

Dynamic Modelling of Electronic Nose Systems

Graham Ellis Searle

A thesis submitted in partial fulfilment of the
requirements for the degree of Doctor of Philosophy in
Engineering

School of Engineering

University of Warwick

September 2002

Contents

Contents	i
Lists of tables and figures	vi
Acknowledgements	xiv
Declaration	xv
Abstract	xvi
Abbreviations	xvii
1 Introduction	1
1.1 Aims of this thesis	3
1.2 Summary of thesis	3
2 Electronic nose system technology	6
2.1 The historical development of electronic noses	7
2.2 Electronic nose hardware	8
2.2.1 Metal oxide semiconductor sensors	10
2.2.2 Conducting polymer sensors	13
2.2.3 Acoustic wave sensors	15
2.2.4 Field effect based sensors	17
2.2.5 Pellistors	18
2.3 Signal and data processing in electronic nose systems	18
2.3.1 Pre-processing	19
2.3.2 Processing of steady-state data	21
2.3.2.1 Forward and inverse linear models	22
2.3.2.2 Nonlinear response models	23
2.3.2.3 Black box models	24
2.3.2.4 Cluster analysis	24
2.3.2.5 Principal components analysis	26

2.3.2.6 K-nearest neighbours	28
2.3.2.7 Soft independent modelling of class analogies	29
2.3.2.8 Discriminant function analysis	30
2.3.2.9 Artificial neural networks	31
2.3.3 Processing of dynamic data	32
2.3.3.1 Exponential response models	33
2.3.3.2 Thermal modulation	34
2.4 Current and future applications for electronic nose systems	35
3 Cyanobacteria strain identification using system identification techniques	38
3.1 Black box modelling	39
3.2 Bacterial strain identification experiment	41
3.3 Static sensor response models	43
3.3.1 FIR models	45
3.3.2 More complex model forms	46
3.4 Dynamic sensor response models	47
3.4.1 FIR models	49
3.4.2 ARX models	50
3.4.3 MAX and ARMAX models	51
3.4.4 Models for dynamic data with extra inputs	53
3.5 An alternative train-test strategy	54
3.5.1 Static sensor response models	54
3.5.2 Dynamic sensor response models	56
3.6 Conclusions	58
4 Cyanobacteria growth phase identification using system identification techniques	60
4.1 Black box modelling	61
4.2 Growth phase identification experiment	62

4.3 Static sensor response models	64
4.3.1 MIMO models for static sensor responses	67
4.4 Dynamic sensor response models	68
4.5 Models for filtered dynamic data	72
4.6 Conclusions	75
5 The experimental electronic nose system	78
5.1 Overview of the system	78
5.2 The sample containers and gas delivery system	80
5.3 The chamber unit	82
5.4 The interface unit	82
5.5 The computer and controlling software	83
5.6 System operation	85
5.7 Example outputs	85
6 Differential equation models for electronic noses	88
6.1 Physical models versus empirical models	88
6.2 Nose system models	89
6.2.1 Odour generation	90
6.2.2 Odour transport	95
6.2.3 Sensor response	102
6.2.4 Interface electronics	103
6.3 Using (forward) models in an odour classification system	104
6.4 Identifiability and estimation of the parameters in ODE models	105
7 A reaction-based model for MOS sensors	108
7.1 Chemical reaction equations	108
7.2 Model assumptions	109
7.3 Mathematical model	112

7.4 Decoupled model with only three equations	115
7.5 Steady states of full model	116
7.6 Initial conditions as steady states from odourless system	117
7.7 Model identifiability	119
7.7.1 Controllability rank criterion	120
7.7.2 Observability rank criterion	123
7.7.3 Identifiability analysis	125
7.8 Conclusions	140
8 Parameter estimation and validation of nose system model	142
8.1 Experimental data	142
8.2 Initial model fitting and parameter estimations	146
8.2.1 Parameter estimation for a single sensor response	148
8.2.2 Model fitting to multiple analyte responses	150
8.3 Parameter estimates over the range of the entire experiment	153
8.3.1 Investigation of correlations between fitted parameters as a partial check for reliable parameter estimation	154
8.3.2 Investigation of the effects of variations in the temperature and humidity on the fitted parameters	156
8.4 Evaluation of the discrimination performance of parameter extraction versus steady state methods	161
8.5 Discussion and conclusions	163
9 Parameter extraction technique applied to cyanobacteria data	166
9.1 Cyanobacteria strain identification	167
9.1.1 Experimental data	167
9.1.2 Initial model fitting and parameter estimations	168
9.1.3 Parameter estimations for responses throughout the whole experiment	173

9.2 Cyanobacteria growth phase identification	176
9.2.1 Experimental data	177
9.2.2 Initial model fitting and parameter estimations	179
9.2.3 Parameter estimations for responses throughout the whole experiment	180
9.3 Discussion	183
9.4 Conclusions	186
10 Conclusions	188
10.1 Black box modelling	188
10.2 Models based on the physical system	190
10.3 Classification by parameter extraction	190
10.4 Final remarks	195
References and bibliography	197
Appendix 1	207
Appendix 2	208

List of tables and figure

Tables:

Table 2.1 Some examples of electronic nose applications

Table 3.1 The linear black box model structures used to analyse the electronic nose data.

Table 3.2 Classification success rates obtained using zero-delay FIR models of orders two or three for static (pre-processed) data.

Table 3.3 Successful classification rates for MISO FIR models having six, seven or eight inputs.

Table 4.1 The success rates for FIR MIMO models of orders between one and four for static (pre-processed) data from the growth phase experiment. The percentages are averages over 10 different random reorderings using two-fold (50%) cross-validation.

Table 4.2 The success rates obtained for 10th order FIR MIMO models for dynamically pre-processed and randomly reordered growth phase data. The models were trained on half of the data and tested on the remaining half, i.e. two-fold (or 50%) cross-validation was employed.

Table 8.1 The ranges selected for the various parameters for curve fits to sensor 1 responses throughout the entire duration of the experiment.

Table 8.2 Matrix of correlation coefficients (given to four significant figures) for the fitted parameters to acetone responses of sensor 1, and the recorded temperature and humidity values for the experiment.

Table 9.1 Comparison of the average classification success rates using FIR models of order one for extracted k_2 values with those using static (pre-processed) data for the classification of cyanobacteria growth phase. The rates given are averages over 100

different random reorderings of the data set. For each reordering, the first 18 data vectors were used to train the model, and the remaining 10 to test it.

Figures:

Figure 2.1 The hardware components of a typical electronic nose system.

Figure 2.2 Configuration of a basic tin dioxide gas sensor.

Figure 2.3 Basic configuration of a delay line surface acoustic wave device.

Figure 2.4 Schematic showing the basic construction of a BAW odour sensor.

Figure 2.5 A typical response of a MOS odour sensor to a square pulse of acetone, x_b is the sensor baseline, and x_a the steady-state response.

Figure 2.6 An example of a dendrogram produced by HCA, similar to those produced from a successful electronic nose experiment. The data points are labelled with their known class, but this information is not used in the cluster analysis.

Figure 2.7 A typical PCA plot of the first and second principal components from an electronic nose experiment to discriminate between three different analytes. Although six MOS sensors were used in the experiment, adequate group separation is observed with just two plotted dimensions.

Figure 3.1 Diagrammatic representation of the electronic nose modelling problem.

Figure 3.2 A 14 hour section of the raw data, showing the response of a single odour sensor during part of the experiment. The voltage is directly proportional to sensor resistance.

Figure 3.3 Plot of successful classification rates against orders of A and B polynomials when ARX model structures (with zero input delays) were tested on the raw data set.

Figure 3.4 Plot of successful classification rates against orders of B and C polynomials when MAX model structures (with zero input delays) were tested on the raw data set.

Figure 3.5 A plot of percentage success versus training and testing day for an FIR model of order 3, using ‘fractional difference’ (not normalised) pre-processed data.

Figure 3.6 A plot of percentage success versus training and testing day for an FIR model of order 12, using the raw dynamical data.

Figure 3.7 Plot of percentage successful classification versus testing day for an FIR model of order 12 for dynamic data. The model was trained on data from day 12 of the experiment and tested on subsequent days.

Figure 4.1 Diagrammatic representation of the electronic nose modelling problem.

Figure 4.2 Plot of the data from the CellFacts instrument for the growth phase identification experiment. The upper plot shows the general increase in biomass (cell counts) with time. The lower plot shows the variation in mean size of the bacteria cells with time. The four growth phases (lag, log, stationary and late stationary), are labelled I to IV in each plot.

Figure 4.3 Plot showing the considerable long term variations in the output voltage from a single sensor over the whole (40 day) period of the growth phase experiment. There is clear evidence of diurnal peaks probably associated with the temperature of the biomass. The temperature of the room fluctuated daily.

Figure 4.4 The upper plot shows a section of the output data from a single sensor for the growth phase experiment. The lower plot shows the same section after dynamic pre-processing to shift each response cycle and thus remove some of the effects of the baseline drift evident in the upper plot.

Figure 4.5 The upper plot shows the output of sensor 1 over the whole period of the experiment. Significant low frequency drift is observable (mostly due to daily

temperature fluctuations). The lower plot shows the sensor 1 data after being passed through a high pass filter with cut-off frequency 60 μ Hz.

Figure 4.6 The upper plot shows the output of sensor 1 over a short period of the experiment. Some high frequency noise is observable. The lower plot shows the sensor 1 data after being passed through a low pass filter with cut-off frequency 7 mHz.

Figure 5.1 A schematic of the main components of the electronic nose system.

Figure 5.2 A photograph of the nose system in operation.

Figure 5.3 A schematic of the gas delivery system.

Figure 5.4 A photograph of the gas delivery system.

Figure 5.5 The system control software in operation.

Figure 5.6 Plot showing the response of sensor 1 to repeated inputs of acetone and isopropyl alcohol over a 13 day period.

Figure 5.7 Plot showing the response of sensor 1 to exposure to first acetone, then isopropyl alcohol.

Figure 6.1 Diagram showing the four main components of a (physical) forward model of an electronic nose system.

Figure 6.2 Plots to show the vital differences between sampling from an (ideal) static headspace (plots (a), (c) and (e)) and sampling from a headspace which is significantly altered by the outflow of gases initiated by the sampling process itself (plots (b), (d) and (f)).

Figure 6.3 Schematic of the gas system of the electronic nose.

Figure 6.4 Plot showing the relationship between the headspace flow rate \dot{Q}_h and the observed time delay in the odour transport system.

Figure 6.5 Plot showing numerical solutions of the diffusion equation at various times. The plots correspond to odour concentration profiles at the odour ‘front’ in the pipework, evolving from a sharp ‘plug flow’ situation at $t = 0$ to a shallower odour gradient at the front at later times.

Figure 7.1 Diagram illustrating the model assumption with regards to the partitioning of the sensor chamber into a main region and a local region around the sensor.

Figure 8.1 A plot of the output of sensor 1 over the entire duration of the experiment, showing the long term variation in the responses of the sensor.

Figure 8.2 A typical response of sensor 2 to a single exposure cycle. The first response is to an input of acetone, the second to isopropyl alcohol. The dotted line shows when the valves were operated on the gas delivery system.

Figure 8.3 A response of sensor 2 to an input of acetone. The flatness of the saturation level shows the concentration of the analyte in the headspace of the sample container was not significantly degraded over the sixty seconds for which it was sampled.

Figure 8.4 The ‘on’ portion of a response of sensor 1 to diluted acetone; x’s are the real data points, the solid line relates to simulated data using the model of Chapter 7. The r.m.s. error of the curve fit was approximately 4.58×10^{-3} V.

Figure 8.5 Simulations produced by simultaneous parameter estimations to the response of sensor 1 to (a) an input of acetone and (b) an input of isopropyl alcohol. The parameters corresponding to the nose system and not the analyte were kept equal for the estimation procedure. The combined r.m.s. error of the two curve fits was approximately 0.175 V.

Figure 8.6 Simulations produced by parameter estimations to the response of sensor 1 to (a) an input of acetone and (b) an input of isopropyl alcohol. The parameters

corresponding to the e-nose system and not the analyte were constrained to lie within the same narrow bounds for each curve fit. The combined r.m.s. error of the two curve fits was approximately 0.0174 V.

Figure 8.7 The variation in the fitted values for k_2 (top), r_{an} (middle), and the associated r.m.s. errors (bottom) for curve fits to responses of sensor 1 over the duration of the experiment. The solid lines represent data corresponding to acetone responses, the dashed lines represent data corresponding to isopropyl alcohol responses.

Figure 8.8 Plot of the fitted values of k_2 against those for r_{an} for (a) acetone response, and (b) isopropyl alcohol responses, all from sensor 1.

Figure 8.9 Plot of ambient temperature versus (a) k_2 , and (b) r_{an} for responses of sensor 1 to acetone.

Figure 8.10 Plots of (a) fitted values of k_2 , and (b) temperature compensated fitted values for k_2 ($k_{2,0}$ as described in equation (8.1)). The x's represent values extracted from acetone responses and the o's represent values extracted from isopropyl alcohol responses, all from sensor 1. Note the reduced within-group variation in the values of k_2 for plot (b) compared with plot (a).

Figure 8.11 Plots of (a) k_2 and (b) temperature compensated k_2 , against response number, showing the slightly improved separation between the two classes after temperature compensation. The solid lines represent data corresponding to acetone responses, the dashed lines represent data corresponding to isopropyl alcohol responses, all from sensor 1.

Figure 8.12 Plots of pre-processed static data from sensor 2: (a) static response, (b) static difference, (c) fractional difference, and (d) relative difference. The solid lines represent data corresponding to acetone responses, the dashed lines represent data corresponding to isopropyl alcohol responses.

Figure 9.1 The output of sensor 2 over the whole experiment, showing the diurnal and longer term variations observed.

Figure 9.2 A response of sensor 3 to exposure to the headspace gas from toxic bacteria sample.

Figure 9.3 The response of sensor 3 to an exposure to the headspace gases from a sample of toxic bacteria sample, together with a model fit to the data. The data points are plotted as x's, the model fit as a solid line. The r.m.s. error for the fit was 8.32453×10^{-3} V.

Figure 9.4 Plot of fitted values for k_2 (top) and the associated r.m.s. errors for curve fits (bottom) to responses of sensor 3 over the duration of the experiment. The solid lines represent data corresponding to responses to the headspace of toxic bacteria, the dashed lines to non-toxic bacteria responses.

Figure 9.5 Plots of (a) k_2 and (b) temperature compensated k_2 , against response number, showing the slightly improved separation between the two classes after temperature compensation. The solid lines represent sensor responses to toxic bacteria exposure, the dashed lines represent data sensor responses to non-toxic bacteria, all from sensor 3.

Figure 9.6 Plot of data from the CellFacts instrument for the growth phase identification experiment. The upper plot shows the general increase in biomass (cell counts) with time. The lower plot shows the variation in mean size of the bacteria cells with time. The four growth phases (lag, log, stationary and late stationary), are labelled I to IV in each plot, respectively.

Figure 9.7 The output of sensor 3 over the whole growth phase identification experiment.

Figure 9.8 The response of sensor 6 to an exposure to the second toxic cyanobacteria sample, together with a model fit to the data. The data points are plotted as x's, the model fit as a solid line. The r.m.s. error for the fit was 6.91129×10^{-3} V.

Figure 9.9 Plot of fitted values for k_2 (top) and the associated r.m.s. errors (bottom) for curve fits to responses of sensor 2 over the duration of the experiment.

Acknowledgements

I would like to acknowledge the financial support of the EPSRC, for their funding of this research project.

I also wish to thank my supervisors, Professor Julian W. Gardner, and Doctor Michael J. Chappell, of the University of Warwick, for their help and guidance throughout my time in the School of Engineering. Without access to their complementary skills and expertise, this project could not have been successful.

More personally, I wish to thank my parents, Colin and Menna Searle. They're great. I love them. Thanks for everything Mum and Dad.

Declaration

I declare that this thesis is entirely my own work.

As detailed in the text, the experimental data used for analysis in Chapters 3, 4 and 9 was gathered by Dr Hyun Woo Shin.

The design and construction of the experimental test system used to gather the data for Chapter 8, was by Professor Julian Gardner, Dr James Covington, and Ian Griffiths. The experiments to gather the data were designed and performed by me.

Parts of the work detailed in Chapter 3 were presented in the following conference papers:

Searle, G. E., Gardner, J. W., Chappell, M. J., Godfrey, K. R. and Chapman, M. J. (2000). 'System identification of electronic nose data for monitoring cyanobacteria in potable water: black box modelling.' *SYSID 2000 Symposium on System Identification*, Santa Barbara, California, 21-23 June 2000, pp. 145-150.

Gardner, J. W., Dow, C. S., Shin, H. W., Searle, G. E. and Chappell, M. J. (1999). 'Dynamical signal processing techniques for bioprocess monitoring.' *The 6th International Symposium, Olfaction and Electronic Noses 99*, Tubingen, 20-22 September 1999, pp. 331-335.

Parts of the work detailed in Chapters 3 and 4 were published in the following journal paper:

Searle, G. E., Gardner, J. W., Chappell, M. J., Godfrey, K. R. and Chapman, M. J. (2002). 'System identification of electronic nose data from cyanobacteria experiments.' *IEEE Sensors Journal*, **2**, 218-229.

This thesis has not been submitted for a degree at any other university.

Abstract

This thesis details research into the modelling of the dynamic responses of electronic nose systems to odour inputs. Most electronic nose systems contain an array of between 4 and 32 odour sensors, each of which respond in varying degrees to a range of different gaseous stimuli. In almost all electronic nose systems in use today, the steady-state responses of the odour sensors are extracted and passed to one of a variety of pattern recognition systems. The primary aim of this thesis is to investigate the use of information contained within the dynamic portion of the sensor response for odour classification.

System identification techniques using linear time-invariant black box models are applied to both extracted steady state and full dynamic data sets collected from experiments designed to assess the ability of an electronic nose system to discriminate between the strain and growth phases of samples of cyanobacteria (blue-green algae). The results obtained are compared with those obtained elsewhere using the same data, analysed with nonlinear artificial neural networks.

A physical model for the electrochemical mechanisms resulting in the measured responses is translated into a mathematical model. This model consists of a system of coupled nonlinear ordinary differential equations. The model is analysed, and the theoretical structural identifiability of the model is investigated and established.

The parametric model is then fitted to data collected from experiments with simple (single chemical species) odours. An odour discrimination method is developed, based upon the extraction of physically significant parameters from experimental data. This technique is evaluated and compared with the previously explored black box modelling techniques.

The discrimination technique is then extended to the analysis of complex odours, again using the cyanobacteria data sets. Successful classification rates are compared with those obtained earlier in the thesis, and elsewhere with neural networks applied to steady state data.

Abbreviations

ANN	artificial neural networks
ARMAX	auto-regressive moving average with exogenous inputs
ARTMAP	adaptive resonance theory map
ARX	auto-regressive with exogenous inputs
BAW	bulk acoustic wave
BP	back-propagation
CP	conducting polymer
CRC	controllability rank criterion
DFA	discriminant function analysis
DWT	discrete wavelet transform
FFT	fast Fourier transform
FIR	finite impulse response
HCA	hierarchical cluster analysis
KNN	k-nearest neighbour
LVQ	learning vector quantisation
MAX	moving average with exogenous inputs
MIMO	multi-input multi-output
MISCAP	metal-insulator-semiconductor capacitor
MISFET	metal-insulator-semiconductor field effect transistor
MISO	multi-input single-output
MLP	multilayer perceptron
MLR	multiple linear regression
MOS	metal oxide semiconductor
ODE	ordinary differential equation
ORC	observability rank criterion
PC	personal computer
PCA	principal components analysis
PIC	programmable integrated circuit

PLS	partial least squares
PTFE	polytetrafluoroethylene
QCM	quartz crystal microbalance
QMB	quartz crystal microbalance
RAM	random access memory
SAW	surface acoustic wave
SIMCA	soft independent modelling of class analogies

Chapter 1

Introduction

The five primary senses of sight, smell, hearing, touch and taste constitute the mechanisms by which we receive information from our surroundings. Without them we would be unable to interact with anyone or anything. Throughout history, scientific advancement has been driven by man's desire to describe and understand the world around him, via the information carried to him via these senses.

The information conveyed in optical or auditory stimuli is relatively easily understood, inasmuch as it is comprised of a pattern of frequencies of either light or sound waves. It is indeed possible to record and reproduce such stimuli. This is not the case for the chemical stimuli which evoke olfactory responses. Whilst it is true that such stimuli are essentially just mixtures of different gases present in the air which we are breathing, it is a highly non-trivial task to objectively and completely describe them. A typical odour to which we might be exposed may be made up of tens, hundreds, or thousands of different chemical compounds (Gardner and Bartlett 1999, p.35). In theory, a given odour could be completely described by the vector of concentrations of all of its chemical constituents, but this is not practical, since there are hundreds of thousands of compounds which are capable of producing an olfactory response, some of which can do so at extremely low concentrations.

The information conveyed to us via our sense of smell does not come to us in the form of a complete chemical analysis of the compounds present in the air around us, nor does it need to be. Mammals use their noses to identify certain odours which are important to them – in finding food, in identifying food which has gone off or is poisonous, in finding a mate, in identifying individuals. These odours are not sharply and precisely defined, rather they are characterised by a pattern of concentrations (or relative concentrations) of a set of compounds, which may have many elements. In seeking to describe an odour, there are thus two approaches which may be used – an analysis which breaks down the gas mixture into each of its chemical constituents and describes it in terms of concentrations for each, or an analysis which more closely mimics the mammalian system, whereby the odour is described by an expression of a much smaller number of key features.

The human sense of smell has been studied quite extensively, and attempts have been made to break down olfactory stimuli into a small number of ‘primary odours’. Typically the number of primary odours has been set at between five and seven (Amoore 1977; Murphy 1987), the idea being that any odour can be described by some combination of intensities of these primary odours. This description of olfactory stimuli in terms of (independent) primary odours is based on our perception of the odours, rather than corresponding directly to the manner in which odours are sensed in the biological system. In the mammalian olfactory system, there are of the order of a hundred million olfactory receptor cells, of about a thousand different types, located high within the nose (Gardner and Bartlett 1999, pp. 37-38). The fact that we can detect more than a thousand different odours suggests that each type of receptor cell responds to more than one different odour. Thus the receptors (or mammalian sensors) are not independent, instead they have overlapping sensitivities to different odours. The responses of these receptor cells are transmitted as electrical signals, along the axon of the receptor cell, to the olfactory bulb situated within the brain. Here the first level of sensory processing occurs, where the signals from each example of a given type of receptor cell are combined into just a few glomeruli. From here, the information is transmitted to the olfactory cortex and higher brain (Nagle *et al.* 1998). Researchers in the field of artificial olfaction have, in recent years, attempted to mimic the mammalian system by using a relatively small number of different sensors, each of which responds in varying amounts to different chemical gas stimuli. A great deal of

research has been conducted into both the sensor hardware technology, and the processing of the data produced by these sensors. A summary of the hardware and signal processing technologies used to date is presented in the following chapter.

