Remote sensing of gas/odor source location and concentration distribution using mobile system. *Sensors and Actuators. B, Chemical*, *49*, 52–57. doi:10.1016/S0925-4005(98)00036-7

Kowadlo, G., & Russell, R. A. (2003). Naive physics for effective odour localisation. In *Proceedings of the Australian Conference on Robotics and Automation*.

Lilienthal, A. J., & Duckett, T. (2003a). A stereo electronic nose for a mobile inspection robot. In *Proceedings of the IEEE International Workshop on Robotic Sensing*.

Lilienthal, A. J., & Duckett, T. (2003b). An absolute positioning system for 100 Euros. In *Proceedings of the IEEE International Workshop on Robotic Sensing*.

Lilienthal, A. J., & Duckett, T. (2003c). Creating gas concentration gridmaps with a mobile robot. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 118–123).

Lilienthal, A. J., & Duckett, T. (2004a). Building gas concentration gridmaps with a mobile robot. *Robotics and Autonomous Systems*, *48*(1), 3–16. doi:10.1016/j.robot.2004.05.002

Lilienthal, A. J., & Duckett, T. (2004b). Experimental analysis of gas-sensitive Braitenberg vehicles. *Advanced Robotics*, *18*(8), 817–834. doi:10.1163/1568553041738103

Lilienthal, A. J., Loutfi, A., Blanco, J. L., Galindo, C., & Gonzalez, J. (2007). A Rao- Blackwellisation approach to GDM-SLAM integrating SLAM and gas distribution mapping. In *Proceedings of the European Conference on Mobile Robotics*, (pp. 126–131).

Lilienthal, A. J., Reggente, M., Trincavelli, M., Blanco, J. L., & Gonzalez, J. (2009). A statistical approach to gas distribution modelling with mobile robots–the Kernel DM+V Algorithm. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, (pp. 570-576).

Lilienthal, A. J., Streichert, F., & Zell, A. (2005). Model-based shape analysis of gas concentration gridmaps for improved gas source localisation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, (pp. 3575–3580).

Nakamoto, T., Ishida, H., & Moriizumi, T. (1999). A sensing system for odor plumes. *Analytical Chemistry News & Features*, *1*, 531–537.

Rechenberg, I. (1973). *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. Fromman-Holzboog.

Reggente, M., & Lilienthal, A. J. (2009a). Using local wind information for gas distribution mapping in outdoor environments with a mobile robot. In *Proceedings of IEEE Sensors*, (pp. 1715-1720).

Reggente, M., & Lilienthal, A. J. (2009b). Three-dimensional statistical gas distribution mapping in an uncontrolled indoor environment. *AIP Conference Proceedings Volume 1137: Olfaction and Electronic Nose - Proceedings of the 13th International Symposium on Olfaction and Electronic Nose* (ISOEN), (pp. 109-112).

Roberts, P. J. W., & Webster, D. R. (2002). Turbulent diffusion. In Shen, H., Cheng, A., Wang, K.-H., Teng, M. H., & Liu, C. (Eds.), *Environmental fluid mechanics-theories and application*. Reston, VA: ASCE Press.

Russell, R. A. (1999). *Odour sensing for mobile robots*. World Scientific.

Schwefel, H.-P. (1981). *Numerical optimization of computer models*. New York, NY: John Wiley & Sons, Inc.

Shraiman, B., & Siggia, E. (2000). Scalar turbulence. *Nature*, *405*, 639–646. doi:10.1038/35015000

Stachniss, C., Plagemann, C., & Lilienthal, A. J. (2009). Learning gas distribution models using sparse Gaussian process mixtures. *Autonomous Robots*, *26*(2-3), 187–202. doi:10.1007/s10514-009-9111-5

Wandel, M. R., Lilienthal, A. J., Duckett, T., Weimar, U., & Zell, A. (2003). Gas distribution in unventilated indoor environments inspected by a mobile robot. In *Proceedings of the IEEE International Conference on Advanced Robotics*, (pp. 507–512).

ADDITIONAL READING

Mobile Robot Olfaction

Lilienthal, A. J., Loutfi, A., & Duckett, T. (2006). Airborne Chemical Sensing with Mobile Robots. *Sensors (Basel, Switzerland)*, *6*, 1616–1678. doi:10.3390/s6111616

Russell, R. A. (1999). *Odour Sensing for Mobile Robots*. World Scientific.

Environmental Monitoring

Gilbert, R. O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Wiley.

Fluid Dynamics

Ferziger, J. H., & Peric, M. (2001). *Computational Methods for Fluid Dynamics*. Springer.

Hinze, J. O. (1975). *Turbulence*. New York: McGraw-Hill.

Evolution Strategies

Schwefel, H.-P. (1995). *Evolution and Optimum Seeking*. Wiley & Sons.

ENDNOTES

[1] Please note that because of calibration issues with gas sensors the sensor response is typically normalised but otherwise modelled directly. To emphasize this procedure we use the term "gas distribution map" while, if the input is given in absolute concentration values, we call the resulting representation a "gas concentration map". Accordingly, in the former case, the term "concentration maximum estimate" does not imply that the absolute concentration value could be precisely estimated. The *location* of the concentration maximum will be estimated correctly, however, assuming that the response of all involved sensors depend on the concentration in the same, monotonous way.

[2] Not to confuse with the kernel width σ in Equation (2.1), for example.