1.1 Aims of this thesis

This thesis describes the work of the author in attempting to improve the performance of an electronic nose system based on one particular type of sensor hardware technology: tin oxide sensors. As will be discussed in the following chapter, the majority of signal processing techniques currently employed (across all types of gas sensors) make use only of the steady state responses of the sensors used. It has been postulated, and indeed demonstrated (Vilanova *et al.* 1996; Wilson and DeWeerth 1995; Hines *et al.* 1999; Gutierrez-Osuna *et al.* 1999; Nakamura *et al.* 1994; Llobet *et al.* 2001a), that discarding the information contained within the dynamic responses of the sensors is not necessarily wise. Hence, the work detailed in this thesis focuses on the development and evaluation of odour classification methods based on models for the dynamic responses of tin oxide sensors to various stimuli (both simple and complex). The performances of the techniques developed here are also compared with existing techniques based on static information only.

1.2 Summary of thesis

In the following chapter, background material relevant to the investigations detailed in this thesis is explored and discussed. In Chapters three and four, very simple linear black box modelling techniques are applied to data from electronic nose experiments involving biological samples that produce complex odours. The experiments were performed by Dr. H. W. Shin at the University of Warwick. The black box modelling techniques employed have previously been applied to inverse modelling of steady state electronic nose data (Marco *et al.* 1996), but here they are also employed for odour discrimination based on the full dynamic responses of the sensors in an array. The classification performance is also compared with the results obtained elsewhere (Shin *et al.* 2000) with the same data, using nonlinear neural network techniques.

Chapter five details the hardware and data-sampling software of an experimental test system constructed in order to facilitate the gathering of high quality dynamical data from simple odours. The majority of the hardware was designed and constructed by Prof. Julian Gardner, Mr. Ian Griffiths, and Dr. James Covington, all of the University of Warwick, with minor modifications made by the author. The software to control the system and record the data was written almost entirely by the author, with some subroutines provided by Dr. James Covington.

Chapter six describes the general modelling approach to be utilised in the subsequent parametric modelling of the electronic nose system. A discussion of the modelling of each of the critical components which make up the system is provided. The practical and theoretical issues of using (forward) models for the system in order to classify odours are explored.

In Chapter seven, a physical model for the interaction between a simple odour (comprising a single chemical species) is described, and translated into a system of ordinary differential equations. The system of equations is reduced, and analysed in terms of its steady states, and typical initial conditions. Subsequent to a suitable reparameterisation, the identifiability of the model is established.

Chapter eight details the development of techniques to utilise the model described in chapter seven for odour classification. Computational numerical methods are used to fit the parameters in the model to the responses of tin oxide sensors in the experimental test system to simple odours. The values of fitted parameters are used to identify the input odours.

In Chapter nine the techniques developed in chapter eight are applied to the biological data previously used in chapters three and four. A comparison is thus made between the simple black box models of Chapters three and four, the results reported elsewhere on the same data using nonlinear artificial neural networks, and the results obtained using the newly-developed physical parameter extraction technique.

Finally, in Chapter ten, the results and findings of the previous chapters are summarised and discussed.

Chapter 2

Electronic nose system technology

Although the idea of developing artificial olfactory systems was suggested as early as 1920 (Zwaardemaker and Hogewind 1920), it was not until the middle of the twentieth century that actual working instruments were reported. The concept of an ‘electronic nose’ was not formally defined until much later, when in 1994, the following definition was proposed, and came to be accepted (Gardner and Bartlett 1994):

“An electronic nose is an instrument, which comprises an array of electronic chemical sensors with partial specificity and an appropriate pattern-recognition system, capable of recognising simple or complex odours.”

So an electronic nose system is comprised of hardware which converts odour inputs to electronic signals, and software which converts those signals into an odour classification. The hardware essentially consists of an array of electrochemical odour sensors with overlapping sensitivities, which respond to a variety of target odours. Aside from some electronic filters to reduce signal noise, these sensor responses are passed directly to some form of software signal processing system. This signal processing system may take many forms. The dynamic sensor output is generally fed through some form of pre-processing, to extract either steady state or dynamic information from the response. This extracted information is then passed to some form of pattern recognition system that attempts to classify the odour input based on the

information from each sensor in the array, combined with some past information (e.g. calibration information or parameters calculated from some training set).

In this chapter, the history of electronic noses is presented, and a selection of the different hardware systems and signal processing techniques used to date are introduced and discussed. Some examples of current and proposed applications of electronic nose instruments are also given.

2.1 The historical development of electronic noses

As mentioned above, the development of artificial olfaction dates back to 1920, when it was suggested by Zwaardemaker and Hogewind that odours might be detected via monitoring the electrical charge on a fine spray of water containing the dissolved odourant substance (Zwaardemaker and Hogewind 1920). However, they were unable to produce a working instrument based on this principle. It was much later that the development of the first operational instruments was reported. During the period from the mid-fifties to the mid-sixties, several groups began to achieve some success, using a number of different systems. Moncrieff's approach, the details of which were published in 1961, used thermistors coated in a variety of different substances to non-specifically detect different odours (Moncrieff 1961). At about the same time, Wilkens and Hartman were developing a system using an array of true electrochemical sensors, making use of redox reactions at an electrode in contact with a porous rod saturated with an electrolyte (Hartman 1954; Wilkens and Hartman 1964). It was discovered in 1953 that the adsorption of gases on the surface of a semiconductor can produce measurable changes in the electrical resistance of the material, though this was not used in artificial olfaction systems until later (Brattain and Bardeen 1953). In 1965, two other groups published work detailing two other methods of detecting odours. One made use of conductivity modulation (Buck *et al.* 1965), the other of modulation of the contact potentials of certain materials in the presence of odours (Dravieks and Trotter. 1965).

Much later, in the 1980s, interest in the field increased and the idea of electronic nose systems, as they are understood now, blossomed. Publications from researchers at The University of Warwick (Gardner and Bartlett 1991) and from a group in Japan (Madou and Morrison 1989; Grate and Frye 1996) heralded the beginning of renewed research activity around the world. The first conference dedicated to electronic noses was held in Iceland in 1991, and resulted in the first book on the subject (Grate *et al.* 1995). Though the artificial olfaction systems developed had by then been colloquially referred to as electronic noses for some time, it was not until 1994 that the definition given at the start of this chapter was suggested and adopted (Gardner and Bartlett 1994). Worldwide there are now many academic research groups and at least a dozen commercial companies active in the field, with a growing number of instruments in use in a wide range of applications. The hardware and software technologies that have been developed in this time, and the applications that have been explored, are detailed in the following sections.

2.2 Electronic nose hardware

The most important components in an electronic nose system must be the odour sensors themselves, which may take many forms; some of which are described in the following sections. However, regardless of the particular type of sensor, in order to detect and recognise an odour there must be hardware to convert the changes which take place in the sensing material into a form which can be processed by computer software. A simplified diagram of the hardware of a typical electronic nose system is given in Figure 2.1 below.

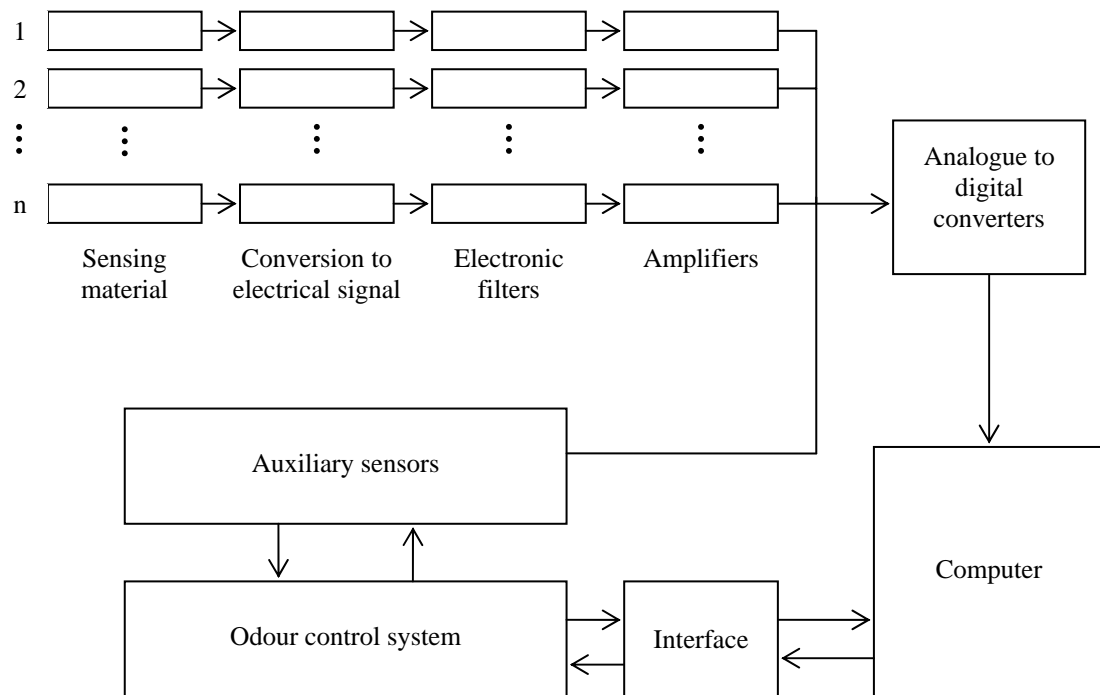


Figure 2.1 The hardware components of a typical electronic nose system.

The first stage in the interface between the sensing material of a sensor and the computer software is the process of detecting the physical or electrochemical changes that occur in the sensing material and converting this to an electrical signal. This can be a non-trivial process, especially given the fact that the changes produced in the sensing materials are often very small. This signal must reflect the physical or electrochemical properties of the sensing material which are altered by the presence of odours. The details of this stage are very specific to the sensor type, so are discussed appropriately in the subsections below. Generally, once the response of the sensor to its chemical surroundings has been converted to an electrical signal, this signal is passed through filters and amplifiers, before undergoing conversion to a digital signal and being processed by the software in a computer.

In addition to the odour sensors, electronic nose systems frequently contain other sensors to monitor important factors such as the temperatures in various parts of the system, the humidity in the carrier gas, the gas flow rates in the odour delivery system, and maybe even the mass flow rates in the gas system for some highly controlled

systems. The outputs of these sensors will also be fed through analogue to digital converters to the processing computer.

Most electronic noses contain, in addition to their sensor systems, hardware to control the odours to which the sensors are exposed. The complexity of this hardware varies greatly. The simplest might be essentially just a manually controlled air pump which sucks air containing the odourant over the sensors, perhaps for a handheld nose unit. The most complex would comprise many components, integrated with auxiliary sensors, in order to tightly control the chemical composition of the input to the sensor array. Such a system might contain strict temperature control of the samples, mass flow control of the carrier gas (which might well be nitrogen or artificial air from a tank), humidity and temperature control of the carrier gas, and mass flow control of the headspace of the samples. Such a system would be capable of ensuring a great deal of control and stability of the interfering factors (i.e. those factors other than the differences between odour samples) which might influence the measured responses of the odour sensors.

2.2.1 Metal oxide semiconductor sensors

Metal oxide semiconductor (MOS) chemoresistive sensors are probably the most commonly-used odour sensors in electronic nose systems. The majority of sensors of this type are based on tin dioxide, which is an *n*-type semiconductor. Variations in the sensitivity of the sensing material to different odourants are achieved through doping the material with catalytic metals such as palladium or platinum (Gardner and Bartlett 1999, p. 74). The use of tin oxide-based sensors was investigated more than 20 years ago, for simple, specific applications (such as the detection of carbon monoxide (Windischmann *et al.* 1979)), and the technology has remained popular as the field of electronic noses has grown. There are a number of reasons for this popularity. A range of tin oxide sensors has been commercially available for over ten years, making the construction of an array with a wide range of sensitivities relatively cheap and simple. The electronic circuitry required to operate the sensors is also relatively cheap and easy to construct, and the sensors themselves are very sensitive to a wide range of organic

compounds. The basic configuration of a Taguchi type tin dioxide gas sensor, such as those made by Figaro Engineering Inc. of Japan, is shown in Figure 2.2.




Figure 2.2 Configuration of a basic tin dioxide gas sensor.

The physical structure and operation of the sensor is essentially very simple. The conductivity of the metal oxide sensor varies when it is exposed to different odours. An electrode at each end of the sensing material is connected to electronic circuitry which measures the conductivity of the material. Aside from the sensing material and these electrodes, the only other significant part of the sensor is an integral heater coil. This is required since the sensors must be operated at elevated temperatures – normally between 250°C and 550°C (Madou and Morrison 1989, p. 83; Gardner and Bartlett 1999 p. 77). This is for two reasons – firstly, at temperatures below 100°C water is absorbed onto the surface of the sensing material and greatly affects the ability of the sensor to detect odours. Secondly, the sensitivity of the sensor to target odours is heavily temperature-dependent, so the temperature is maintained at a high level to ensure maximum sensitivity. However, this temperature requirement is the reason for one of the main drawbacks of this type of sensor – a high power consumption (typically around 1W for a Taguchi type sensor).

The mechanisms by which exposure to odours produces changes in the conductivity of MOS sensors depend upon the exact type of sensor, and the nature of the analyte species. The mechanisms have been studied in some detail and the electrochemical

reactions that can occur for a wide range of analytes are known (Madou and Morrison 1989; Williams 1991; Windischmann *et al.* 1979; Strässler and Reis 1983). The basic reaction mechanism (seemingly valid for simple combustible gases such as methane or carbon monoxide) has two stages:

- ionisation of oxygen and binding of ions to the sensor material, and
- combustion of analyte, leading to liberation of electrons, and resulting in increased conductivity.

The level of ionisation of the oxygen species depends upon the operating conditions (temperature, pressure etc.). Also, depending upon the analyte in question (and again the operating conditions), there may be subsequent stages where the combustion product undergoes further reactions. A particular example of the mechanism above is used to form the basis of a mathematical model for MOS sensor dynamics in Chapter 7.

The relationship between the change in the number of free electrons in the sensor material and the resulting change in the sensor conductivity is not simple. Were the sensor made from a single crystal, the change in the number of free electrons would directly produce the change in conductivity observed. However, the majority of commercial devices consist of a sintered polycrystalline material, in which the mobility of the charge carriers is crucial to the conductivity. The conductivity in such a device is given by:

$$\sigma = \mu_n en', \quad (2.1)$$

where σ is the conductivity, μ_n is the mobility of the electron, e is the charge on an electron and n' is the concentration of charge carriers (Gardner and Bartlett 1999, p. 75). In a sintered polycrystalline device, the observed change in conductivity is too great to be explained by the change in n' alone, and instead is attributed to large changes in μ_n with changes in n' . This change in mobility is explained by a lowering of the potential barrier height between adjacent grains in the polycrystalline material, making hopping between grains significantly easier and thus the overall electron mobility (and thus conductivity of the device) greater.

The potential barrier is dependent on the operating temperature, as are the reaction rates in the mechanism for the modulation of the number of charge carriers. This explains the strong temperature dependence of the observed sensor responses. MOS sensors are also influenced by changes in the humidity of the carrier gas, and can be 'poisoned' by irreversible or slowly reversible reactions with various species. This poisoning can lead to what is known as 'long term drift' whereby the response of a sensor is not exactly reproducible over long periods of time. In some applications, permanent poisoning is possible, in others the effect is more akin to a memory effect as the poisoning species are slowly desorbed from the sensor material. Regardless of the mechanism causing the drift, it can lead to a need for frequent recalibration or a need for a signal processing technique which produces classifications which are as independent of drift as possible, or which are capable of automatic recalibration.

The experimentally observed dynamic responses of MOS sensors in fixed operating conditions are generally well approximated by first or second order models. Examples of these dynamic responses can be seen in section 2.3.1, and in most succeeding chapters. The relationship between the analyte concentration and steady-state sensor response is generally considered to be a power law (so-called Freundlich isotherm), although this can sometimes be successfully approximated by a linear relationship over a suitable range of concentrations (Windischmann *et al.* 1979; Gardner and Bartlett 1999, p. 76).

More recently, micro-machined hotplate sensors have been produced, using a much smaller amount of tin oxide sensing material. These sensors overcome many of the traditional disadvantages of MOS sensors, since they have much lower power consumptions (of the order of mW) and low production costs since they can be batch fabricated (Al-Khalifa *et al.* 2001).

2.2.2 Conducting polymer sensors

Conducting polymer (CP) sensors are broadly similar in construction to MOS sensors, the crucial difference being that the sensing material is made from an organic polymer,

rather than a metal oxide semiconductor. The use of conducting polymers in odour sensors has been investigated since at least as far back as the early 1980s (Persaud and Pelosi 1991). Like MOS sensors, the presence of target odours is monitored with CP sensors via changes in the electrical conductivity of the sensing material. However, the mechanisms by which the conductivity of the sensor material is changed are different. In fact, the mechanisms by which the conductivity is altered are still not entirely known. There have been many candidate explanations, however it seems most likely that the true answer is some combination of the five possible mechanisms listed below (Gardner and Bartlett 1999, p. 84):

- The analyte molecule could affect charge transfer between the polymer and the electrode.
- The analyte might lead to oxidation or reduction of the polymer chain, thus imprisoning or liberating charge carriers.
- The analyte might affect the mobility of charge carriers along the polymer chains, through interactions with the charge carriers themselves.
- The analyte might interact with counterions within the polymer, and thus affect the mobility of the charge carriers along the polymer chains.
- The analyte might affect the mobility of charge carriers between polymer chains.

Although the relative importance of each of the mechanisms suggested above has not been positively established, empirical evidence suggests that the overall mechanism of the interaction behaves similarly to a simple binding action, where analyte molecules dynamically adsorb and desorb from a fixed number of binding sites in the polymer material. A simple mathematical model for a simple odour undergoing this kind of reversible binding produces first order response dynamics which fit well with the response curves observed, and Langmuir-type saturation curves for the steady state response versus odour concentration which again match very well with experimental data (Gardner and Bartlett 1999, p. 85).

Unlike metal oxide sensors, conducting polymer sensors operate at room temperature, thus they generally do not require a heater. This results in a much lower power consumption than MOS sensors. Conducting polymer sensors also typically have fast

response times (around two to 20 seconds). They can be fabricated fairly cheaply, and very small arrays of devices can be made (e.g. the array of 32 CP sensors used in the Cyranose R320 handheld unit made by Cyrano Sciences, Pasadena, California). These advantages make CP sensors popular for a range of applications, and the subject of much further research. However, the sensitivity of CP sensors is generally substantially inferior to that of MOS sensors, and they are currently prone to long-term drift, and are highly sensitive to the interfering influence of humidity variations. Thus they are far from the perfect sensors, but their utility is likely to continue increasing as advances are made in polymer technology and applicable signal processing techniques.

2.2.3 Acoustic wave sensors

Unlike MOS and CP sensors, acoustic wave sensors do not use changes in the electrochemical properties of an odour-absorbing material to detect the presence of analytes. Acoustic wave sensors consist of a piezoelectric substrate, such as quartz, which is coated with a material that absorbs the target odourous species. The presence of the target gas species is detected via the effect of the sorbed molecules on the propagation of acoustic waves (usually at frequencies between 1 and 500 MHz) through the piezoelectric material. Depending on the coating and analyte in question, the adsorption of the analyte into the sensing material may affect the acoustic properties of the sensor via a mass change of the sorbent coating, or via changes in the elastic properties of the coating, or indeed a combination of the two. The choice of coating is critical to operations of the sensor. A range of different coating materials have been used, from surface attached molecules (Grate *et al.* 1995), through monolayer films (Moriizumi 1998) to layers of polymeric films (Slater *et al.* 1992). The most commonly used coatings are ‘rubbery’ polymers, into which odours permeate more quickly than harder polymers. A range of different polymers can be used to produce a set of sensors with varying sensitivities to a range of target analytes. Indeed the choice of sensor can be made based upon knowledge of the solvation properties of the candidate polymer coatings and the target species.

There are two main types of acoustic wave sensors: those utilising surface acoustic waves (known as SAW devices), and those using bulk acoustic waves (BAW devices).

In SAW devices, a Rayleigh wave is set up on a plate of piezoelectric material coated with a sorbent material, and the propagation of this wave on the surface of the material is monitored, usually in terms of a shift in frequency from the initial wave to the received wave. There are two common configurations for SAW devices. In the first (known as a delay line SAW device), an a.c. voltage is applied to one pair of interdigitated electrodes on the surface of a plate of piezoelectric material. The surface wave propagates parallel to the fingers of the electrodes, across an area coated with sorbent material to a second pair of interdigitated electrodes, which receives the wave and transmits its properties (usually frequency or delay) via a suitable interface to the processing system. A diagram showing this sensor configuration is given in Figure 2.3. In the second SAW sensor format (known as a resonator SAW device), only a single pair of interdigitated electrodes is used to both transmit and receive the wave. Again the wave is transmitted perpendicular to the electrodes, across a sensing area, but then is reflected by a groove or ridge in the crystal surface, back across the sensing area to the electrodes (Gardner and Bartlett 1999, pp. 96-97; Grate and Frye 1996, pp. 40-45).




Figure 2.3 Basic configuration of a delay line surface acoustic wave device.

Bulk acoustic wave devices (often referred to as quartz crystal microbalances (QCM or QMB)) have a simpler construction. They essentially only consist of a slice of a single

crystal of quartz, typically around 1cm in diameter, with an electrode on either face, as illustrated in Figure 2.4.




Figure 2.4 Schematic showing the basic construction of a BAW odour sensor.

The required properties for the sorbent material which is coated onto one of the electrodes are the same as those for SAW devices. However, the method of detecting the changes in the properties of the coating, effected by the adsorption of the analyte into the material, is slightly different. In BAW devices, the presence of an analyte is detected via changes in the resonant frequency of the sensor.

Both SAW and BAW devices suffer from being strongly influenced by variations in temperature (mainly due to thermal expansion of the polymer coating) and humidity. They also have relatively slow response times (typically 100 to 1000s) and associated recovery times. SAW devices are generally five to ten times more sensitive to mass changes caused by the adsorption of analyte molecules, and are thus currently more popular for electronic nose research (Gardner and Bartlett 1999, pp. 93-98).

2.2.4 Field effect based sensors

Field effect devices which are sensitive to exposure to gases have existed since the early 1970s (Lundström *et al.* 1991). There are two main types: metal-insulator-semiconductor field effect transistors (MISFETs) and metal-insulator-semiconductor capacitors (MISCAPs) (Gardner and Bartlett 1999, pp. 99-99). More recently, research

has been conducted (successfully) into the use of conducting polymers, rather than metals, as the gate material (Covington *et al.* 2000). The basic principle of these sensors is that certain gaseous species can react with the gate material, and produce measurable changes in the properties of the transistor or capacitor (Madou and Morrison 1989; Gardner and Bartlett 1999).

2.2.5 Pellistors

Pellistors are sensors for detecting combustible (flammable) gases, and have been in production for over 35 years (Jones 1991). They measure the energy liberated when a combustible gas is oxidised. A typical pellistor consists of a platinum wire coil, surrounded by a bead of porous alumina, which contains precious metal catalysts. The wire coil is used both as a heater, to raise the temperature of the pellistor to approximately 500°C (Gardner and Bartlett 1999, p. 104), and as a resistance thermometer, to measure changes in the temperature of the coil that are brought about as a result of oxidation processes occurring within the sensor bead. Pellistors are fairly sensitive and have fast response times (typically around 20s). However their response is entirely non-specific (to different combustible gases) and they have a fairly high power consumption (around 350mW (Jones 1991)). They are not currently in widespread use in electronic nose systems, but current research may produce pellistors that are more suitable for this application.

2.3 Signal and data processing in electronic nose systems

The signals produced by the sensor array are generally passed, via filters and analogue to digital converters to a computer in order to be processed and transformed into information that is useful to the user. There is no single 'right' way to process the data from the array; the best choice depends upon the hardware used, the application considered, and the information required.

2.3.1 Pre-processing

The first stage in the software signal processing procedure is usually the extraction of the required information from the sensor responses, and after possibly transforming it some way, to collect it together in a suitable form to pass on to the pattern recognition system. This stage is known as ‘pre-processing’.

The majority of electronic nose systems available commercially and in industrial use make use of the static responses of the odour sensor arrays (i.e. the steady state values of the sensor outputs), ignoring the dynamic response of the sensor (Hines *et al.* 1999; Di Natale *et al.* 1995; Gutierrez-Osuna *et al.* 1999; Nakamura *et al.* 1994). A typical odour sensor output will have an essentially stable (over the short-term) baseline. When exposed to an odour, the output of the sensor will either rise or fall (depending on the choice of sensor and analyte), and then, assuming that the concentration of the odour remains stable for long enough, the sensor output will re-establish a new stable equilibrium value. A typical response of a metal oxide semiconductor (MOS) sensor is plotted in Figure 2.5, along with the baseline and static response outputs.




Figure 2.5 A typical response of a MOS odour sensor to a square pulse of acetone, x_b is the sensor baseline, and x_a the steady-state response.

The two steady state sensor outputs (in air or carrier gas and in odour) of the sensor are extracted from the dynamic sensor response, and may be manipulated in a number of ways before being passed to the pattern recognition system. Some commonly used pre-processing algorithms are given below (Gardner and Bartlett 1999):

- Absolute response: x_a
- Difference: $(x_a - x_b)$
- Relative difference: $\frac{x_a}{x_b}$
- Fractional difference: $\frac{(x_a - x_b)}{x_b}$.

Where x_b is the sensor output baseline, and x_a is the steady state sensor output when exposed to an analyte. Pre-processing of static information is discussed in more detail in Chapter 3. Often logarithmic versions of the above algorithms are also used in an attempt to linearise the response (Gardner 1991). These values are calculated for each sensor, and then used to produce a vector (one element per sensor in the array) that describes the static response of the array.

The choice of appropriate static pre-processing algorithm depends on a number of factors, not least of which are the sensors used and the likely range of operating conditions. For example, if the sensors exhibit a drift or temperature dependence of baseline, but not magnitude of response, then the ‘difference’ algorithm might be suitable. For many sensors though, the ‘fractional difference’ algorithm is thought to provide the best reduction in temperature-dependence (Gardner 1991).

Pre-processing for analysis of dynamic sensor response signals can take a number of forms. One approach is to simply extract the relevant portions of the dynamic responses, and pass them unchanged to a classification system that uses the full dynamic response. Variations on this include processing algorithms which are similar to those used to reduce the effects of temperature and drift, such as a ‘dynamic difference’ algorithm where the baseline is subtracted from each sampled data point over the response. Such techniques are used in Chapter 4. Another possibility is the use of filters which produce a discrete approximation to the derivative of the sensor output

at each time point, again to attempt to counteract the interfering effects of sensor drift and temperature influences. Alternative methods of pre-processing dynamic data mostly involve the extraction of parameters which characterise the dynamics of the response (e.g. the parameter estimation techniques developed in Chapter 8, or the estimation of a ‘constant’ in (Vilanova *et al.* 1996)).

In addition to each of the pre-processing methods described above (which act on the outputs of individual sensors), there may also be some form of array pre-processing incorporated into the signal processing system. This can take a number of forms, such as normalisation of the array output, where each element of the array response vector is scaled by the length of the vector, so that the response vector always has length one. Often the responses of each sensor are also transformed in order to remove the possibility of emphasis being placed on sensors due to the magnitudes or variances in their responses. For example, subsequent to being exposed to a large training set, a transformation might be established for each sensor in order to map the outputs of each to the range zero to one, and possibly to map the likely variances to be the same for each sensor.

2.3.2 Processing of steady-state data

As mentioned above, the majority of data processing techniques used for electronic nose systems operate on parameters extracted from the steady-state responses of the sensors in the array. In this section, a selection of the most popular techniques for modelling these steady-state responses, and converting them to odour classifications is introduced and discussed. Some of the techniques described below are directly applicable as classification tools, and some are more often used for visualisation and evaluation of the odour discrimination potential of the multi-dimensional data collected from nose systems. Some can be regarded as straightforward supervised learning techniques, where a training (or calibration) set is used in order to estimate parameters linking the odour inputs to the nose system with the resulting pre-processed sensor outputs (such as black-box models, and artificial neural network techniques). Others are unsupervised methods where no information regarding the classes to which different response vectors should belong is supplied, and the separation of the data into

classes is performed purely based on the response vectors themselves (e.g. cluster analysis or principal components analysis)(Gardner and Bartlett 1999).

2.3.2.1 Forward and inverse linear models

The simplest forward model for the static response of an array of sensors to an odour would be that each sensor responds linearly to each component of the odour. So if an odour was described by a vector c of the concentrations of m gas components, the response vector x of an array of n sensors would be given by

$$x = \mathbf{A}c, \quad (2.2)$$

where \mathbf{A} is an $n \times m$ array. Using this model, the coefficients of \mathbf{A} are calculated from calibration experiments (assuming sufficient experiments can be carried out with different combinations of the potential gas constituents). Thus this is a supervised learning technique. Once an estimate for \mathbf{A} is obtained, provided that the number of sensors n is greater than the number of potential components m , calculation of an estimate for the concentration vector of an unknown input odour is simply a matter of using linear algebra to produce a least squares solution to equation 2.2. This sort of model is attractive because of its simplicity and ease of implementation, and its ability to provide a quantitative estimate of the concentrations of each of the gas components. However, its effectiveness is clearly limited by the assumptions that the responses of the odour sensors to different gases are linear and independent. For most sensors, even after pre-processing, these assumptions are only valid at low concentrations and over a limited range of concentrations. There is also no feature of the model that can take into account any ‘memory’ of recent array exposures, so that the initial calibration is required to hold until later calibration experiments are performed. Thus the technique relies wholly on the pre-processing to provide counteraction of temperature or long-term drift effects. As mentioned above, the number of sensors in the array must be greater than or equal to the number of candidate components in the gas mixture. With current hardware, this restricts the technique to applications where the odours are composed of a small number of species, since the majority of instruments have arrays of 32 or less sensors. Complex odours, such as those produced by biological samples

or foods and beverages often contain many more species than this. In these circumstances it is more appropriate to use the more qualitative techniques described below, where the output is a selection from a small number of classes, rather than an estimate of the concentrations of each of the chemical species which might be present.

A more commonly-used variation on the above model for sensor responses is an inverse linear model. For this technique, linearity is assumed for an inverse model of the sensor responses. This means that the gas input concentrations are assumed to depend linearly on the sensor responses, i.e.

$$c = \mathbf{B}x, \quad (2.3)$$

where again c is the (m -dimensional) vector of gas component concentrations and x is the (n -dimensional) vector of (pre-processed) sensor outputs. \mathbf{B} is an $m \times n$ array. Once \mathbf{B} is obtained from calibration experiments, the process of converting a response vector to an estimate of the constituents of the input is trivial. The estimation of the coefficients of \mathbf{B} are usually obtained via multiple linear regression (MLR) or partial least squares (PLS) (Gardner and Bartlett 1999, pp. 144-145; Gutierrez-Osuna 2000). For details of MLR and PLS see, for example, Beebe *et al.* 1998. The use of this model is subject to essentially the same limitations as the forward linear response model.

2.3.2.2 Nonlinear response models

Since the (static) responses of odour sensors are generally nonlinear with respect to analyte concentration, the above linear models are only valid over fairly small concentration ranges. As mentioned in section 2.3.1, one method used to overcome this problem and extend the usefulness of the model is to use a pre-processing algorithm which attempts to linearise the responses (e.g. a logarithmic algorithm for sensors which respond via a power law). Another method is to extend the linear (forward or inverse) response models to different functional forms. This can be done either by simply extending to higher order polynomial response functions (Sundgren *et al.* 1990), or by assuming a response to be of a functional form such as a power law (Hierold and Muller 1989). Since the coefficients for the function(s) postulated for the

model must be calculated from a calibration experiment, this too is a supervised learning technique.

2.3.2.3 Black box models

Linear time-invariant black box models are not dissimilar in structure to the linear response models discussed in section 2.3.2.1. However, there are two crucial differences. The first is that they are often used (as inverse system models) to produce a classification of the odour input, rather than a quantitative estimate of the concentrations of each postulated gas component. The second is that instead of attempting to produce this classification based solely upon the response vector x of the sensor array for a given response, the structure of the models allows for previous response vectors (i.e. to previous odour exposures) and indeed classifications into account. This is again a supervised learning technique.

For a detailed description and examples of the use of black box models, see chapters 3 and 4, and Searle *et al.* 2002.

2.3.2.4 Cluster analysis

Cluster analysis is an unsupervised multivariate statistical technique, the object of which is to split a set of multivariate data into groups according to how ‘close’ they are in multivariate space. As such it is frequently used to assess the discrimination power of a nose system for a given application. If cluster analysis is performed on a test set of known odours, and the cluster analysis grouping of the data matches the desired set of classes then it can be asserted that the sensor array and pre-processing used possess the ability to discriminate between the classes (Gardner 1991).

The most commonly-used form of cluster analysis is known as hierarchical cluster analysis (HCA). Since cluster analysis is concerned with forming groups based on similarity of data vectors, a measure of distance between points in multi-dimensional space is required. This is known as the distance metric, and can be selected from a

number of options. Normally the data vectors are pre-processed in such a way that the size and variation along each dimension of the data set is approximately equal, so most commonly the n -dimensional Euclidean metric is used:

$$d(x, y) = \sqrt{\sum_{i=1}^n (x_i - y_i)^2}, \quad (2.4)$$

however, variations on this are sometimes used, especially where the data have not been extensively pre-processed. A way of measuring the distance between a point and a group of points is also required, as is a method of measuring the distance between two groups of points. These are usually defined in terms of the point-to-point distance metric, and can be based upon the nearest neighbour distances (known as ‘single linkage’), furthest neighbour distances (known as ‘complete linkage’), or many other variations including the use of ‘centres’ of groups to measure from (Beebe *et al.* 1998).

Given a choice of distance metric and linkage method, there are two extremes from which the HCA can start, and work towards the other. In one method the analysis begins by clustering the data so that every data point belongs to the same cluster, then successively splits the set into more and more groups until eventually each data point lies alone in a cluster separate from the other data points. In the other method, the analysis starts with p clusters (where p is the number of data points) and gradually increases the between-groups distance threshold, collecting point into clusters, until the distance threshold is so great that all data points lie in a single cluster. The HCA produces a graphical output of the natural clustering of the data, using the selected distance metric and linkage method, called a dendrogram. An example of a dendrogram is given in Figure 2.6. For examples of the application of HCA to electronic nose data, see Gardner 1991 and Aishima 1991a. A more thorough explanation of the technique of HCA is given in Beebe *et al.* 1998.




Figure 2.6 An example of a dendrogram produced by HCA, similar to those produced from a successful electronic nose experiment.

2.3.2.5 Principal components analysis

Principal components analysis (PCA) is another unsupervised pattern recognition technique, which is particularly suitable for data where the measured parameters are not independent, exactly as is the case for most arrays of odour sensors. With an array of n sensors, each (static) data point is an n -dimensional vector. For $n > 3$, such data are difficult to portray in a graphical manner that is easy for the user to interpret. However, since the outputs of the sensors are generally not uncorrelated, and the number of odour classes is often relatively small, the differences in the data points can often be seen by viewing the data in fewer than n dimensions. Clearly one way of doing this would be to simply plot the outputs of two or three sensors against each other, but a much more effective method is provided by PCA.

Principal components analysis transforms an n -dimensional data set to a different n -dimensional data set, in such a way that each data point of the new data set (the

principal components) is a linear combination of the old data set (the outputs of the n sensors in the array). The coefficients of the transformation (known as the ‘loads’) are chosen so that the first coordinate (first principal component) has maximal variance. Thus, the first principal component captures the most-varying facets amongst the data points. The second principal component is obtained by using loads which again maximise the variance in the resulting data, but with the added constraint that the second principal component must be uncorrelated with the first. This process is continued until the n -th principal component is calculated. If the sensors in the array were in fact selective and uncorrelated, then the PCA would be of no benefit, since the resulting data set would be equivalent to the original. However, with real electronic nose data, it is often found that 90% or more of the variance in the data is captured by the first two or three principal components (Gardner and Bartlett 1999). This suggests that in some sense the true dimensionality of the data set is very much less than n . In this case, if the PCA plot shows separation of the data into the desired class groups then all is well, and principal components ‘scores’ (the values of the transformed data coordinates) for responses to unknown odours may be used to classify those odours, by placing them on the PCA plot of a training set, and seeing which group of known odours they lie closest to. In this situation the PCA is being used effectively as a tool to produce a supervised rule-based classification technique. If good separation of the desired groups is not observed, then perhaps a different choice of sensors is required. An example of a PCA plot is given in Figure 2.7.




Figure 2.7 A typical PCA plot of the first and second principal components from an electronic nose experiment to discriminate between three different analytes. Although six MOS sensors were used in the experiment, adequate group separation is observed with just two plotted dimensions.

Unlike cluster analysis, PCA can show trends and patterns in the data such as long term drift (which might show up as a consistent drifting of the data points on the PCA plot with time), or dependence on temperature variations. Examples of the use of PCA to visualise and analyse electronic nose data are abundant in the literature (Marco *et al.* 1998; Gardner 1991; Llobet *et al.* 1999b; Shin *et al.* 2000). For a more thorough description of PCA, see Beebe *et al.* 1998.

2.3.2.6 K-nearest neighbours

K-nearest neighbours (KNN) is a supervised technique whereby a classification of an unknown data point is made based on which class the data point is nearest to in multidimensional space. Ordinarily the space in which this is carried out would be the space of pre-processed output sensor vectors, although it could instead be performed on a space of principal components scores, in order to reduce the computational load

and facilitate visualisation of the problem. As with HCA, a distance metric for the multidimensional space must be selected; the most common again being the Euclidean distance, although others (such as the Mahalanobis distance) are sometimes used (Beebe *et al.* 1998). When a data point resulting from sensor exposure to an unknown odour is presented, the decision on which class it should belong to is based, as the name suggests, on the classes of the data point's K nearest neighbours. Normally the class allocation is made in a democratic fashion – each of the data point's K nearest neighbours 'vote' that the data point belongs to their class, and the class with the most votes wins. The choice of how many neighbours to consider depends mainly on the nature of the training set. If the classes are well separated then a small value of K will suffice, however if the classes overlap then more care must be taken. Also, if there are not equal numbers of training vectors for each class, then care must be taken that the choice of K will not lead to a bias towards the class(es) with large numbers of training points (Beebe *et al.* 1998). An advantage of a well-implemented KNN classification system is that the 'voting' system can provide a quantitative estimate of the level of confidence in a classification (via the proportion of polled votes that were for the winning class), although this must be carefully handled to ensure that outliers aren't incorrectly placed in a given class with a high confidence, merely due to the fact that they are more like that class than the others, even though in fact they are nothing like any of the training classes. A disadvantage of KNN is the care required in its use (in terms of selecting K and parameters relevant to the prediction of classification confidence) and its susceptibility to sensor drift, making it likely that careful and frequent recalibration will be needed. KNN is described in more detail in Beebe *et al.* 1998.

2.3.2.7 Soft independent modelling of class analogies

Soft independent modelling of class analogies (SIMCA) is another supervised learning method, based on principal components analysis. The principal of SIMCA is simple and relatively easy to implement. A training set consisting of a number of example data points for each of a number of classes is required. For each class, PCA is performed on the data points from that class, independently of the data points from the other classes. An appropriate number d_{ci} of principal components is selected, and a

multi-dimensional box is drawn around the selected class in its PCA space. Note that the number of principal components used need not be the same for each class (Beebe *et al.* 1998).

Since the transformation from sensor space to principal components space is linear, these multidimensional boxes in PCA space correspond to scaled and rotated hypercubes in sensor space. Given a data vector from an unknown sample, the class membership of the data point is determined by a combination of the distance of the data point from the d_{ci} -dimensional subspace (of sensor space) defined by the principal component loadings, and the distance of the data point from the box boundaries defined within that subspace. In practice this is equivalent to projecting the data point separately into the PCA space of each of the classes. If the data point is projected to a point that is within a distance of some critical value (chosen based on the within-class variance of the training data) from the box boundary, then it is considered to be a member of that class.

Clearly using this method it is likely that the class boxes will not occupy all of sensor space, and it is possible that some class boxes will overlap. Thus an unknown data point may be placed in zero, one, or many classes. How these situations are dealt with must depend upon the application in question. A more comprehensive discussion of the SIMCA technique is presented in Beebe *et al.* 1998.

2.3.2.8 Discriminant function analysis

Discriminant function analysis (DFA) is a supervised pattern recognition technique. The aim of DFA is to find ‘discriminant functions’ of the (pre-processed) sensor outputs, each of which map the multidimensional sensor space to a single dimension, in such a way as to maximise the separation between the classes.

Linear discriminant function analysis has much in common with principal components analysis. In linear DFA, each discriminant function is formed simply as a linear combination of the sensor outputs. Usually the coefficients of the first discriminant function are selected to maximise the ratio of the inter-class variance to intra-class

variance. The coefficients of the second discriminant function are then chosen similarly, with the added constraint that the second discriminant function is uncorrelated with the first, and so on for subsequent discriminant functions (Gardner and Bartlett 1999). Thus linear DFA and PCA are essentially the same except that for PCA, the functions are chosen to maximise the spread of the whole data set (so no information about the known classes is used) whereas in DFA the functions are chosen specifically to maximise the spread between the known classes of training vectors. For examples of the application of linear DFA to electronic nose data, see Aishima 1991b and Romain *et al.* 2000.

Nonlinear discriminant functions have also been applied to electronic nose data (e.g. quadratic functions, (Mason 1996)), however the problem of calculating the parameters is potentially much more difficult. For a thorough description of discriminant function analysis, see Krzanowski 1988.

2.3.2.9 Artificial neural networks

Artificial neural networks (ANN) are structures which are based loosely upon the structure of the brain / nervous system. They are an attractive choice for the analysis of electronic nose data for a number of reasons. The electronic nose hardware is designed to mimic the biological olfactory ‘hardware’ and thus produces broadly similar outputs, so it seems appropriate to use a processing system that also mirrors the biological system (Gardner *et al.* 1990). Artificial neural networks can be chosen and designed so that they can handle the nonlinear nature of the sensor data, and some can adapt themselves to correct for drift in the outputs (Distante *et al.* 2000; Marco *et al.* 1998).

Over the last twenty years or so, neural networks have been researched extensively, and a great variety of structures have been developed. Many of these have been applied to electronic nose data (Hines *et al.* 1999). Essentially neural networks are structures with a set of inputs, corresponding in this context to the (pre-processed) sensor outputs, a set of outputs (normally one per odour class), and a series of connections between them, possibly via extra (hidden) layers of ‘neurones’, each of which is weighted. The

range of possible structures is great, and there are many methods of training the neural network, which in practice means calculating the parameters (weights) for the network. The most commonly-used artificial neural networks for electronic nose data include Hamming networks, Kohonen networks, learning vector quantisation (LVQ), multilayer perceptron (MLP) and back-propagation (BP) networks. Fuzzy methods have also been used with significant success, such as the Fuzzy ARTMAP (Shin *et al.* 2000). For details of the application of these neural network techniques to electronic nose data, see Gardner and Bartlett 1999, and for a more general overview of neural networks, see Haykin 1999. Artificial neural networks are currently the most popular tool for odour classification, often used in conjunction with PCA as a visualisation tool. Their popularity is due to the fact that when used carefully they frequently outperform other pattern recognition methods. However, their strengths are derived from their complexity (in terms of structure and numbers of parameters which define them), which is also their weakness, inasmuch as reliable performance may only be attained after a large training set is used. This is an unavoidable aspect of trying to describe a complex system using model structures that do not incorporate any knowledge of the mechanisms defining the system.

2.3.3 Processing of dynamic data

As mentioned in section 2.3.1, generally the pre-processing stage is used to extract information relating only to the steady-state responses of the sensor outputs. However, it has been shown that the dynamic response of odours can also carry important information (Di Natale *et al.* 1995; Gutierrez-Osuna *et al.* 1999; Nakamura *et al.* 1994). Dynamic information has been used in a number of ways to improve the performance of nose systems.

Some studies have reported that the dynamic response of the sensor array contains important discriminatory information which can be passed to the pattern recognition system in order to improve the classification performance (Gutierrez-Osuna *et al.* 1999). For slowly responding sensors, or quickly changing environments, the early part of the dynamic sensor response can be used to predict the final (steady-state) value, and thus reduce the time required for calibration and subsequent analysis of a

sample (Di Natale *et al.* 1995). A more recent development is the use of thermal modulation of a single sensor to produce a dynamic response that is analysed using wavelet transforms and neural networks. This technique potentially reduces the number of sensors required in an array.

Physical models for the conduction mechanisms of some types of odour sensors have been converted into mathematical models, and subsequently analysed, but these are not yet generally used to enable odour classification via dynamic responses. An example for MOS sensors is the reaction-diffusion partial differential equation model suggested in Gardner 1989. For CP sensors, a simple reversible binding model gives rise to a first order exponential response model, and a Langmuir isotherm steady-state response, though the time constant observed is still infrequently used.

2.3.3.1 Exponential response models

The pre-processed steady-state sensor outputs essentially capture the magnitudes of the sensor response to an odour stimulus. In order to describe the dynamics of a sensor response efficiently, it is natural to seek a parameter or set of parameters which capture(s) the key features of the shape or geometry of the sensor response. For many classes of odour sensors (e.g. MOS sensors), the response curves observed in experiments fit multi-exponential curves very well (Gutierrez-Osuna *et al.* 1999; Samitier *et al.* 1994). Frequently a first-order model produces acceptable fits, but more recently a bi-exponential response model has been shown to give a better approximation (Endres *et al.* 1995). The one or two (or more for higher order multi-exponential models) time constants are extracted from the dynamic responses of each sensor, and then odour classification is performed on these values, possibly in conjunction with the steady-state values. This technique is useful in that it can increase the separation between odours which produce differing dynamic characteristics, but similar steady-state values. Thus, for a given application, it can increase classification performance or decrease the number of sensors required. However, due to the lack of a firm physical basis for the response model, the effects of interfering influences such as temperature and humidity are ignored, or can only be empirically modelled. Thus in order to use such a response model effectively over a range of conditions, a highly

nontrivial calibration must be performed, possibly quite frequently (to allow for sensor poisoning or other long-term affects).

2.3.3.2 Thermal modulation

As mentioned in section 2.2.1, MOS sensors are operated at elevated temperatures, usually via an internal heating element. The conductance and odour-sensitivity of the sensor material depends quite strongly upon the temperature at which it is operated. This characteristic is common to most types of odour sensors, in varying degrees. However, the combination of a high temperature gradient between the sensor operating temperature and its surroundings with the developing technologies for producing very small sensors based on micro-machined hotplates mean that very rapid changes in the temperature of the sensing material can be achieved quite easily (compared with, say, CP sensors which are operated at approximately room temperature). It is possible to heat these devices to 500°C and then cool them back to room temperature in a few milliseconds (Al-Khalifa *et al.* 2001). The working temperature of the sensor is cycled (with a period typically of between a few seconds and a few minutes, depending mainly on the response characteristics of the sensor material / analytes) and the resistance of the sensor recorded. The heater driving signal is frequently sinusoidal (Llobet *et al.* 2001a), though slower triangular cycles have been used with more traditional Taguchi MOS sensors, with their somewhat slower thermal dynamics (Kohler *et al.* 1999).

The processing of the resulting dynamic data may be approached in a number of ways. In Kohler *et al.* 1999, the sensor outputs at each time point sampled (256 points were recorded for a single thermal cycle) were used to produce a 256-dimensional response space. The five calibration responses (for five different target analytes) were transformed to an orthonormal base, and unknown samples were directly back-transformed to give a quantitative estimate of the concentration of each of the five target analytes. An alternative method was employed by Llobet *et al.* 2001a and Al-Khalifa *et al.* 2001, whereby the sensor response to a sinusoidal thermal input was analysed using fast Fourier transforms (FFT) or discrete wavelet transforms (DWT). This produced good results, especially when combined with a 'Fuzzy ARTMAP'

artificial neural network that takes the Fourier or wavelet coefficients as inputs (Llobet *et al.* 2001b). This is still a fairly new method, so it remains to be seen how robust the classification performance will be to interfering influences and long-term drift. So far the method has been used on single sensors, which will clearly limit the number of odours that can be discriminated, but there is no reason why an array of sensors shouldn't be operated in this manner.

2.4 Current and future applications for electronic nose systems

The range of potential applications for electronic nose systems is truly vast. Any situation where the identification of gases or odours is required constitutes a potential application. To date the most common commercial applications have been in the food and beverage industries, but it has been suggested that electronic nose systems might find uses in many other fields. Examples of other potential areas of application include packaging quality control, perfumery, environmental monitoring, biotechnology, explosives / illegal drug detection at airports etc., early detection of biological or chemical weapon deployment in battlefields, medical diagnosis and quality control in the automotive industry. Table 2.1 lists some applications which have been reported in the literature, together with the types of nose systems used.

Area of application:	Aim of analysis:	Sensors used:	Reference:
Alcoholic beverages	Discrimination between different beverages	6 MOS	Aishima 1991a
Coffee	Discriminate between two varieties of bean	6 MOS	Aishima 1991b
Tomatoes	Assess quality of tomatoes, and discriminate between two cultivation methods	8 QMB	Di Natale <i>et al.</i> 1998
Milk	Detection of 'burned' aroma in UHT milk samples	8 QMB	Di Natale <i>et al.</i> 1998

Water quality	Discrimination of strain and growth phase of cyanobacteria	6 MOS	Gardner <i>et al.</i> 2000
Banana testing	Non-destructive ripeness determination	4 MOS	Llobet <i>et al.</i> 1999b
Cow's breath	Diagnosis of illness	6 MOS	Llobet <i>et al.</i> 1999a
Bacteria / fungal growth	Detection and classification	15 MOS	Schiffman <i>et al.</i> 2000
Olive oil	Classification of three olive oils	8 CP	Bargagna <i>et al.</i> 2000
Pollutant gases	Identification and concentration estimation	4 MOS	Reich <i>et al.</i> 2000
Eye infections	Detection of bacteria	32 CP	Boilot <i>et al.</i> 2000
Beer	Detection of taints	12 CP	Gardner <i>et al.</i> 1994
Fish	Assessment of freshness	4 MOS	Olaffson <i>et al.</i> 1992
Cheese	Maturity assessment	20 CP	Persaud and Travers 1996
Cow teats	Detecting hygiene condition prior to milking	20 CP	Mottram and Persaud 2000
Industrial chemicals	Detection and recognition of industrial chemicals	32 CP	Furlong and Stewart 2000
Automotive industry	Detect and classify sources of 'new car smell'	6 QMB	Garrigues <i>et al.</i> 2000
Explosives detection	Detection of military grade TNT	5 MOS	Pardo <i>et al.</i> 2000
Ham	Control of drying process	8 MOS	Horrillo <i>et al.</i> 2000

Table 2.1 Some examples of electronic nose applications.

As the technology matures, the selection of new applications for electronic nose systems is increasingly being driven by commercial interests. As the hardware and data processing technologies improve, the range of applications for which nose systems can be useful will expand. This will be due to an increase in the effectiveness of the systems when analysing complex odours (e.g. biomedical applications), and a decrease in the cost of producing commercial devices, for both simple and complex applications.

Chapter 3

Cyanobacteria strain identification using system identification techniques

In this chapter, black box modelling techniques from the field of system identification are employed to analyse data from a biological experiment. The aim of the experiment was to assess the ability of an electronic nose system to discriminate between two species of cyanobacteria (blue-green algae). Establishing such a capability would not only be directly useful (for environmental agencies or water companies) but would also help to provide an indication of the potential of electronic nose technology for applications in the biomedical arena, such as bioprocess monitoring and medical diagnosis.

As mentioned in Chapter 2, most of the well established techniques of analysing data gathered from electronic nose experiments involve the use of only the steady state (static) responses of the sensors. The data are ‘pre-processed’ to extract the steady state information from the dynamic data produced by the system. This steady state information is then used by a pattern recognition system to classify the odour. See Chapter 2 for more details of current techniques. It has been suggested that there could be significant discriminatory information contained within the traditionally discarded dynamic data (Hines *et al.* 1999; Wilson and DeWeerth 1995; Vilanova *et al.* 1996; Nakamura *et al.* 1994). This suggestion is investigated via a comparison of success rates obtained using similar black box model structures for both steady state (static)

data and dynamic data from the same experiment. The models used are linear time invariant black box models. Such models have been shown to be effective when applied to steady state data (Marco *et al.* 1996), but their application to dynamic data is novel. The success rates achieved are also compared with those obtained elsewhere (Shin *et al.* 2000) using neural network techniques applied to the same experimental data.

3.1 Black box modelling

The models considered here are, in effect, inverse black box models of the electronic nose system. A forward model of the system would have inputs corresponding to the odour input $y(t)$ to the system, and outputs corresponding to the electrical resistances of the sensors responding to those odours $u(t)$. Our models work in the opposite direction, taking the sensor resistances as inputs, and producing a classification of the odour as an output. See Figure 3.1 for a pictorial representation of this.




Figure 3.1 Diagrammatic representation of the electronic nose modelling problem.

Generally the model output could take many forms. For example, if the target odours were mixtures of many chemical components then a multicomponent model could have many outputs, each giving an estimated concentration of one of the components. However, for the qualitative problems considered here, the output required is a simple classification of the input odour.

For the experiment considered here, the bacteria must simply be classified as toxic or non-toxic. For an example of a more complex classification problem, see the following

chapter. These simple classifications can be obtained from black box model outputs in a number of ways, requiring the use of either single or multi-output black box models.

The *System Identification Toolbox* (Ljung 1995) within *Matlab* (Version 5) provides the user with many functions for creating, evaluating and using black box models. The general form for the discrete time multi-input, single-output (MISO) models considered is given below (Ljung 1987):

$$\begin{aligned}
 A(q^{-1})y(t) = & \frac{B_1(q^{-1})}{F_1(q^{-1})}u_1(t-k_1)+\dots+\frac{B_i(q^{-1})}{F_i(q^{-1})}u_i(t-k_i)+\dots \\
 & +\frac{B_n(q^{-1})}{F_n(q^{-1})}u_n(t-k_n)+\frac{C(q^{-1})}{D(q^{-1})}e(t),
 \end{aligned} \tag{3.1}$$

where y is the discretised model output, u_i is the i -th input (for $i = 1 \dots n$), e is white noise and k_i is the delay from input u_i to the system. The functions $A(q^{-1})$, $B_i(q^{-1})$, $C(q^{-1})$, $D(q^{-1})$ and $F_i(q^{-1})$ are polynomials in the backwards shift operator q^{-1} that is defined by:

$$q^{-1}x(t) = x(t-1). \tag{3.2}$$

Not all of the polynomials A , B_i , C , D and F_i are used in a particular model. The models considered in this chapter are listed in Table 3.1, together with details of the polynomials used. The model structures considered were selected due to their simplicity and previously reported success (Marco *et al.* 1996). Finite impulse response (FIR) models only use the present and past values of the inputs (the u_i 's) in order to produce an output. Auto-regressive with exogenous inputs (ARX) models also use past values of the (simulated) outputs. The addition of a moving average term to FIR and ARX models to produce moving average, with exogenous inputs (MAX) and ARMAX models, respectively, corresponds to the inclusion of a C polynomial to the model structure (see Equation 2.1). This allows more effective modelling of the noise characteristics of the system. Note that the polynomials $D(q^{-1})$ and $F_i(q^{-1})$ do not appear in Table 3.1. They were not used in any of the model structures considered, but are included in Equation 3.1 for generality.

Model structure:	Polynomials used:
FIR	B_i
ARX	A, B_i
MAX	B_i, C
ARMAX	A, B_i, C

Table 3.1 The linear black box model structures used to analyse the electronic nose data.

3.2 Bacterial strain identification experiment

The experiment was designed to evaluate the ability of an electronic nose system to discriminate between two strains of cyanobacteria (blue-green algae), one toxic and the other non-toxic. The ability of such a system to classify quickly and accurately the strain of bacteria present in an algal bloom could clearly be useful to environmental agencies, monitoring reservoirs and lakes.

The headspaces of separate cultures of the two strains of cyanobacteria, grown in a nutrient medium (BG11), were sampled periodically by an electronic nose system over 40 days (Shin et al. 2000). The nose system used consisted of six commercial metal oxide resistive odour sensors (Alpha MOS, France), and two other sensors to monitor ambient temperature (LM35CZ, National Instruments) and humidity (MiniCap 2, Panametrics) (Shin et al. 2000). The repeated exposure cycle was as follows:

- 23 min 20 s - medium only
- 2 min - medium with toxic *microcystis aeruginosa* PCC 7806 strain
- 23 min 20 s - medium only
- 2 min - medium with non-toxic *microcystis aeruginosa* PCC 7941 strain.

The outputs from the sensors were sampled every 10 seconds, producing 350,358 data vectors corresponding to 1,150 exposure cycles. A plot of a section of the data

showing the response of a single odour sensor is given in Figure 3.2. The fluctuations in the baseline signals are attributed to variations in the ambient air.

The data obtained from the experiment have previously been pre-processed to extract static parameters and subsequently analysed using artificial neural networks with considerable success (Shin et al. 2000). Here we use system identification techniques to analyse the same data in order to compare the efficacy of linear time-invariant black box models, for both static and dynamic data, with the static data based non-linear neural network methods.




Figure 3.2 A 14 hour section of the raw data, showing the response of a single odour sensor during part of the experiment. The voltage is directly proportional to sensor resistance.

Two distinct classes of models were analysed: in Section 3.3 models for the steady state responses of the sensors (static models), and in Section 3.4 models for the full temporal data (dynamic models).

3.3 Static sensor response models

Previous work on the analysis of data from electronic nose experiments has mostly concentrated on using pre-processing algorithms to extract the steady state features from the sensor signals. Accurate results have been achieved using only this steady state information (Shin et al. 2000). There are also other benefits to using this steady state data, some of which are listed below:

- Information compression - the full (dynamic information intact) data-sets can be very large, especially when the data sampling rate is high. The pre-processing compresses this information to a single value per sensor per odour exposure cycle.
- Baseline drift removal - the gas sensors employed are susceptible to poisoning effects which produce long term systematic drift in the baseline (sensor resistance values in air). Some of the pre-processing algorithms help to counteract these effects, an example is the ‘difference’ algorithm below.
- Reduction of temperature dependence - the resistances of the odour sensors are highly dependent on the temperatures of the various components of the nose system. Some physical models for the effects of temperature variations on the sensor responses suggest that an appropriate choice of pre-processing algorithm should reduce the effects of the temperature variations on the data being analysed (Gardner 1991).

These pre-processed data are traditionally then passed on to pattern recognition algorithms or artificial neural networks for classification (Gardner and Bartlett 1994; Gardner and Bartlett 1999).

In this section black box models for these pre-processed (or ‘static’) data are analysed. The four pre-processing algorithms tested are:

- Absolute response: V_b
- Difference: $(V_b - V_m)$

- Relative difference: $\frac{V_b}{V_m}$
- Fractional difference: $\frac{(V_b - V_m)}{V_m}$

where V_b is the sensor voltage when exposed to the headspace of the bacteria sample (final value at end of exposure period) and V_m is the sensor voltage when exposed to the headspace of the medium (immediately before the sensor is exposed to bacteria). The pre-processed sensor outputs for each of the six gas sensors are considered to be the inputs to MISO black box models.

The pre-processed values are heavily influenced by the ‘intensity’ of the odour to which the sensors are exposed. For certain quantitative applications this can be useful, since one might wish to obtain an estimate of the concentration of a gas in some sample. However, for this experiment, the aim is only to discriminate between the two classes of odour (the toxic bacteria and the non-toxic bacteria). For this reason, the effects of normalising the data vectors in order to try and remove the ‘concentration’ information contained within the data are investigated. The normalisation process used for this experiment is simply to take each 6-dimensional data vector and to divide it by its Euclidean length, so that each data point is projected onto the surface of the 6-dimensional unit hypersphere.

The evaluation of each model structure and pre-processing algorithm was carried out using the following procedure:

- The chosen pre-processing algorithm was applied to the data to produce an array (2300×6) of input data for the inverse models, normalised if required.
- An output vector was formed using knowledge of the experimental details. A classification output of ‘toxic’ was encoded as +1, and ‘non-toxic’ as -1.
- The data set was split at the half-way point. The first half to be used for training of the models, and the second to be used for testing.
- A MISO black box model was trained, i.e. *Matlab* was used to estimate the polynomial coefficients in the selected model structure.

- The model was used to produce a simulated output vector from the second half of the data set.
- The elements of the simulated output vector were converted to classifications simply by taking their sign. Finally this classification vector was compared with the known classification vector in order to evaluate the success of the model. Note that this method forces the classification system to always produce a classification and does not allow for a 'not known' result.

3.3.1 FIR models

Finite impulse response (FIR) models are the simplest of the black box model structures considered here. A zero-delay FIR model forms a prediction of the system output at time t by simply taking a linear combination of the system inputs at time t and at times $t - 1, t - 2, \dots, t - b_i$, where b_i is the order of the polynomial B_i in the model (see Section 2). Depending on the nature of the system being modelled, it may be appropriate to include delay terms (the k_i parameters in equation 3.1). These could be different for each of the model inputs, thus for a 6-input 1-output FIR model, the structure of the model is determined by the orders of the six B polynomials, and the six delays k_i . However, due to the fundamental similarity between our six inputs (i.e. same type of semiconducting oxide gas sensor), our investigation was restricted to models using the same order for all six B polynomials, and the same delay for all six inputs.

Early investigations showed that the most appropriate delay was zero. This was expected with this static data since a single data point corresponds to a whole exposure cycle. Hence to introduce a non-zero delay would be to expect the model to classify an odour using none of the data recorded during the sampling of that odour.

The most effective orders for the B polynomials were either two or three for each of the eight pre-processed data sets (for all four different algorithms and normalised/raw data). The results obtained using models with these optimal structures are given in Table 3.2.

The normalisation process carried out on the pre-processed data had the effect of reducing the success rates slightly. As mentioned before, the aim of the normalisation procedure is to remove some of the ‘intensity’ of odour information to leave mainly ‘type’ of odour information. The fact that this process reduced the ability of the models to classify correctly the odours in the test data set would seem to indicate that the models for pre-processed data (raw) were using a small amount of this ‘intensity’ information to produce the classifications. In a field-based application, it would be impossible to control the intensities of the signals, so for such purposes one might expect normalised data processes to yield greater success rates.

The fact that the absolute response algorithm produced a noticeably poorer success rate than the other three algorithms could be attributed to the fact that this algorithm is the most affected by the long term baseline drift evident in the data. The normalisation process reduces these effects somewhat, enabling the normalised absolute response models to perform comparably with the other models using normalised data.

Pre-processing (static) algorithm used:	Percentage success rate using	
	Pre-processed data:	Normalised pre-processed data:
Absolute response	98.8	97.6
Difference	99.7	97.4
Fractional difference	99.5	97.0
Relative difference	99.5	98.4

Table 3.2 Classification success rates obtained using zero-delay FIR models of orders two or three for static (pre-processed) data.

3.3.2 More complex model forms

The FIR models in the previous section produced very high classification success rates, however the inclusion of either A or C polynomials to the model structure to form auto-regressive with extra inputs (ARX) or moving average with extra inputs (MAX)

models increased the success rates obtained, perhaps because the more complex model structures are better able to handle the drift evident in the data.

By selecting appropriate orders for the B and A polynomials in an ARX model, a successful classification rate of 100% was obtained using each of the eight pre-processed data-sets. The orders required were mostly one for the B polynomial and between one and three for the A polynomial.

The extension of the FIR model structure to a MAX model structure produced similar success rate increases. Success rates of 100% were achieved for 7 of the 8 data sets, with 99.9% achieved on the remaining data set. The orders required for these models were mostly 2 for the B polynomials and between 0 and 3 for the C polynomial. The MAX models generally required more parameters (polynomial coefficients) than the ARX models, making the ARX models preferable in terms of simplicity and computational efficiency. Further extension of the model structures to form ARMAX models was not found to produce any significant improvement over the ARX and MAX models.

It should be noted that for the training and testing of these static models the temporal order of the data was maintained. It might be considered more realistic to randomly reorder the data vectors to avoid the possibility of the models 'learning' the characteristics of the particular experiment rather than the characteristics of the different odours. In order to enable fair comparison with the dynamic models in the next section (for which random reordering would not be as straightforward) the data sets used in this section were not reordered. However, it was noted that, when the static data sets were randomly reordered (100 times, average results taken), the success rates were only reduced by a few percent (typically around 4%).

3.4 Dynamic sensor response models

The models for the static data used in the previous section achieved considerable success, but it has been suggested (Hines *et al.* 1999; Wilson and DeWeerth 1995; Vilanova *et al.* 1996) that there is also useful information contained within the

dynamic sensor response which the pre-processing algorithms discard. In this section we form linear time-invariant black box models for the full dynamic data in order to investigate this possibility. We also investigate the effects of normalising the sample vectors (as in Section 3.3).

The model structures tested are the same as those in the previous section. The crucial difference is that, for the models for static data, a model of order d , say, takes into account the last d exposure cycles, whilst for the models for dynamic data, a model of order d takes into account the last d sampled data points, i.e. the last $10 \times d$ seconds.

The evaluation of each model structure was carried out using the following procedure:

- If required, the data-set was normalised to form an array ($350,358 \times 6$) of input data.
- As before, an output of ‘toxic’ was encoded as +1, and ‘non-toxic’ as -1. Here, however, an output of ‘medium’ was required for the periods where no bacteria were sampled. This was encoded as 0. An output vector was accordingly formed.
- The data-set was split into halves. Data from the first 20 days of the experiment were used for training the models, and data from the remaining 20 days were used for testing.
- A MISO black box model of the chosen structure was trained on the first half of the data using *Matlab*.
- The model obtained was used to produce a simulated output vector from the second half of the data.
- The simulated output was sampled at the time point just before the system switched from bacterial odour input to medium input, for each exposure cycle. Thus producing a single numerical output for each exposure cycle.
- The vector of numerical outputs was converted to a vector of encoded classifications by taking the sign of each element. This vector was then compared with the known sequence of classifications in order to evaluate the success of the model.

3.4.1 FIR models

The first model structures considered were FIR models. As with the models for static data, the set of structures considered was restricted to those having the same order polynomials for each of the inputs, and similarly the same delay on each input. FIR models for both the raw and normalised data-sets with B polynomials of orders from 1 to 25 and having delays from 0 to 5 data points (corresponding to 0 to 50 seconds) were formed and evaluated. Note that the orders of each of the six B polynomials were the same for any given model. The optimal structures for each data-set were chosen as a compromise between a high successful classification rate, and a ‘simple’ model structure, i.e. one having few parameters requiring estimation.

With the models for static data in the previous section, the optimum delay was found to be zero, as expected. However, for the models for dynamic data in this section, a non-zero optimum delay would not be unexpected. The delay required in the model reflects the physical characteristics of the system in question. The sample vessels were connected to the sensor chamber via a system of pipes and valves, thus one might expect there to be a noticeable time lapse between the switching of the valves and the arrival of the new odour at the sensor array. However, for our data-set, no delay was observed. This could be explained by the relatively low dead volume in the pipework (V_d) compared with the volumetric flow-rate of the pump (\dot{Q}_p) producing a physical delay that was short enough to be undetectable at the 0.1 Hz sampling frequency used (i.e. $\frac{V_d}{\dot{Q}_p}$ much less than 10 seconds).

Using the raw data-set, the optimum model order was found to be 12. This corresponds to an input information ‘window’ of 120 seconds being used by the model in order to calculate an estimated output at a given time. The successful classification rate achieved by this model on the test data was 91.3%. A slightly higher success rate (91.8%) could be achieved using a model of order 22, but the lower order model was deemed more suitable due to the insignificant (and unlikely to be reproducible) difference and the greatly reduced computing costs associated with using a model with 72 polynomial coefficients rather than 132.

The models for the normalised data-set achieved greater success than those for the raw data. The optimum model order was found to be 11, producing a success rate of 99.3%. The fact that the normalisation process improved the success rates for the dynamic models contrasts with the situation for the static models, where the normalisation process reduced slightly the classification success.

As discussed in Sections 3.3 and 3.3.1, the normalisation process reduces the 'intensity' information within the data and increases the significance of the 'odour type' information. For this dynamic data-set, it would appear that the models for raw data were partially classifying the odours based on correlations in the training set between the intensity of the signal and the correct classification. Thus when the models were applied to the testing set, where these correlations were no longer valid due to long term drifts in sensor responses, incorrect classifications were made. With the intensity information removed by the normalisation process, the models were able to make their classifications based on the type of odour and thus generalise more successfully from the training to testing sets.

3.4.2 ARX models

The inclusion of an A polynomial to the black box structure to form ARX models allows the model to consider past values of the model output as well as past and present values of the inputs. The delays on the inputs were set to zero, since (as discussed in the previous section) the delay required is a reflection of the physical characteristics of the system and so should be model-independent. Models with orders from 0 to 15 for the A polynomial and from 1 to 15 for the B polynomials were formed and evaluated for both the raw and normalised data sets. The percentage successful classification rates obtained on the test data set by each of the 240 models for the raw data are plotted in Figure 3.3.

It is evident from Figure 3.3 that the order of the B polynomials is a more significant factor in the success of the model than the order of the A polynomial. For low order B polynomials (around 3-7), the inclusion of a low order A polynomial increased the

success rate slightly. However the greatest success rate overall for this raw data set was obtained using the order 12 FIR model of the previous section (91.3%).

The results for the normalised data set are similar in that the greatest success was observed when no A polynomial was included in the model structure, so that the maximum success rate (99.3%) was obtained with an 11th order FIR model.




Figure 3.3 Plot of successful classification rates against orders of A and B polynomials when ARX model structures (with zero input delays) were tested on the raw data set.

3.4.3 MAX and ARMAX models

In an attempt to further improve the success rates obtained, the addition of a C polynomial (a moving average term) to the FIR and ARX model structures to form MAX and ARMAX model structures was investigated. For MAX models, the order of the A polynomial is zero, so the investigation was restricted to zero delay models with

orders from 1 to 15 for the B polynomials and from 0 to 15 for the C polynomial. The percentage successful classification rates obtained on the test data set by each of the 240 models for the raw data are plotted in Figure 3.4.




Figure 3.4 Plot of successful classification rates against orders of B and C polynomials when MAX model structures (with zero input delays) were tested on the raw data set.

Once again the optimum models only involve the B polynomials, indicating that the extension from an FIR model to an MAX model is not beneficial. Similar results were obtained using models for the normalised data set.

The addition of an A polynomial (an auto-regressive term) to the MAX models to form ARMAX models did not increase the success of the models. This suggests that (recent - 10 to 150 seconds previous) past output values do not provide any useful information that can be used to produce a more accurate value for the current output. This makes sense when the drift seems to vary diurnally rather than by the minute - probably associated with changes in the temperature of the bacterial samples.

3.4.4 Models for dynamic data with extra inputs

As mentioned in Section 3.2, the sensor chamber contained temperature and humidity sensors in addition to the six gas sensors. In the previous sections, only the outputs from the gas sensors were used as inputs to the black box models. In this section, the outputs from the temperature and/or humidity sensors are included as inputs to our models. The outputs of the gas sensors used are known to be affected by variations in their operating temperature and ambient humidity. The biological activity and hence odour production level of the sample could also be expected to vary with the ambient temperature. The outputs from the temperature sensor indicated significant diurnal variations (up to ten Kelvin) in the ambient temperature throughout the experiment. Thus it would not be unreasonable to expect the inclusion of temperature and/or humidity information to produce some improvement in the classification performance of these models.

FIR models of order 12 for the raw data set and 11 for the normalised data set (the extra inputs were not included in the normalisation process) were produced and evaluated. The results obtained are shown in Table 3.3. The results for the normalised data show too little variation to draw any firm conclusions. The results for the raw data do show slight variation, and suggest that the temperature and humidity signals have negligible effect here. If anything, the inclusion of the humidity signal degraded the raw data somewhat.

Extra inputs:	Percentage successful classification:	
	Raw data	Normalised data
None	91.3	99.3
Temperature and humidity	92.6	99.4
Temperature	91.1	99.3
Humidity	86.3	99.4

Table 3.3 Successful classification rates for MISO FIR models having six, seven or eight inputs.

3.5 An alternative train-test strategy

So far all of the models considered have been trained and tested on very large data sets. In practice it would be more likely that the system would be trained using a relatively short calibration experiment, and then subsequently used for a substantial length of time before requiring recalibration. Clearly it is best, from an operator's point of view, that calibration experiments are short and infrequent.

With this in mind the performances of the previously considered model structures have been evaluated using a rather different train-test strategy. Each model was trained on day i (for $i = 1$ to 39) and then tested on each subsequent day j ($i < j \leq 40$).

3.5.1 Static sensor response models

First the models for static (pre-processed) data are considered. Here an FIR model structure of order three was selected (found to be optimal in Section 3.3.1), and models trained using each day in turn. Both raw and normalised data were used. All four previously considered forms of pre-processing algorithm were employed, thus 320 models were formulated. Each model was then tested on data for each of the days subsequent to its training day, and the percentage correct classification on the test data calculated. Encouragingly, it was found that the success of the models did not diminish significantly as the testing day became distant from the training day. As a fairly typical example of what was found, a plot of percentage success rate versus training day and testing day for fractional difference (not normalised) pre-processed data is given in Figure 3.5.




Figure 3.5 A plot of percentage success versus training and testing day for an FIR model of order 3, using ‘fractional difference’ (not normalised) pre-processed data.

Note that the data points corresponding to models tested on days prior to the day of the training represent a somewhat unrealistic scenario and are thus omitted from Figure 3.5.

The performance of the models can be seen to be very impressive, even many days after the calibration experiment. A model trained on day two maintains a 100% success rate for many (36) of the following 38 days. This might indicate that a well-performed, day-long calibration experiment could be expected to be sufficient training for a model to perform well over a period in excess of 40 days.

For the pre-processed data which were not normalised, the choice of pre-processing algorithm made little difference (although the ‘difference’ algorithm was slightly less successful than the others). However, for the normalised pre-processed data, there was a clear difference. The ‘absolute response’ and ‘relative difference’ algorithms

performed similarly to their not normalised counterparts, but the ‘difference’ and ‘fractional difference’ algorithms performed significantly worse.

3.5.2 Dynamic sensor response models

For the models for dynamic data an FIR model structure of order 12 was selected for the raw data, and order 11 for the normalised data, producing two sets of 40 models. These orders were chosen as a result of the findings in Section 3.4.1. Models were trained using each day in turn, and again tested on the data from each subsequent day of the experiment. A plot of the percentage success versus the training and testing days is given in Figure 3.6.




Figure 3.6 A plot of percentage success versus training and testing day for an FIR model of order 12, using the raw dynamical data.

It might be expected that, given a model trained on day i , the performance of the model when tested on day j would steadily decrease as the difference between j and i becomes

large. However, whilst such a trend is clearly visible for models trained on certain days (e.g. day 10), this is not always found to be the case. In fact, plotting percentage success against time between training and testing phases shows no significant correlation. Again this might suggest that infrequent calibration could be employed with some success.

It is evident from Figure 3.6 that models trained on days prior to day 8 were not very successful when tested on later days. This could be attributed to an unknown experimental discontinuity occurring approximately one week into the experiment.

It is interesting to note that some of the models produced success rates of up to 100% even many days after the training day. An example of the performance of one of the more successful training days is given in Figure 3.7, for which the success rate stays above 89% right up to day 39. This compares well with equivalent models trained on the whole data set.




Figure 3.7 Plot of percentage successful classification versus testing day for an FIR model of order 12 for dynamic data. The model was trained on data from day 12 of the experiment and tested on subsequent days.

Similar results were obtained using models trained and tested on normalised data. So for all of the data sets considered, the model structures used were capable of producing acceptable classification success rates many days after the model training was carried out.

3.6 Conclusions

The models for the static strain identification data were able to produce very high success rates: 98.4% for normalised, 99.7% for non-normalised data, using 50% cross validation (using MISO FIR models of order two or three). These success rates increased to 100% on the extension of the FIR models to ARX or MAX model structures.

The models for static data work best with a ‘memory’ of the last three or so response cycles in order to make a classification. In practical applications it may not be acceptable to wait for three complete exposure cycles to obtain a classification. In contrast, the models for dynamic data, even though they use higher order polynomials, require only a single exposure cycle to produce a classification (since for the dynamic data, 10 data points corresponds to 100 seconds, rather than 10 exposure cycles).

The models for the dynamic data produced maximum success rates of 91.3% for the raw data (using an FIR model of order 12) and 99.3% for the normalised data. This showed that the normalisation process can be an effective tool for improving the classification success rates when long term drift might otherwise adversely affect the system performance. This might be attributable to the fact that the normalisation process reduces the ‘strength’ of smell information in the data, and thus forces the model to learn (and subsequently make classifications based on) the ‘type’ of smell encountered.

The success rates of the models for static data, and those for normalised dynamic data, compare well with the success rates achieved using the best artificial neural networks (Shin et al. 2000) where a successful classification rate of 100% was obtained using a Fuzzy ARTMAP applied to static data from the same experiment.

The success of the black box models might be improved by the use of non-linear system identification techniques. It should also be noted that the criteria used for parameter estimation were not precisely the same as that used for the evaluation of the resulting models. Thus further work to produce a more appropriate parameter estimation algorithm may improve the success rates obtained.

It has been shown that simple linear black box (inverse) models for an electronic nose system can be successfully employed for strain classification of cyanobacteria. The models performed as well as the previously employed artificial neural network techniques, with the advantage that they require less computing power to implement. Thus such modelling techniques could be more appropriate for use in applications where computing power is limited, such as in a handheld instrument.

Chapter 4

Cyanobacteria growth phase

identification using system identification techniques

In this chapter, the techniques discussed in the previous chapter are applied and extended to the problem of classifying the growth phase of cyanobacteria. The experiment considered here was designed to test the ability of the electronic nose system to discriminate between the different growth phases of a single strain of bacteria. The levels of biological activity of the bacteria vary as the culture progresses through its life cycle, which consists of four distinct growth phases, known as the ‘lag’, ‘log’, ‘stationary’ and ‘late stationary’ (or ‘death’) phases (Shin et al. 2000). This presents a more challenging problem for the system since the odours given off by the bacteria may change only slightly over the course of their life cycle. Thus the classes to be resolved are expected to be closer together (and less sharply defined) than those for the simpler strain identification experiment of the previous chapter.

The knowledge of which growth phase a sample of bacteria is currently in provides a useful indication of the likely future progress (viability) of the bacteria. The phase also defines the rates at which the cells take in substances from around them. Thus information about the current growth phase could be used to predict accurately the

dosage levels of antibiotic required to challenge the bacteria. If utilised in medical applications, such predictions would yield substantial benefits to the healthcare industries.

4.1 Black box modelling

The purpose of the electronic nose system is to produce a classification or identification of unknown gases or odours to which it is exposed. Thus the processing of the sensor outputs to produce a classification of the odour input can be thought of as an inverse black box model for the nose system. This is illustrated in Figure 4.1.




Figure 4.1 Diagrammatic representation of the electronic nose modelling problem.

For this experiment the bacteria must be classified according to which of four possible growth phases the bacteria are currently in. Quantitative information about the strength or concentration of the odour inputs is unnecessary. These simple classifications can be obtained from black box models outputs in a number of ways, requiring the use of either single or multi-output black box models.

The *System Identification Toolbox* (Ljung 1995) within *Matlab* (Version 5) provides the user with many functions for creating, evaluating and using black box models. In this chapter, the model structures used were almost without exception multi-input, multi-output (MIMO) finite impulse response (FIR) models. This is in contrast with the simpler strain identification experiment considered in the previous chapter, where FIR, ARX, MAX and ARMAX models were considered. The model structure considered was selected due to its simplicity (and thus computational efficiency) and

previously reported success (Marco *et al.* 1996). FIR models effectively just take a linear combination of the present and past values of the inputs (the u_i 's) in order to produce an output.

For example, MIMO FIR models with n inputs and m outputs have the general form (Ljung 1987):

$$\mathbf{y}(t) = \mathbf{B}(q^{-1})\mathbf{u}(t) + \mathbf{e}(t), \quad (4.1)$$

where \mathbf{y} and \mathbf{e} are m -dimensional vectors, \mathbf{u} is an n -dimensional vector and \mathbf{B} is an $m \times n$ matrix. The elements of \mathbf{B} are polynomials in the backwards shift operator q^{-1} , which is defined by:

$$q^{-1}x(t) = x(t-1). \quad (4.2)$$

Thus \mathbf{B} can be written:

$$\mathbf{B}(q^{-1}) = \begin{pmatrix} b_{11}(q^{-1}) & b_{12}(q^{-1}) & \cdots & b_{1n}(q^{-1}) \\ b_{21}(q^{-1}) & b_{22}(q^{-1}) & \cdots & b_{2n}(q^{-1}) \\ \vdots & \vdots & \ddots & \vdots \\ b_{m1}(q^{-1}) & b_{m2}(q^{-1}) & \cdots & b_{mn}(q^{-1}) \end{pmatrix}, \quad (4.3)$$

where for $i = 1 \dots m$ and for $j = 1 \dots n$:

$$b_{ij}(q^{-1}) = b_{ij}^1 q^{-k_{ij}} + \dots + b_{ij}^{d_{ij}} q^{-(d_{ij}+k_{ij})}, \quad (4.4)$$

where d_{ij} is the order of the polynomial b_{ij} , and k_{ij} is the delay from input j to output i .

4.2 Growth phase identification experiment

The experiment considered here involved a single (toxic) strain of cyanobacteria, monitored over a 40 day period. As well as electronic nose data, information

concerning the mean size of the bacterial cells and the biomass present in the cultures was recorded using a CellFacts instrument (Microbial Systems Ltd.). This enabled the identification of the four distinct growth phases through which the bacteria pass during their life cycle. Neural networks have been successfully used (by Shin et al. 2000) with the electronic nose data to classify the bacteria into each of the four growth phases, obtaining classification success rates of up to 95.1%. In this section we apply MIMO linear black box models to the same data in order to investigate the ability of such techniques to tackle this challenging problem.

The experiment in question was intended not only to test the ability of an electronic nose to discriminate between the different growth phases of a cyanobacteria strain, but also to investigate the reproducibility of the measurements and success rates. For this reason the experimental system consisted of three vessels, two containing nominally identical cultures of toxic *microcystis aeruginosa* PCC 7806 in nutrient medium (BG11), and one reference vessel containing only the nutrient medium. The headspaces of these vessels were connected via a system of pipes and computer-operated valves to an electronic nose system. The nose system used consisted of six commercial metal oxide resistive odour sensors (Alpha MOS, France), and two other sensors to monitor ambient temperature (LM35CZ, National Instruments) and humidity (MiniCap 2, Panametrics) (Shin et al. 2000). The repeated exposure cycle was:

- 50 min - medium only
- 5 min - medium and toxic *microcystis aeruginosa* PCC 7806 strain sample 1
- 50 min - medium only
- 5 min - medium and toxic *microcystis aeruginosa* PCC 7806 strain sample 2.

The sensor outputs were again sampled every 10 seconds, producing 361,698 data vectors corresponding to 548 full exposure cycles.

The information collected using the CellFacts instrument was used to produce a 'correct' classification vector for the data, identifying the four growth phases of the bacteria using understanding of the biological processes involved. It should be noted that the boundaries between the growth phases were by no means sharp, making the

manual identification of the phases subjective, and thus the noise data classification problem highly non-trivial. Some of the data from the CellFacts instrument are plotted in Figure 4.2.




Figure 4.2 Plot of the data from the CellFacts instrument for the growth phase identification experiment. The upper plot shows the general increase in biomass (cell counts) with time. The lower plot shows the variation in mean size of the bacteria cells with time. The four growth phases (lag, log, stationary and late stationary), are labelled I to IV in each plot.

4.3 Static sensor response models

In this chapter, as in the last, linear black box models for both static and dynamic data are considered. This section deals with models for the static data. The pre-processing algorithms used were the same as those in the previous chapter, namely:

- Absolute response: V_b
- Difference : $(V_b - V_m)$
- Relative difference: $\frac{V_b}{V_m}$
- Fractional difference: $\frac{(V_b - V_m)}{V_m}$.

However, the details of how the black box techniques should be applied to the data was less clear-cut. In the previous chapter, the basic aim of the modelling was to produce an algorithm for classifying an odour into one of two classes (toxic or non-toxic), so the classification output was encoded simply as +1 for a classification of 'toxic', and -1 for 'non-toxic'. In this experiment, the odour must be classified into one of four classes corresponding to the current growth phase of the bacteria. Hence the simple numerical encoding of the output used in the previous chapter is no longer appropriate.

Two methods of encoding the classification output of the model into a numerical output were considered. The first was an analogue of the method used in the previous chapter, whereby the four growth phases were each allocated a numerical label (a (possibly zero) integer) between -2 and +2 and a multi-input, single-output model was formulated using the given phase labelling system. All of the different labelling combinations possible were tested to find the optimum labelling system for each model structure. This method was time consuming, though moderately successful, producing a maximum success rate of approximately 66% when trained on half of the static data, and tested on the second half. However, the success of this method relies on the labelling combination chosen – if the classes (and the corresponding data vectors) were such that in some sense they could be sorted into an order, then this method could work well. However, if the differences in the classes do not lend themselves well to such a one-dimensional ordering, then other methods could be expected to produce better results.

The second method considered was the use of multi-output models to encode the different classification. An inverse model with four outputs (and six inputs, as before) was chosen, corresponding to the four different possible classifications of the odour.

Thus a ‘perfect’ classification of ‘growth phase 1’ would be a 4-dimensional vector with a +1 in the first co-ordinate and zeros elsewhere. This method was found to be more successful than the single-output model method. Though it should be mentioned that analogous 2-output models were tested for the data in the previous chapter but were out-performed by the single-output models previously used. This is unsurprising since there were only two classes in that case which were fairly well separated.

The effects of normalising the sample data vectors were also investigated. The evaluation of each model structure and pre-processing algorithm was carried out as follows:

- The chosen pre-processing algorithm was applied to the data to produce an array (1096×6) of input data for the inverse models, normalised if required.
- An output array (1096×4) was formed using the data from the CellFacts instrument. A classification output of ‘growth phase i ’ was encoded as a 4-dimensional vector with +1 in the i -th position and zeros elsewhere.
- The data vectors corresponding to the ‘Toxic 1’ bacteria culture were separated from those corresponding to the ‘Toxic 2’ culture, thus forming two separate data-sets (each with 548 input and output vectors).
- Each data-set was randomly reordered and then split into two halves. The first to be used for training and the second for testing of the models.
- A MIMO black box model was trained (see Section 4.1 for details of the MIMO model structure used).
- The model was used to produce a simulated output array from the second half of the data-set.
- Each output vector (row in the output array) was converted to a growth phase classification by choosing the output with the largest (positive) value.
- The resulting classification vector was compared with the actual growth phase data to evaluate the success of the model.

Note that, for this growth phase data set, the data were randomly reordered (unlike those in the previous chapter). This was done for two main reasons:

- a) To give a more realistic indication of how well the models could perform on real world data.
- b) To enable cross validation testing. Unlike the experiment in the previous chapter, the different classes are not distributed evenly (with time) through the data-set. Thus training the model on the first half of the data and testing on the second half would be nonsensical, since only growth phases 1 and 2 would be ‘seen’ by the model training algorithm, so it would be impossible to get useful results when testing the model on data corresponding to phases 3 and 4.

4.3.1 MIMO models for static sensor responses

As with the strain identification experiment of the previous chapter, fairly low order models were found to produce the best compromise between model simplicity and success rate. The exact order used varied across the 16 pre-processed data sets (choice of two bacteria cultures, four algorithms, each one subsequently normalised or left unchanged) varied between one and four.

For the Toxic 1 bacteria the most successful of the pre-processing algorithms was the relative difference algorithm, not normalised, producing a success rate of 78.6% using a model of order four. For the Toxic 2 bacteria, the success rate was highest (82.3%) using the static difference algorithm, again not normalised, this time using a model of order three. For the results using other pre-processing algorithms see Table 4.1.

Data set:	Percentage successful classification using pre-processing algorithm:			
	Absolute response	Difference	Fractional difference	Relative difference
Toxic 1	66.1	75.7	78.1	78.6
Toxic 1 (normalised)	70.2	68.5	73.9	77.5
Toxic 2	67.3	82.3	81.4	78.9
Toxic 2 (normalised)	67.5	80.3	75.2	77.7

Table 4.1 The success rates for FIR MIMO models of orders between one and four for static (pre-processed) data from the growth phase experiment. The percentages are averages over 10 different random reorderings using two-fold (50%) cross-validation.

Notice that, as with the strain identification experiment discussed in the previous chapter, the normalisation process reduced the success rate slightly in most cases. This can be attributed again to the odour intensity being related to the cell count and hence growth phase. Also note that the absolute response pre-processing algorithm was markedly less successful than the other three algorithms.

4.4 Dynamic sensor response models

The MISO models for dynamic data from the bacterial strain identification experiment considered in the previous chapter achieved considerable success. In this section we consider similar MIMO models applied to the bacterial growth phase data.




Figure 4.3 Plot showing the considerable long term variations in the output voltage from a single sensor over the whole (40 day) period of the growth phase experiment. There is clear evidence of diurnal peaks probably associated with the temperature of the biomass. The temperature of the room fluctuated daily.

As with the models for static data in Section 4.3, the survey of MIMO model structures is restricted by the available computing algorithms and time constraints to FIR MIMO models of various orders. However, unlike the case for the strain identification experiment in the previous chapter, the appropriate method for preparing the data for the models is not obvious. As mentioned in Sections 4.2 and 4.3, the experimental procedure was such that the data-set obtained was effectively two data-sets interleaved. For the pre-processed (static) data this posed no real problem, however for the dynamic data, things are less straightforward. As with the static data, it was desirable to treat the data from the two cultures separately, but this necessitated splitting and reforming the data-set into two halves. The sensor values drifted significantly over the course of the experiment (see Figure 4.3) so, whilst splitting up the data-set into two, it was decided to shift each response segment to remove baseline drift in an analogous

manner to the ‘difference’ pre-processing algorithm for the static data. The effects of this ‘dynamic pre-processing’ can be seen in Figure 4.4.




Figure 4.4 The upper plot shows a section of the output data from a single sensor for the growth phase experiment. The lower plot shows the same section after dynamic pre-processing to shift each response cycle and thus remove some of the effects of the baseline drift evident in the upper plot.

Since the data were already being split up into individual response cycles, it was decided to also randomly reorder these response cycles to better simulate real-world application and facilitate fair comparison between the performances of the static and dynamic models.

MIMO FIR models of various orders for these dynamically pre-processed data-sets were formulated, and the effects of two different normalisation methods were investigated. The evaluation of each model was carried out as follows:

- The dynamic pre-processing method was applied to the original growth phase data-set to produce two separate dynamic input data-sets (each being 180840×6) corresponding to the Toxic 1 and Toxic 2 cultures.
- An output array (180840×4) was formed for each of the two cultures using the classifications gained from the CellFacts information.
- The response cycles of each data-set were randomly reordered.
- If required, a normalisation process was applied to the input data.
- A MIMO FIR model was trained on the first half of the set in question using *Matlab*.
- The model obtained was used to produce a simulated output from the second half of the data-set.
- The simulated output was sampled at the appropriate points, and the output vectors obtained were converted to a sequence of (274) classifications.
- The sequence of classifications obtained was compared with the correct classifications to evaluate the success of the model.

The above process was repeated ten times with different random reorderings, and the results averaged.

One of the two normalisation processes investigated was the same as used in the previous chapter, where each data vector was scaled so as to have unit length. This constrains the position of the input vector in sensor space to move about on the surface of the six dimensional hypersphere (there were six sensors used), removing to some extent information regarding the intensity (or strength) of the odour and forcing the system to utilise instead information regarding the type (or note) of odour. The other normalisation process considered was to normalise (i.e. scale to unit length) only the input data vectors corresponding to the times when the bacterial headspace was sampled, leaving the data vectors corresponding to the headspace samples from the (non-bacterial) control vessel (containing only the nutrient medium) in their original form. The latter normalisation process consistently produced superior success rates to those obtained using the former method.

FIR model orders between 1 and 20 were investigated and generally a model of order around 10 was deemed an appropriate compromise between model complexity and success rate. The success rates obtained with order 10 models are given in Table 4.2.

	Percentage success rate for dynamically pre-processed data-set with		
	no normalisation:	normalisation of whole input data-set:	normalisation only during bacterial sampling:
Bacteria culture:			
Toxic 1	40.4	45.3	61.5
Toxic 2	52.2	61.9	76.6

Table 4.2 The success rates obtained for 10th order FIR MIMO models for dynamically pre-processed and randomly reordered growth phase data. The models were trained on half of the data and tested on the remaining half, i.e. two-fold (or 50%) cross-validation was employed.

4.5 Models for filtered dynamic data

The dynamic pre-processing in the previous section achieved some success, but still considerably less than the equivalent static models or the neural networks used by (Shin et al. 2000). In this section the effects of filtering the dynamic data to remove the effects of both high frequency noise and low frequency drift / temperature are investigated.

The *Signal Processing Toolbox* within *Matlab* contains functions facilitating the use of Butterworth filters to achieve the desired results. Initial estimates and investigations showed that low cut-off frequencies in the range 20 μ Hz to 0.2 mHz and high cut-off frequencies in the range 1 mHz to 50 mHz were appropriate. Low and high pass filters with cut-off frequencies in the ranges described above were tested using FIR MIMO models (order 10), with the intention of subsequently using a bandpass filter with the optimum low and high cut-off frequencies. The procedure for evaluating the efficacy of the different filtering strategies was similar to that used in the previous section. The

differences were that, prior to the random reordering process, the data were passed through a Butterworth filter (order 3), and subsequent to the reordering, only one tenth of the data was used for training (instead of one half), and one tenth for testing. This was to reduce processing time and enable a thorough search of the cut-off frequencies to be used.

The results obtained using the models for filtered dynamic data were not as clear as might have been expected. The high pass filter was intended to remove the low frequency drift, associated mainly with diurnal temperature fluctuations. It was successful in doing this (see Figure 4.5). However, since the data were subsequently dynamically pre-processed in such a way as to correct for simple baseline drift, the high pass filtering appeared to have little overall effect on the successful classification rates of the systems. This should not be assumed to be a typical case. If the baseline drift were not removed by the particular form of dynamic pre-processing chosen, then the high pass filter could be expected to improve the success rates of the models.




Figure 4.5 The upper plot shows the output of sensor 1 over the whole period of the experiment. Significant low frequency drift is observable (mostly due to daily temperature fluctuations). The lower plot shows the sensor 1 data after being passed through a high pass filter with cut-off frequency 60 μ Hz.

The low pass filter was intended to remove high frequency noise from the data. Examining the data before and after the filtering process showed that it did this effectively (see Figure 4.6). However, the results obtained did not show a significant increase in successful classification rate with filtered data compared with those obtained using non-filtered data. From this it could be concluded that the signal to noise ratio in the data analysed was sufficiently large that the high frequency noise was not a limiting factor for the success rates of the models. Again, this may not always be the case. With a different experimental set up operating under different conditions (especially in a field-based rather than laboratory-based application), there might be more significant high frequency noise present in the data. In such a situation it would not be unreasonable to expect the filtering to have a pronounced positive effect on the success of the models.




Figure 4.6 The upper plot shows the output of sensor 1 over a short period of the experiment. Some high frequency noise is observable. The lower plot shows the sensor 1 data after being passed through a low pass filter with cut-off frequency 7 mHz.

4.6 Conclusions

The models for the growth phase identification experiment data were significantly less successful than those for the simpler strain experiment detailed in the previous chapter. The static data from the experiment were randomly reordered to simulate real-world applications (and also to avoid the fact that if the data had not been reordered then the model would have encountered a training set consisting of all the phase 1 data points, followed by all the phase 2 data points etc.). Successful classification rates of up to 82.3% were obtained with the static data. This compares with 95.1% obtained elsewhere (Shin et al. 2000) using an LVQ artificial neural network.

The models for dynamic data were considerably less successful than those for static data. Although a maximum success rate of 76.6% was achieved using one particular data-set with one of the normalisation techniques considered, this success was not repeated using the alternative data-set (where a maximum of 61.5% success was achieved). The use of high and low pass filters to remove low frequency drift and high frequency noise present in the dynamic data was considered but not found here to produce significant improvements in classification performance. It is noted however, that for other experimental systems operating under different conditions, the filtering process may prove useful.

The lack of success of linear black box techniques (for both static and dynamic data) to identify growth phase in comparison to non-linear neural network techniques (Shin et al. 2000) may be attributed to the fact that the models used here are linear in nature, whilst the processes being modelled are clearly not. Only FIR models were investigated for the growth phase data due to the computational demands of more complex model structures when dealing with very large data-sets. It is possible that more complex model structures (such as ARX, MAX, ARMAX etc.) may be capable of producing better results, although these are still linear model structures so will still have some of the same limitations.

Possible future work includes the application of non-linear system identification techniques to similar data, and work to produce a parameter estimation algorithm that uses the same criterion as that used for evaluation of the models.

In conclusion, it has been shown that simple linear black box (inverse) models for an electronic nose system can be employed for growth phase classification of cyanobacteria, with moderate success. However, for the complex problem considered in this chapter, the simple models could not produce results that were able to compete with those achieved with artificial neural techniques. This contrasts with the results of the previous chapter where it was shown that for simpler classification problems the black box models could match the non-linear neural network techniques. The black box models still hold the advantage that they require less computing power to implement, and thus could be attractive for use in applications where processor power is at a premium (e.g. a handheld instrument with a 8-bit PIC microcontroller). Future

refinements of the techniques could make them suitable even for challenging classification problems such as the one considered here.

Chapter 5

The experimental electronic nose system

The data for the previous chapters came from experiments involving very complex odours (i.e. odours containing many different chemical components). For the empirical black-box modelling techniques considered thus far, this complexity presents no difficulty. However, for the physical models of the following chapters, which seek to describe the ways in which the analytes interact with the sensor materials, it is preferable to work (at least initially) with simpler analytes. For this reason an experimental electronic nose system was constructed, enabling experiments to be carried out on demand with a range of analytes, and under a range of operating conditions. This facilitated the collection of a range of high quality data sets for validation of mathematical models for single analyte or simple mixture sensor responses. This chapter details the nose system used.

5.1 Overview of the system

A simple schematic of the electronic nose system is given in Figure 5.1. The system was essentially composed of the following components:

- sample containers

- computer-controlled gas delivery system
- sensor chamber containing four metal oxide gas sensors, plus a temperature sensor and a humidity sensor
- an external unit containing a gas flow sensor and a sensor to monitor ambient temperature
- interface electronics, and
- a personal computer (PC) to run software which controlled the operation of the gas delivery system and recorded the outputs of the interface electronics.

Each component is discussed in more detail in the following sections. The hardware system was designed and assembled in-house, by Professor Julian Gardner, Doctor James Covington, Ian Griffiths, and with small modifications by myself. For a photograph of the whole system in operation, see Figure 5.2.




Figure 5.1 A schematic of the main components of the electronic nose system.



Figure 5.2 A photograph of the nose system in operation.

5.2 The sample containers and gas delivery system

The samples were held in up to four glass containers, held in a dry-block heater. Note that the analytes used were invariably sufficiently volatile at room temperature, that heating of the samples was not required. The lids of the sample containers were each drilled with two holes. To each sample pot lid, a check (i.e. one-way) valve was fitted to the inlet hole (to prevent headspace degradation and exposure of operators to potentially hazardous fumes). PTFE tubing (approximately 1 mm internal diameter) connected the outlet hole to a four-way valve block situated on top of the interface unit. The valve block consisted of a block of PTFE, drilled out to accept four solenoid valves (connected to, and controlled by, the computer via the interface unit). The block had four inlet holes – one each for the four outlet tubes from the sample containers, and a single outlet hole. Further tubing connected the outlet hole (through a throttle valve and a gas flow sensor) to a Y connector above the sensor chamber. The other branch of the Y connector was left open to allow mixing of the analyte gas and air. At the base of the sensor chamber was a pump sucking the gases through the system. The

exhaust of the pump was vented via pipes to the atmosphere. A schematic of the gas delivery system is given in Figure 5.3 and a photograph is given in Figure 5.4.

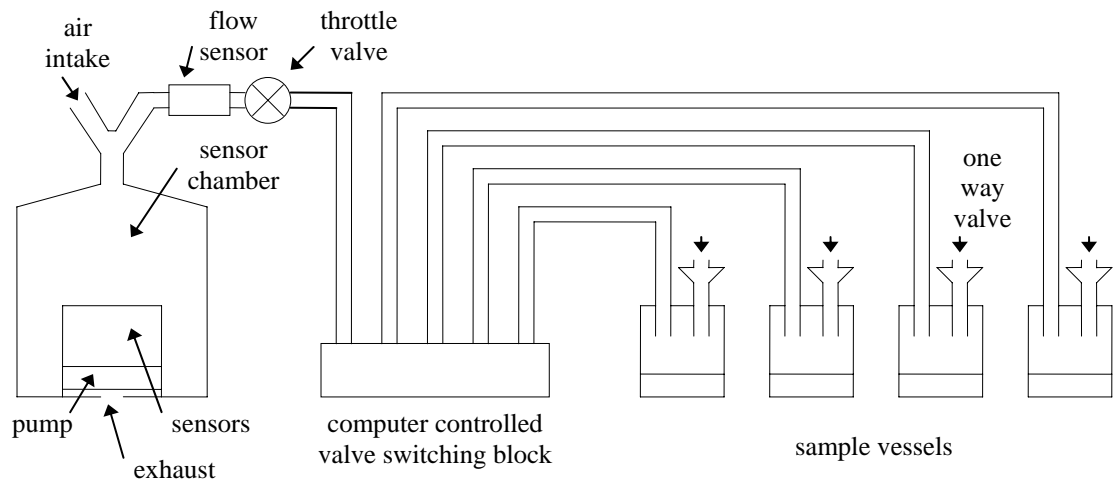


Figure 5.3 A schematic of the gas delivery system.




Figure 5.4 A photograph of the gas delivery system.

5.3 The chamber unit

The chamber unit contained a brass chamber (approximately 6 cm in diameter, 9 cm high) mounted on a PTFE base. Mounted on the PTFE base were four tin oxide gas sensors, a temperature sensor and a humidity sensor. The sensors used were:

- Alpha MOS P.40.1
- Alpha MOS P.A.2
- Alpha MOS T.70.2
- Alpha MOS T.30.1
- Temperature: LM35CZ
- Humidity: Panametrics MiniCap 2

It should also be noted that the chamber was also capable of holding two conducting polymer gas sensors, though these were not present for the experiments conducted here. An air pump was mounted below the chamber, drawing gases through it, and exhausting through tubing leaving the unit.

The chamber unit also contained interface electronics for the sensors in the chamber – including filters to suppress noise on the sensor outputs, and power supplies for the sensor systems and the pump.

5.4 The interface unit

The interface unit was manufactured in-house at the University of Warwick to a design by Professor Julian Gardner, James Covington and Ian Griffiths. The unit contained the following:

- two power supplies, one for digital electronics and one for analogue electronics,
- a valve interface card, which connected directly to the valve block described in

section 5.2,

- a sensor interface card, which contained interface electronics for up to six conducting polymer gas sensors (not used for the experiments considered in this thesis), and
- further inputs and outputs to carry signals and power to / from the PC, the chamber unit, and the flow sensor / ambient temperature sensor unit.

Mounted above the interface unit was a gas flow sensor (Honeywell AWM3100V) and a temperature sensor to monitor the ambient temperature (LM35CZ).

5.5 The computer and controlling software

The personal computer (PC) contained an Intel Pentium 133MHz processor, 96Mb RAM, and communicated with the interface unit via a National Instruments LPM-16 interface card. The PC used the Microsoft *Windows 98* operating system.

The software to automatically control the operation of the valves, and record the outputs of the sensors was written by myself (with a portion of the program adapted from software previously written by James Covington) using *Labview 5*. A photograph of the output display of the software is given in Figure 5.5.



control8.vi
 C:\elec-nose\graham\control8.vi
 Last modified on 03/12/01 at 16:09
 Printed on 14/10/02 at 17:03

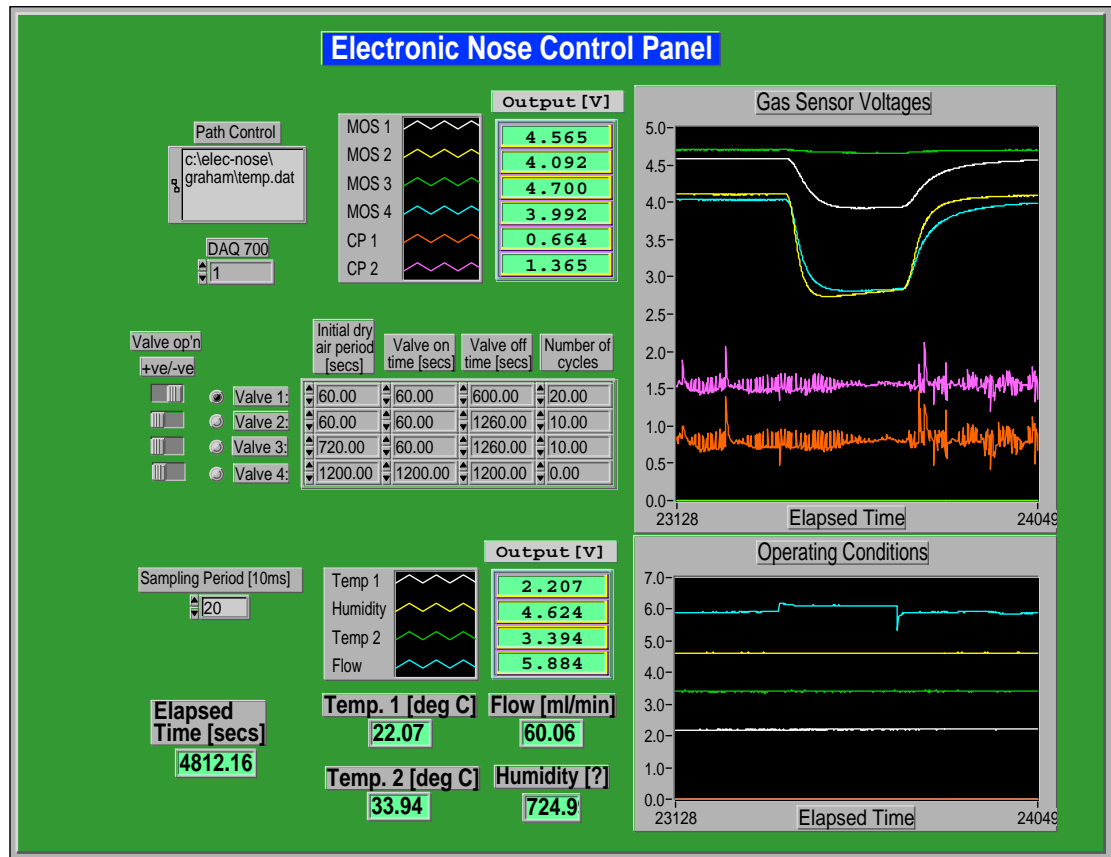


Figure 5.5 The system control software in operation.

5.6 System operation

The software was designed to enable the continuous operation of the nose system for repeated sampling of the headspaces of a maximum of four samples in rotation, with one of the four samples usually being left empty to provide the baseline. The standard mode of operation then consisted of the following automatically repeated cycle:

- Baseline sampling of ambient air via sample container number one, typically for several minutes (up to an hour) to ensure a stable baseline reading
- Sampling of sample container number two for a short period of time (typically around 30 seconds to one minute).
- Long baseline sample.
- Sampling of sample container number three.
- Long baseline sample.
- Sampling of sample container number four.

5.7 Example outputs

In Figure 5.6, the response of sensor 1 to repeated inputs of acetone and isopropyl alcohol is plotted over a 13 day experiment. This shows the long term variations observed. In Figure 5.7, data from a shorter time period are plotted, showing the different responses to the two analytes.




Figure 5.6 Plot showing the response of sensor 1 to repeated inputs of acetone and isopropyl alcohol over a 13 day period.




Figure 5.7 Plot showing the response of sensor 1 to exposure to first acetone, then isopropyl alcohol.

Chapter 6

Differential equation models for electronic noses

The work detailed in the previous chapters involved data-based models, the structures of which were selected on a purely empirical basis. There was no physical justification or motivation for the models used – they were chosen purely for their simplicity of implementation and observed efficacy when tested on experimental data. In this and the following chapters, a contrasting approach to tackling the modelling problem is employed. Physical models for electronic nose systems are proposed, then translated into mathematical models (systems of nonlinear ordinary differential equations). These models are then analysed and utilised in a classification system.

In this chapter, a general framework for (mechanistic) mathematical models of the whole electronic nose system is proposed. The practical usefulness of such models is discussed, and the concept of identifiability of such models is introduced.

6.1 Physical models versus empirical models

The potential advantages of physically-motivated mathematical models over data-based empirical models are numerous. Experimenting exhaustively with different

model structures can only provide limited information concerning the reasons why some structures perform better in certain circumstances than others. Thus when considering a new experimental situation, the modeller might not be able to predict which models will be most suitable without performing extensive computational experiments. If the models being used are based on a physical model for the system, then subsequent changes to the system or the experiment can be incorporated as simple modifications to the existing models. So physical models can be more versatile in their application.

There is also a limit as to how well the types of data-based models that have been hitherto employed can describe a complex system such as those encountered here. The investigations in the previous chapters have utilised only linear time-invariant black-box models for the electronic nose data. The systems in question are essentially nonlinear in their behaviour. Thus it would be naïve to expect a linear model to perform well over a range of situations. Nonlinear black-box modelling techniques could be applied with potentially more success than their linear counterparts, but there would still be the issue of deciding upon the functional forms to employ, which would have to be done exhaustively without any physical knowledge to provide guidance.

The use of an accurate physical model to produce a mathematical model for the system could be considered a short-cut to the best (or at least a much better) modelling outcome. The systems to be modelled are of course not actually black-boxes, where the processes occurring are hidden and unknown to the modellers, rather they are physical and chemical processes which are (to some extent) understood and so can be described by differential or algebraic equations. This knowledge can be employed to produce models which are not only superior in terms of their performance in use but also in their extensibility and versatility, and which contain parameters of physical significance.

6.2 Nose system models

With the black-box models of the previous chapters it was possible to deal exclusively in what were referred to as inverse models for the system (see Section 3.1), which are

then immediately usable as an odour classification system. However, with physical models it is not (normally) possible to jump straight to an inverse model – instead forward models for the system are proposed and validated. These forward models can then be utilised within an odour classification system in a number of ways which are discussed later in this chapter.

An electronic nose system, although essentially a simple device, consists of a number of different components which determine the relationship between the odour input at one end of the system, and the data output (either raw data or an odour class in the case of a complete classifying nose system) at the other end. A forward dynamic model for an electronic nose system must include models of the four major components of the nose system detailed in Figure 6.1. These submodels must then be combined into a full system model that describes the behaviour of the nose system.




Figure 6.1 Diagram showing the four main components of a (physical) forward model of an electronic nose system.

The roles of these four main components in both the physical nose system and models for it are discussed below.

6.2.1 Odour generation

The first component in a model for an electronic nose system must deal with how the odour to be examined is produced. The particular application or laboratory operating conditions will determine the way in which the odour to be sampled is generated and maintained.

The simplest model for the odour generation stage of the process would be one in which the odour sample is drawn from an effectively limitless supply of a sample that produces a constant odour concentration. Thus the strength and composition of the odour is constant with respect to both time and other variables such as the ambient temperature or temperature of the sample (which may or may not be the same).

The above physical model may be sufficiently realistic for some applications, but for many it would not. In environmental monitoring, the sample is often drawn from an effectively very large, well-mixed headspace (e.g. the air outside). Thus the removal of the (small) sample will have little or no effect on the concentration of the remaining odour gases in the headspace. However, in laboratory tests, samples are often held within fairly small sealed containers, so the removal of the sample can significantly affect the concentration of the odour remaining in the sample container (see Figure 6.2 for plots illustrating this issue). If the flow rate of the sample delivery system is small relative to the volume of the headspace above the sample, and the sample itself is fairly volatile at the operating temperature used, then this may not be a problem, since the headspace may remain almost entirely saturated (known as a static headspace). However, this will not always be the case, and if it is not, then the resulting dynamic variations in the concentration of the gases in the headspace must be modelled. In the extreme, the concentration falls to a lower limit known called the dynamic headspace (see Figure 6.2d).




Figure 6.2 Illustrative plots of artificial data to show the vital differences between sampling from an (ideal) static headspace (plots (a), (c) and (e)) and sampling from a headspace which is significantly altered by the outflow of gases initiated by the sampling process itself (plots (b), (d) and (f)).

In most applications, over short timescales, the ambient temperature and temperature of the sample (often more or less the same, certainly related to some extent) would not be expected to vary significantly. However, for longer term experiments, variations in the temperature might be significant, and may strongly affect the rate of evaporation and saturation concentration of the odour gases. In this case, a model for the odour generation which includes these variations must be used.

In order to show the range of parameters or variables upon which the odour headspace concentrations depends we state that a general form for a model for the odour generation stage of the process might be as follows:

$$C_h = f_1\left(t; T_s, T_a, \dot{Q}_h, \underline{p}_a, \underline{p}_s\right), \quad (6.1)$$

where C_h is the concentration of analyte in the headspace of the sample vessel (more specifically: at the gas sample extraction point in the sample vessel), f_1 is some function, T_s is the temperature of the sample, T_a is the ambient temperature, \dot{Q}_h is the gas flow rate through the headspace of the sample, \underline{p}_a is a vector of parameters relating to the physical properties of the analyte and \underline{p}_s is a vector of parameters relating to the physical dimensions of the sample vessel. The parameters T_s , T_a and \dot{Q}_h may be time varying, whilst the parameter vectors \underline{p}_a and \underline{p}_s will be constant for a given analyte and physical system.

So, for example, the simplest usable special case of equation 6.1 would be the case where the headspace is static:

$$C_h = k_1, \quad (6.2)$$

where k_1 is a constant. For this model, the parameters T_s , T_a and \dot{Q}_h are assumed to be either constant for the duration of the experiment, or at least to have no discernible effect on the value of C_h . As mentioned above, these assumptions may be valid given certain operating conditions. This simple model corresponds to the ideal (static) headspace scenario depicted in Figure 6.2(a).

If the analyte is not very volatile at the operating temperature, or the flow rate is high relative to the volume of the sample vessel then the headspace of the sample will not remain saturated once the gas system valves are operated and the odour gas(es) are drawn from the headspace. In this case, a simple dynamic model for the odour generation may be employed to describe how the headspace concentration might vary with time as the odour is sampled:

Assume that the headspace is initially saturated (concentration C_{hs}), and the flow rate through the headspace is zero. Assume that at $t = 0$, the system valves are switched over and the flow rate jumps instantly to a new constant rate \dot{Q}_h . The volume of the

headspace is assumed to remain fixed throughout (so the volume of liquid sample is not substantially changed by the evaporation from it). The temperature of the sample (including headspace) is also assumed to be constant for the duration of the experiment. The rate of evaporation is assumed to be proportional to the difference between the concentration of odour gas in the headspace and the reference saturation concentration C_{hs} (with constant of proportionality k_e). The gases in the headspace are assumed to be well mixed at all times. Then, the concentration of the odour gases C_h in the headspace could be modelled by the expression:

$$C_h(t) = C_{hs} \left\{ \frac{k_e}{\dot{Q}_h + k_e} + \left(1 - \frac{k_e}{\dot{Q}_h + k_e} \right) \exp\left(-\{\dot{Q}_h + k_e\}t\right) \right\}. \quad (6.3)$$

See Appendix 1 for the derivation of this expression as the solution of a first order ODE. The form of this expression matches the shape of the curve given in Figure 6.2(b).

Note that the assumptions above are highly valid for the experimental test system used to collect data for the model validation in the subsequent chapters. The choice between the models of equation (6.2) and (6.3) must be made based mainly on considerations of desired model simplicity, and volatility of analyte at the operating temperature.

However, it should also be noted that without some external means of identifying the parameters C_{hs} and k_e (e.g. from the chemistry literature, or via mass spectrometry experiments) the above model may be of limited practical use since the nose system cannot be used to measure C_h directly. Hence the parameters would need to be estimated concurrently with those in the other submodels of the full system model being considered.

For further extension of the model to allow greater versatility of application, it would be necessary to incorporate the effects of temperature variations on the parameters. Clearly the saturated headspace concentration value C_{hs} and the evaporation coefficient k_e could be expected to depend quite significantly on the temperature of the sample, and also perhaps the ambient temperature (since this would dictate the temperature of

the carrier gas). For example, it could be assumed that the saturated headspace concentration C_{hs} is approximately given by

$$C_{hs} = A \exp\left(-\frac{H_{m,e}}{RT_s}\right), \quad (6.4)$$

where A is a constant, $H_{m,e}$ is the molar enthalpy of evaporation of the analyte, R is the gas constant, and T_s is again the (absolute) sample temperature.

6.2.2 Odour transport

The second component in a model of an electronic nose system must deal with how the odour is transported from the sample to the odour sensors. In some systems this might be simple to deal with, since the sensors might be physically very near to the sample, and the transfer of the odour from the sample to the sensors may be almost instantaneous. In other systems (such as the system described in Chapter 5), the odour gases are transported to the sensors via a system of pipes and computer-controlled valves, culminating in a sensor chamber having a significant volume. A system such as this will certainly introduce a time delay into the system, which must be modelled, and will be dependent on the gas flow rate through the system. A more accurate model for the odour transport system would include mixing and diffusion effects in the pipe (at the odour ‘front’) and in the sensor chamber.

Thus a general expression for the odour transport stage might be:

$$C_a = f_2\left(t, C_h(t); \dot{Q}_h, \dot{Q}_c, T_a, T_s, T_c, \underline{p}_a, \underline{p}_g\right), \quad (6.5)$$

where C_a is the concentration of analyte in the air immediately above the sensor surface, $C_h(t)$ is the concentration of analyte drawn from the sample vessel (see previous section), \dot{Q}_h is the gas flow rate through the headspace of the sample vessel, \dot{Q}_c is the gas flow rate through the sensor chamber, T_a is the ambient temperature, T_s is the sample temperature, T_c is the chamber temperature, \underline{p}_a is a vector of parameters

relating to the physical properties of the analyte and \underline{p}_g is a vector of parameters relating to the physical dimensions of the various components of the gas delivery system. The parameters \dot{Q}_h , \dot{Q}_c , T_a , T_s and T_c may be time varying, whilst the parameter vectors \underline{p}_a and \underline{p}_g will be constant for a given analyte and physical system.

The simplest special case of equation (6.5) would be that where C_a is equal to C_h , i.e.

$$C_a(t) = C_h(t). \quad (6.6)$$

This would be unrealistic except under very specific conditions (for example when the temperatures are constant and the flow rates are constant and high (so that there is no detectable delay induced by the time taken for the odour front to travel along the gas system), and the physical characteristics of the system were such that no significant diffusion effects could be observed as the odour front arrives at the sensor array). A more realistic (and still simple) model would be to incorporate a time delay to reflect the physical characteristics of the system thus:

$$C_a(t) = C_h(t - \tau). \quad (6.7)$$

Now, equation (6.7) more accurately describes the characteristics of the odour transport system but the value of the delay τ must depend not only on fixed (for a given experiment) aspects of the system such as the lengths of pipework and analyte being sampled, but also on variables (or at least time-varying parameters) such as the flow rates and temperatures in the system. If diffusion at the odour front in the pipework was neglected, then the time delay τ might be expected to be inversely proportional to the gas flow rate through the system. However, due to the specific physical characteristics of the system in question, the situation is not quite so straightforward, as explained below. See Figure 6.3 for a diagram of the gas system used in the electronic nose system (described in greater detail in Chapter 5).

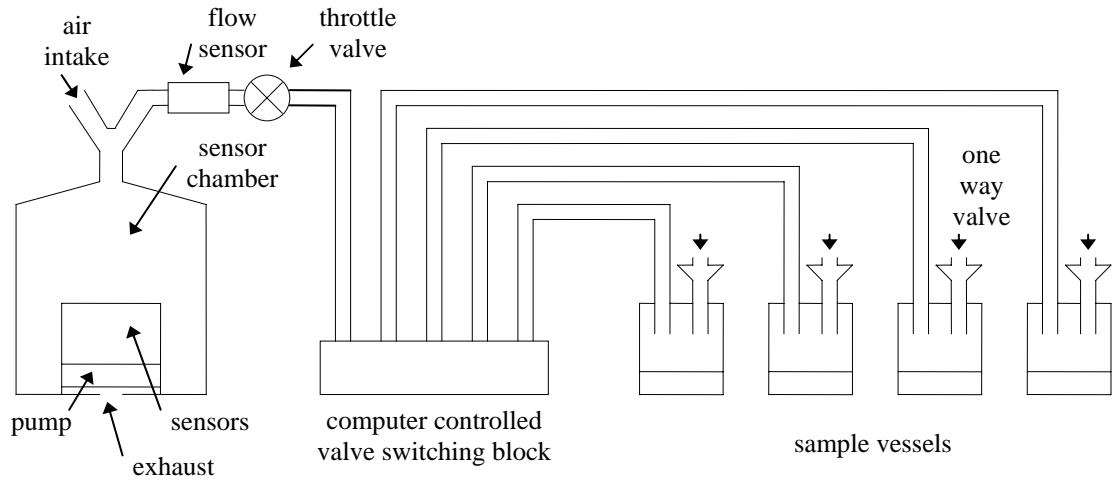


Figure 6.3 Schematic of the gas system of the electronic nose.

There are three different sections of the gas system, which have different gas flow rates. Firstly there is the main air intake (upper left in Figure 6.3), the flow rate through which is denoted by \dot{Q}_i . Next there is the gas flow through the sensor chamber itself, denoted by \dot{Q}_c , and lastly the flow rate through the sample delivery system (and hence headspace of sample), \dot{Q}_h .

Thus the time delay τ must be split into two components, the time taken for the odour front to travel through the headspace section (from the sample vessel to the ‘Y’ piece above the sensor chamber), τ_h , and the time taken for the odour front to pass through the chamber section (from the ‘Y’ piece to the sensors themselves), τ_c . Then the total time delay is given by

$$\tau = \tau_h + \tau_c. \quad (6.8)$$

Now, if the volume of the headspace section (V_h) and the (effective) volume of the chamber section (V_c) are known, then the time delay τ can be written in terms of these and the relevant flow rates:

$$\tau = \frac{V_h}{\dot{Q}_h} + \frac{V_c}{\dot{Q}_c}. \quad (6.9)$$

To (at least partially) confirm this relationship, an experiment was performed using the experimental test rig detailed in Chapter 5. The headspace flow rate \dot{Q}_h was varied from 10 to 95 ml/min and the corresponding delays (values of τ) were recorded. Note that the headspace flow rate (\dot{Q}_h) was varied by the adjustment of the throttle valve shown in Figure 6.3. The chamber flow rate (\dot{Q}_c) was assumed constant. This assumption is valid since estimates for τ_c vary only between about 0.86s and 0.94s as \dot{Q}_h varies between zero and 100 ml/min, due to the nature of the gas system (i.e. when the throttle valve is closed, more gas is simply sucked in through the air intake instead of through the odour delivery system). Thus the assumed relationship between the time delay τ and the headspace flow rate \dot{Q}_h reduces to

$$\tau = \frac{V_h}{\dot{Q}_h} + \tau_c. \quad (6.10)$$

The data recorded are plotted in Figure 6.4, along with a curve fit of the form in equation (6.10), produced using *Matlab*. The estimated value for V_h is 69.41ml, and the value for τ_c is 0.8709s. The curve fits the experimental data reasonably well, indicating that for the particular set-up (in terms of equipment used and flow rates selected), a very simple model for the relationship between the flow rate and time delay is acceptable.




Figure 6.4 Plot showing the relationship between the headspace flow rate \dot{Q}_h and the observed time delay in the odour transport system.

Equations (6.6) and (6.7) above are based on the assumption that the odour gas travels through the system via ‘plug flow’, i.e. a sharply defined odour ‘front’ travels through the pipework, valves and sensor chamber at a rate determined by the flow rate of the carrier gas. Thus it is assumed that any diffusion of this odour front within the odour transport system is negligible. At high flow rates this may seem sensible, since the transport of the odour front due to the pumping of the carrier gas will be occurring much faster than any diffusion effects. However, at low flow rates, this may not be the case.

The diffusion in the pipework at the (travelling) odour front can be modelled using the one-dimensional diffusion (partial differential) equation:

$$\frac{\partial u}{\partial t} = c^2 \frac{\partial^2 u}{\partial x^2}, \quad (6.11)$$

where u is the odour concentration, t is time, x is position along the pipework (measured from an origin which is in fact moving along the pipe at constant speed, corresponding to the position of the odour front if it were travelling via plug flow) and c is a diffusion coefficient, which will depend upon the analyte, the carrier gas and the temperatures of the various components of the system. Note that any friction with the walls of the pipework is neglected, as is any mixing through eddy currents / mixing etc., and the pipe is assumed to be of infinite length in both directions.

Then, with the initial condition

$$u(x,0) = \begin{cases} 1 & x \leq 0 \\ 0 & x > 0 \end{cases} \quad (6.12)$$

the solution of equation (6.11) (for $t > 0$) is given by [Kreyszig, 1988, p. 673]

$$u(x,t) = \frac{1}{\sqrt{\pi}} \int_{-\infty}^{-\frac{x}{2c\sqrt{t}}} \exp(-z^2) dz. \quad (6.13)$$

Or, equivalently

$$u(x,t) = \frac{1}{2} \left(1 + \operatorname{erf} \left(-\frac{x}{2c\sqrt{t}} \right) \right), \quad (6.14)$$

where erf is the error function (and $x, t > 0$).

Some plots of equation (6.14) are shown in Figure 6.5, for various times, starting at $t = 0$, where the odour profile (plot of odour intensity (or concentration) against distance along the pipe) is sharply defined, and evolving to more rounded profiles as time goes on.




Figure 6.5 Plot showing numerical solutions of the diffusion equation at various times. The plots correspond to odour concentration profiles at the odour ‘front’ in the pipework, evolving from a sharp ‘plug flow’ situation at $t = 0$ to a shallower odour gradient at the front at later times.

The model for diffusion in the pipework could be used to more accurately model a simple delay in the system (in situations where a simple (i.e. not involving diffusion) model does not suffice), or alternatively to provide a dynamic model for the odour concentration arriving at the sensor materials, which could be considered as an input to the models described in Section 6.2.3 below for the sensor responses.

Note that for complex odours (i.e. those containing more than one chemical species), each component may have a slightly different diffusivity. So each component could have a slightly different odour profile at the odour ‘front’. However, for the purposes of the modelling here, each species will be assumed to have the same diffusion constant. This is for simplicity and due to the difficulty of accurately estimating the diffusion constants.

6.2.3 Sensor response

The third component of a full system model would be a model for the response of the sensor array when exposed to the odour gases. This must consist in some way of a model for each of the sensors in the array, although each model may have an identical structure. The models for the responses of each sensor could, in their simplest form, depend only upon the concentration and composition of the odour being presented to the sensor. However, such a simple model could not be expected to function effectively in situations where other important variables are not held constant. So a comprehensive (and thus versatile) model for the sensor response must include such variables as the sensor operating temperature, the temperature of the sample gas, the humidity and flow rate of the sample carrier gas, and perhaps even the history of the sensor (i.e. the odours which it had previously been exposed to, some of which may have caused irreversible or slowly reversible changes in the sensor properties).

A general expression for the sensor response (variation in conductance) might thus be given by:

$$\underline{G} = \underline{f}_3 \left(t, C_a(t); \dot{Q}_c, T_a, T_s, T_c, \underline{p}_a, \underline{p}_g, \underline{p}_s \right), \quad (6.15)$$

where \underline{G} is the vector of sensor conductances, $C_a(t)$ is the concentration of analyte in the air immediately above the sensor surface, \dot{Q}_c is the gas flow rate through the sensor chamber, T_a is the ambient temperature, T_s is the sample temperature, T_c is the chamber temperature, \underline{p}_a is a vector of parameters relating to the physical properties of the analyte, \underline{p}_g is a vector of parameters relating to the physical dimensions of the various components of the gas delivery system, \underline{p}_s is an array of parameters relating to the physical characteristics and past history of the sensor array. The parameters \dot{Q}_c , T_a , T_s and T_c may be time varying, as may the parameter array \underline{p}_s , whilst the parameter vectors \underline{p}_a and \underline{p}_g will be constant for a given analyte and physical system.

Possible forms for the function \underline{f}_3 are many, and are the subject of subsequent chapters.

6.2.4 Interface electronics

The final component of the full system model for an electronic nose system is the model for the interface electronics. This section provides the link between the electrochemical changes occurring at the odour sensors and the data produced and recorded by the computer controlling the system.

Typically (and in the case of the particular system considered here) the interface electronics are essentially very simple. They are designed to produce voltage outputs (which are sent to a data acquisition card in a PC) which are proportional to the resistances or conductances of the sensors. Electronic filters are used to reduce noise in the circuits.

A general model for this final stage of the system could be given by the expression

$$\underline{V} = \underline{f}_4(t, \underline{G}; \underline{p}_e), \quad (6.16)$$

where \underline{V} is the vector of voltage outputs from the system, \underline{G} is the vector of sensor conductances and \underline{p}_e is a vector of parameters relating to the design and set-up of the interface electronics.

The simplest (and perhaps only useful) model for the interface electronics would assume that the signal outputs of the (whole) system are directly proportional to the resistances of the sensors (and thus to the reciprocal of the conductances of the sensors), so that

$$V_i(t) = \frac{1}{k_i G_i(t)}, \quad (6.17)$$

for $i = 1$ to n , where n is the number of sensors in the array, and the k_i 's are constants.

This may not be exactly true (since there are filters present in the electronic interface, as well as sources of noise), but there may not be much (if anything) to be gained from attempting to model the characteristics of the interface system. It is reasonable to assume that the delays and phase shifts of the electronics are negligible compared with the effects of the gas transport and sensor responses.

6.3 Using (forward) models in an odour classification system

The modelling process is not performed out of interest alone. The aim of this project is to produce algorithms or methods that can be utilised within an odour classification system. So whilst it is tempting to write down complex models incorporating terms and parameters to incorporate every conceivable mechanism affecting the system outputs, this is of little or no practical use. The models that are produced must in some way be usable in practice.

There are a number of ways in which (forward) parametric models might be made useful for classification. The obvious way would be to invert the forward models to produce inverse models like those utilised in Chapters 3 and 4. These inverse models could then be used directly to produce odour classifications from nose system outputs. This is a fine idea in theory, though in practice only the very simplest of models tend to be invertible. There are two main approaches that could be taken to tackling the problem of producing a physically-motivated inverse model.

The first approach would be to produce a system of differential equations describing the forward dynamics of the system, to solve the equations, then invert the solution to produce an analytical solution to the inverse modelling problem. This approach is unlikely to yield success here, since the models that will be produced will necessarily be systems of nonlinear ordinary differential equations. It is rare for such systems to be analytically solvable, and even rarer that solutions can be inverted.

The second approach would be to construct the model in an inverse orientation from the start, attempting to treat each of the components of a system model in turn, and model each backwards. This would be perfectly possible for some of the individual components of the full system (e.g. the interface electronics), but would be very difficult indeed for anything but the very simplest sensor response models.

Inverting a forward model to produce an inverse model is not the only way of using it in a classification system. A more promising method is via parameter estimation techniques. Some of the parameters in a physical model will relate to the nose system itself (and thus be the same regardless of the analyte tested), and others will depend also on the interactions between the analyte and the sensors. Those in the latter category might be used to distinguish between different odours, provided that they can be extracted from system output data.

Finally, even if it proves difficult or impossible to use forward models directly in the ways discussed above, the modelling process itself may provide insights into the ways in which the known physical mechanisms affect the outputs from nose systems. These insights may be useful for modification of the existing data / signal processing techniques in use.

6.4 Identifiability and estimation of the parameters in ODE models

As discussed in the previous section, the reliable estimation of the parameters in an ODE model is vital to the model's usefulness in a classification system. The question of whether the parameters in a model can be estimated from experimental data (assuming that perfect noise free data are available) is addressed by the field of structural identifiability analysis. Once the parameters in a model are shown to be identifiable (i.e. they can in theory be estimated from the experiment performed), the problem of numerically estimating them in practice arises.

Once a mathematical model for the electronic nose system is postulated, it is possible to use various different computer software packages, such as *Facsimile* (for details see <http://www.mcpa-software.com>), *Berkeley Madonna* (for details see

<http://www.berkeleymadonna.com>) or *Matlab* (for details see <http://www.mathworks.com>) to produce estimates for the parameters which best fit a sample of experimental data. However, these packages use optimisation / search algorithms to find a set of parameters that at least locally produces the best fit to the data. Thus they will invariably produce an answer, however there is no way of knowing whether or not this is the (globally) best parameter set to fit the data, or indeed if a single best parameter set exists. This is where identifiability analysis comes in.

The concept of identifiability was first formalised in 1970 (Bellman and Åström 1970). Since then it has been formalised and studied by many. Broadly speaking, an identifiability analysis is some form of analysis which seeks to determine whether the parameters in a model could be estimated from particular experiments (assuming that perfect noise-free data were available from such experiments). There are a variety of methods employed to determine the identifiability of a model. Examples include the Laplace transform method (Godfrey and Chapman 1990) (for linear systems only), the Taylor series approach (Godfrey *et al.* 1982) (for linear and nonlinear systems) and the similarity transform approach (Vajda *et al.* 1989) (for linear and nonlinear systems).

Regardless of the method used, the aim is to classify the model into one of three possible cases (Godfrey 1983):

1. For any (with the exception of at most countably many) possible parameter set, the input-output behaviour produced by the model is unique to that parameter set. Thus there exists a single set of parameters which produce the input-output behaviour that can be observed from the experiment. Thus the parameters of the model can all be estimated uniquely from (perfect) experimental data. The model is then said to be globally identifiable.
2. There exists a countable set of parameter sets which would produce the input-output behaviour which could be observed from the experiment. In this case the model is said to be locally identifiable.
3. There exist uncountably many parameter sets which would produce the input-output behaviour which could be observed from the experiment. The model is then said to be unidentifiable.

Clearly, if a model is unidentifiable, then it is of no practical use. Similarly if it is only locally identifiable then it may be of little use, although sometimes slight modifications to the model (Chappell and Gunn 1998) or the experiment (or indeed the consideration of extra information) may rectify the situation. Thus an identifiability analysis of a model really should be considered an essential prerequisite to experiment design and parameter estimation. To proceed with a model when it is not known whether or not it is identifiable is largely pointless, since no real confidence can be placed in any parameter estimations produced.

Note that many forms of identifiability analysis do not take into account any limitations which may exist on the specific experiments which may be performed (in terms of the available inputs to the model). Rather they frequently assume the availability of all bounded measurable input functions. In the system considered here, this is definitely not the case. However, it is often possible to modify the analysis or make use of corollaries to the main theorems on structural identifiability to take this into account.

Chapter 7

A reaction-based model for MOS sensors

In this chapter, a realistic mathematical model is derived from an electrochemical reaction scheme. The reaction scheme used is known to be an accurate description of the mechanism for a simple analyte (in fact, this mechanism is for carbon monoxide (Windischmann *et al.* 1979)). Thus it is still a simplification of more complex reactions that occur with different analytes or combinations of analytes, especially so in the case of complex odours (which may contain a great number of different chemical species).

7.1 Chemical reaction equations

The electrochemical processes occurring within the bulk of a metal oxide film are described by the following reaction equations (Windischmann *et al.* 1979; Gardner and Bartlett 1999):



where O is oxygen, e^- is an electron, $[\cdot]$ is a vacant oxygen ion binding site (chemisorption) in the sensor material, $[O^-]$ is an occupied oxygen ion binding site, A is the analyte (target molecule) and AO is the product of a reaction between A and O (e.g. CO_2 if A is CO). The first stage in the reaction mechanism is known as electron abstraction, and the second as electron donation.

7.2 Model assumptions

In order to formulate a useful mathematical model based on the electrochemical reaction scheme, it is necessary to make certain assumptions about the sensors, the surrounding equipment (e.g. the gas delivery system, interface electronics etc), and the experiment.

In the context of the whole system model approach discussed in Chapter 6, the necessary modelling assumptions can be attributed to the various components of the full model.

Gas / odour generation and transport assumptions:

- Both O_2 and A are introduced to the chamber at piecewise constant rates (i.e. the gas flow rate into the chamber and its composition are constant (so the headspace composition is assumed static)).
- A fixed fraction of the gas in the chamber is removed per unit of time elapsed (i.e. the gas flow rate out of the chamber is fixed).
- The volume of the gas chamber is fixed (so the absolute numbers of gas molecules or concentrations of the gases may be used interchangeably).

Sensor model assumptions:

- The diffusion of the analyte A into the sensing material is sufficiently fast as to be considered instantaneous, and thus neglected.

- The volume of sensor material is fixed (so the absolute numbers of the species present in the bulk or concentrations of these species may be used interchangeably).
- The number of chemisorption sites available to the oxygen is fixed.

Interface electronics assumptions:

- The mechanism by which a change in the number of conducting electrons produces a measured change in the resistance of the sensor material requires modelling itself. This mechanism is described in section 7.3.
- The measured value is directly proportional to the resistance of the sensor material. i.e. the effects of any noise, filtering effects, delays etc. induced by the interface electronics are ignored.
- The only species observed is the number of conducting electrons in the sensor material.

The standard method of converting the above physical model into a mathematical model would involve applying the law of mass action to the electrochemical reaction scheme above. This would require an implied assumption that all of the gas molecules in the chamber are always available for reaction at the sensor surface / within the sensor material, and that the molecules are always well-mixed. This is clearly not an accurate description of what can occur in the chamber, since the sensor material occupies only a tiny fraction of the sensor chamber volume, and the gases are passed into the chamber at one end, and exhausted at the other.

There are a number of ways of dealing with this problem in the mathematical model. The most accurate would be to incorporate the spatial effects into the model with partial differential equations (possibly even stochastic equations) to describe the diffusion of molecules in the chamber and quantify the ‘availability’ of the molecules based upon their proximity to the sensor. This would produce a very complex model which would be difficult to work with. The objective is to produce a much simpler model, which could be feasibly used in practice without requiring excessive computing hardware and software.

The simplest way of dealing with the potential inaccuracy of the model produced by blindly applying the law of mass action would be simply to ignore it, and proceed regardless. However, even ignoring the distastefulness of applying assumptions which are known to be wrong, it rapidly becomes clear that for any sensible estimates for the parameters, the system of differential equations produced is too stiff to be simulated reliably. Hence the model is very difficult to use in practice.

A more plausible method would be to model the mixing and availability in the chamber via the following assumptions:

- The chamber can be thought of as being split into two different regions: the main chamber itself, and a smaller ‘local region’ in the immediate vicinity of the sensor material.
- Only molecules within the local region are available for reaction with / at the sensor material.
- Gas molecules within each region are well-mixed at all times.
- Gas molecules arriving into, and departing from, the chamber, do so only to / from the main chamber.
- Molecules pass from each region to the other at rates proportional to the concentration difference of the species in question between the two regions.
- The main chamber is very large compared with the local region, thus the concentration of gases in the main chamber is effectively unaffected by exchange of molecules with the local region.

These assumptions are illustrated in Figure 7.1. Note that so far no assumptions have been made about the shape of the concentration profile within the chamber as the experiment progresses. By the assumptions above, this is dependent only on the gas flow dynamics from the headspace to the chamber. This is intentional to permit flexibility later on.




Figure 7.1 Diagram illustrating the model assumption with regards to the partitioning of the sensor chamber into a main region and a local region around the sensor.

7.3 Mathematical model

Let:

$$\begin{aligned}
 \bar{x}_1 &= \# A \text{ in the local region,} \\
 \bar{x}_2 &= \# e^- \text{ in the sensor bulk,} \\
 \bar{x}_3 &= \# \frac{1}{2} O_2 \text{ in the local region,} \\
 \bar{x}_4 &= \# [O^-] \text{ in the local region,} \\
 \bar{x}_5 &= \# AO \text{ in the local region,} \\
 \bar{x}_6 &= \# [\cdot] \text{ in the sensor bulk and} \\
 \bar{x}_7 &= \# AO \text{ in the main chamber,}
 \end{aligned} \tag{7.3}$$

where # denotes 'number of'.

Now let the volume of the local region be vol_{lr} , the concentration of analyte in the main chamber be C_{an} and the concentration of oxygen $\left(\frac{1}{2}O_2\right)$ be C_{ox} . Then the number of analyte molecules passing into the local region from the main chamber in unit time is given by

$$\delta_1 \left(C_{an} - \frac{\bar{x}_1}{vol_{lr}} \right), \quad (7.4)$$

and the number of oxygen atoms passing into the local region from the main chamber per unit time is given by

$$\delta_2 \left(C_{ox} - \frac{\bar{x}_3}{vol_{lr}} \right), \quad (7.5)$$

where δ_1 and δ_2 are diffusion rate constants. The diffusion rates of the two species could be slightly different though they are assumed to be approximately equal. This allows simplification by setting

$$r_{an} = \delta_1 C_{an},$$

$$r_{ox} = \delta_2 C_{ox} \text{ and}$$

$$q = \frac{\delta_1}{vol_{lr}}. \quad (7.6)$$

Applying the law of mass action to the local region yields the differential equations describing the system (7.1) and (7.2):

$$\begin{aligned}
\dot{\bar{x}}_1 &= r_{an} - k_2 \bar{x}_1 \bar{x}_4 - q \bar{x}_1 \\
\dot{\bar{x}}_2 &= -k_{1f} \bar{x}_2 \bar{x}_3 \bar{x}_6 + k_{1b} \bar{x}_4 + k_2 \bar{x}_1 \bar{x}_4 \\
\dot{\bar{x}}_3 &= r_{ox} - k_{1f} \bar{x}_2 \bar{x}_3 \bar{x}_6 + k_{1b} \bar{x}_4 - q \bar{x}_3 \\
\dot{\bar{x}}_4 &= k_{1f} \bar{x}_2 \bar{x}_3 \bar{x}_6 - k_{1b} \bar{x}_4 - k_2 \bar{x}_1 \bar{x}_4 \\
\dot{\bar{x}}_5 &= k_2 \bar{x}_1 \bar{x}_4 + r_{cl} \bar{x}_7 - q \bar{x}_5 \\
\dot{\bar{x}}_6 &= -k_{1f} \bar{x}_2 \bar{x}_3 \bar{x}_6 + k_{1b} \bar{x}_4 + k_2 \bar{x}_1 \bar{x}_4 \\
\dot{\bar{x}}_7 &= -r_{cc} \bar{x}_7 + q \bar{x}_5.
\end{aligned} \tag{7.7}$$

The conductance of the sensor material is determined by the number of charge carriers, and their mobility, thus

$$G = \mu_e n e \kappa \tag{7.8}$$

where G is the conductance of the sensor, μ_e is the mobility of the conducting electrons, e is the charge on each electron, n is the number of available (i.e. charge-carrying) electrons and κ is a constant relating to the dimensions of the sensing material.

The variable n in equation (7.8) clearly relates directly to one of the species in the reaction scheme, and thus to one of the state variables in the corresponding mathematical model. However, the relative changes in n effected by exposure of the sensor to analytes are very small. The magnitude of the changes in the resistances of the sensors cannot be explained by the change in the number of charge carriers, thus they must be attributed to changes in the mobility of the electrons (Gardner 1991; Heiland 1982). The mobility of the electrons is strongly affected by even small changes in the number of charge carriers, via changes in the width of the depletion region at each grain boundary in the sensor material (Gardner 1991; Williams 1991). Thus the important term in equation (7.8) is μ_e , which is a function of the state variable x_2 . The n term in (7.8) is approximately constant. The recorded quantity is proportional to the resistance of the sensor. Consequently, the observation function can be written

$$h(x(t, p), p) = \frac{1}{\mu_e(x_2) n e \kappa}, \tag{7.9}$$

The function μ_e , modelling the dependence of the electron mobility on the free electron concentration, could take many functional forms. Frequently, it is expressed as an exponential function of the intergranular barrier height. The function μ_e will be approximated by a linear function:

$$\mu_e(x_2) = \bar{\alpha} + \bar{\beta}x_2 \quad (7.10)$$

This then gives equation (7.9) the form

$$h(x(t, p), p) = \frac{1}{\alpha + \beta x_2}, \quad (7.11)$$

for suitable α and β .

7.4 Decoupled model with only three equations

The system of six ODEs (7.7) can be reduced by the application of conservation rules (or equivalently by inspection of the equations):

- Conservation of charge: $\dot{\bar{x}}_2 + \dot{\bar{x}}_4 = 0$, so $\bar{x}_2 + \bar{x}_4 = c_{eox}$, a constant.
- Conservation of chemisorption sites: $\dot{\bar{x}}_4 + \dot{\bar{x}}_6 = 0$, so $\bar{x}_4 + \bar{x}_6 = N$, a constant.

This decouples two of the six ODEs. Notice that \bar{x}_5 and \bar{x}_7 appear nowhere in the differential equations for any of the other \bar{x}_i 's, and that neither are observed quantities.

Thus the equations for \bar{x}_5 and \bar{x}_7 are largely irrelevant, so can be ignored.

There remains a system of only three non-linear ODEs:

$$\begin{aligned} \dot{\bar{x}}_1 &= r_{an} - k_2(c_{eox} - \bar{x}_2)\bar{x}_1 - q\bar{x}_1 \\ \dot{\bar{x}}_2 &= k_{1b}(c_{eox} - \bar{x}_2) - k_{1f}\bar{x}_2\bar{x}_3(N - c_{eox} + \bar{x}_2) + k_2\bar{x}_1(c_{eox} - \bar{x}_2) \\ \dot{\bar{x}}_3 &= k_{1b}(c_{eox} - \bar{x}_2) - k_{1f}\bar{x}_2\bar{x}_3(N - c_{eox} + \bar{x}_2) + r_{ox} - q\bar{x}_3. \end{aligned} \quad (7.12)$$

7.5 Steady states of full model

To find the steady states for the model (7.12), set $\dot{\bar{x}}_1 = \dot{\bar{x}}_2 = \dot{\bar{x}}_3 = 0$ and then solve for \bar{x}_1 , \bar{x}_2 and \bar{x}_3 to give the steady states x_1^* , x_2^* and x_3^* .

Setting $\dot{\bar{x}}_1 = 0$ gives:

$$x_2^* = c_{\text{cox}} - \frac{1}{k_2} \left(\frac{r_{\text{an}}}{x_1^*} - q \right). \quad (7.13)$$

Setting $\dot{\bar{x}}_3 = 0$ and substituting for x_2^* gives:

$$x_3^* = \frac{r_{\text{ox}} + \frac{k_{\text{lb}}}{k_2} \left(\frac{r_{\text{an}}}{x_1^*} - q \right)}{k_{\text{1f}} \left(c_{\text{cox}} - \frac{1}{k_2} \left[\frac{r_{\text{an}}}{x_1^*} - q \right] \right) \left(N - \frac{1}{k_2} \left[\frac{r_{\text{an}}}{x_1^*} - q \right] \right) + q}. \quad (7.14)$$

Now, setting $\dot{\bar{x}}_2 = 0$ gives:

$$0 = \dot{\bar{x}}_1 + \dot{\bar{x}}_2 - \dot{\bar{x}}_3 = r_{\text{an}} - qx_1^* - r_{\text{ox}} + qx_3^*, \quad (7.15)$$

which rearranges to

$$x_3^* = \frac{r_{\text{ox}} - r_{\text{an}}}{q} + x_1^*. \quad (7.16)$$

Equating (7.16) with (7.14) gives

$$\frac{r_{\text{ox}} - r_{\text{an}}}{q} + x_1^* = \frac{r_{\text{ox}} + \frac{k_{\text{lb}}}{k_2} \left(\frac{r_{\text{an}}}{x_1^*} - q \right)}{k_{\text{1f}} \left(c_{\text{cox}} - \frac{1}{k_2} \left[\frac{r_{\text{an}}}{x_1^*} - q \right] \right) \left(N - \frac{1}{k_2} \left[\frac{r_{\text{an}}}{x_1^*} - q \right] \right) + q}, \quad (7.17)$$

which rearranges to give the cubic equation

$$\begin{aligned}
& \left\{ q + k_{1f} \left(c_{\text{eox}} + \frac{q}{k_2} \right) \left(N + \frac{q}{k_2} \right) \right\} x_1^{*3} + \\
& \left\{ \left(\frac{r_{\text{ox}} - r_{\text{an}}}{q} \right) \left[q + k_{1f} \left(c_{\text{eox}} + \frac{q}{k_2} \right) \left(N + \frac{q}{k_2} \right) \right] - k_{1f} \frac{r_{\text{an}}}{k_2} \left(c_{\text{eox}} + N + \frac{2q}{k_2} \right) + q \frac{k_{1b}}{k_2} - r_{\text{ox}} \right\} x_1^{*2} + \\
& \left\{ -k_{1f} \frac{r_{\text{an}}}{k_2} \left[\left(\frac{r_{\text{ox}} - r_{\text{an}}}{q} \right) \left(c_{\text{eox}} + N + \frac{2q}{k_2} \right) - \frac{r_{\text{an}}}{k_2} \right] - k_{1b} \frac{r_{\text{an}}}{k_2} \right\} x_1^* + \\
& k_{1f} \left(\frac{r_{\text{an}}}{k_2} \right)^2 \left(\frac{r_{\text{ox}} - r_{\text{an}}}{q} \right) = 0.
\end{aligned} \tag{7.18}$$

The solutions to equation (7.18) (and hence steady states of the system (7.12)) can be found, but are very large expressions and thus difficult to analyse. However, knowledge of the likely relative sizes of some of the parameters can give us some information about the nature of the steady states. The parameter r_{ox} will be greater than r_{an} , thus the constant coefficient (the last term in equation (7.18)) will be positive. The coefficient of x_1^{*3} is positive (since each of the parameters is positive). The signs of the coefficients of x_1^{*2} and x_1^* are not obvious. Hence Descartes rule of signs (Hall and Knight 1950, pp. 459-460) yields only that there can be either two or zero positive real roots of the cubic equation (7.18). Hence there are potentially two feasible steady states of the system (7.12).

7.6 Initial conditions as steady states from odourless system

In a typical experiment, the system is first allowed to reach an equilibrium whilst only air (the odour carrier gas) is pumped through the sensor chamber. Then valves are actuated to present the required odour to the sensors. The initial conditions for the system of ODEs (7.12) thus come from the steady states of the equivalent ODE system where there is no analyte present (i.e. the steady state where oxygen is present in the gas flow, but no odour, so that the chemical reaction described in equation (7.1b) is no longer occurring).

In the absence of analyte, equations (7.7) become:

$$\begin{aligned}
 \dot{\bar{x}}_2 &= -k_{1f}\bar{x}_2\bar{x}_3\bar{x}_6 + k_{1b}\bar{x}_4 \\
 \dot{\bar{x}}_3 &= r_{ox} - k_{1f}\bar{x}_2\bar{x}_3\bar{x}_6 + k_{1b}\bar{x}_4 - q\bar{x}_3 \\
 \dot{\bar{x}}_4 &= k_{1f}\bar{x}_2\bar{x}_3\bar{x}_6 - k_{1b}\bar{x}_4 \\
 \dot{\bar{x}}_6 &= -k_{1f}\bar{x}_2\bar{x}_3\bar{x}_6 + k_{1b}\bar{x}_4.
 \end{aligned} \tag{7.19}$$

Again (7.19) can be reduced using the conservation rules $\bar{x}_2 + \bar{x}_4 = c_{eox}$ and $\bar{x}_4 + \bar{x}_6 = N$. Then there remains a system of only two coupled differential equations:

$$\begin{aligned}
 \dot{\bar{x}}_2 &= -k_{1f}\bar{x}_2\bar{x}_3(N - c_{eox} + \bar{x}_2) + k_{1b}(c_{eox} - \bar{x}_2) \\
 \dot{\bar{x}}_3 &= r_{ox} - k_{1f}\bar{x}_2\bar{x}_3(N - c_{eox} + \bar{x}_2) + k_{1b}(c_{eox} - \bar{x}_2) - q\bar{x}_3.
 \end{aligned} \tag{7.20}$$

Now, to find the steady states of (7.20), set $\dot{\bar{x}}_2 = \dot{\bar{x}}_3 = 0$. The steady states x_2^0 and x_3^0 are:

$$\begin{aligned}
 x_3^0 &= \frac{r_{ox}}{q} \\
 x_2^{0\pm} &= \frac{1}{2} \left\{ c_{eox} - N - \frac{k_{1b}q}{k_{1f}r_{ox}} \pm \sqrt{\left(c_{eox} - N - \frac{k_{1b}q}{k_{1f}r_{ox}} \right)^2 + 4 \frac{k_{1b}c_{eox}q}{k_{1f}r_{ox}}} \right\}.
 \end{aligned} \tag{7.21}$$

However, $x_2^{0-} < 0$ (since $k_{1f}, k_{1b}, c_{eox}, r_{ox}$ and $q > 0$), so there is only one feasible steady state, given by:

$$\begin{aligned}
 x_3^0 &= \frac{r_{ox}}{q} \\
 x_2^0 &= \frac{1}{2} \left\{ c_{eox} - N - \frac{k_{1b}q}{k_{1f}r_{ox}} + \sqrt{\left(c_{eox} - N - \frac{k_{1b}q}{k_{1f}r_{ox}} \right)^2 + 4 \frac{k_{1b}c_{eox}q}{k_{1f}r_{ox}}} \right\}.
 \end{aligned} \tag{7.22}$$

These values (plus the initial condition $x_1^0 = 0$) provide the initial conditions for the full model in (7.12).

7.7 Model identifiability

An identifiability analysis of the model given by equations (7.12) using the Taylor series approach is difficult due to the complexity of the expression for the initial condition for \bar{x}_2 . Instead the similarity transform approach for non-linear systems is used, as described in Chappell *et al.* 1990.

In order to apply Corollary 1 of Vajda *et al.* 1989 later, it is required that the system has zero initial conditions. Thus, it is necessary to translate the state variables \bar{x}_2 and \bar{x}_3 .

Set

$$\begin{aligned}x_1 &= \bar{x}_1, \\x_2 &= \bar{x}_2 - x_2^0 \text{ and} \\x_3 &= \bar{x}_3 - x_3^0.\end{aligned}\tag{7.23}$$

The system can now be written in its general form as:

$$\begin{aligned}\dot{x}(t, p) &= f(x(t, p), p) + u(t)g(x(t, p), p) \\y(t, p) &= h(x(t, p), p) \\x(0, p) &= x_0(p)\end{aligned}\tag{7.24}$$

where $x = (x_1, x_2, x_3)$, $p = (k_{1f}, k_{1b}, k_2, r_{an}, r_{ox}, N, c_{ox}, q, \alpha, \beta)$. The functions f , g and h are given by:

$$\begin{aligned}
f(x(t, p), p) &= \begin{pmatrix} -k_2(c_{\text{eox}} - x_2 - x_2^0)x_1 - qx_1 \\ k_{1b}(c_{\text{eox}} - x_2 - x_2^0) - k_{1f}(x_2 + x_2^0)(x_3 + x_3^0)(N - c_{\text{eox}} + x_2 + x_2^0) + k_2x_1(c_{\text{eox}} - x_2 - x_2^0) \\ k_{1b}(c_{\text{eox}} - x_2 - x_2^0) - k_{1f}(x_2 + x_2^0)(x_3 + x_3^0)(N - c_{\text{eox}} + x_2 + x_2^0) - qx_3 \end{pmatrix} \\
g(x(t, p), p) &= \begin{pmatrix} r_{\text{an}} \\ 0 \\ 0 \end{pmatrix} \\
h(x(t, p), p) &= \frac{1}{\alpha + \beta x_2}.
\end{aligned} \tag{7.25}$$

(Note that the α in equation (7.25) is really different from that in equation (7.11) (in fact it takes the value $\alpha + \beta x_2^0$), but the symbol is kept the same for notational simplicity.)

For typical electronic nose experiments the function u might take the form of a step input for a particular duration of time, i.e.

$$u(t) = \begin{cases} 1 & 0 < t \leq T \\ 0 & T < t \end{cases}, \tag{7.26}$$

for some $T > 0$.

The similarity transform approach requires that the system in question be both controllable and observable (i.e. minimal). Lemma 1 in Vajda *et al.* 1989 states that if for a particular value p^* of the parameter vector p the system becomes linear, and if this linear system is controllable and observable, then so is the corresponding nonlinear system. For the system above, it is possible to choose a p^* so that the system is linear, however, this linear system is neither controllable nor observable. So the minimality of the system must be checked using the controllability and observability rank criteria (Chappell *et al.*, 1990).

7.7.1 Controllability rank criterion

The following definitions and results are taken directly from Chappell *et al.*, 1990.

For vector fields φ^1 and φ^2 , the Lie bracket $[\varphi^1, \varphi^2]$ is defined by:

$$[\varphi^1, \varphi^2](x) = \frac{\partial \varphi^2(x)}{\partial x} \varphi^1(x) - \frac{\partial \varphi^1(x)}{\partial x} \varphi^2(x), \quad (7.27)$$

where $\frac{\partial \varphi^i(x)}{\partial x}$ denotes the Jacobian matrix of φ^i , $i = 1, 2$. For a piecewise constant control u^i , the vector field φ^i is defined by:

$$\varphi^i(x) = f(x) + u^i g(x), \quad i = 1, 2, 3, \dots \quad (7.28)$$

Consider the Lie algebra \mathfrak{f} , which has elements which can be represented by finite linear combinations of elements of the form

$$[\varphi^1, [\varphi^2, [\dots [\varphi^{i-1}, \varphi^i] \dots]]]. \quad (7.29)$$

Note that g is in \mathfrak{f} . Let $\mathfrak{f}(x)$ denote the space of vectors spanned by the vector fields of \mathfrak{f} at x . The system of equations (7.24) is said to satisfy the controllability rank criterion (CRC) at x_0 if the dimension of $\mathfrak{f}(x_0)$ is n (n is the number of state variables, so in our case, $n = 3$) (Vajda *et al.* 1989)

To check that the CRC is satisfied for our system, consider piecewise constant controls $u^1 \neq u^2$, then (since for our system, g is constant):

$$\begin{aligned} [\varphi^1, \varphi^2](x) &= \frac{\partial f}{\partial x}(f(x) + u^1 g) - \frac{\partial f}{\partial x}(f(x) + u^2 g) \\ &= (u^1 - u^2) \frac{\partial f}{\partial x} g \\ &= r_{an} (u^1 - u^2) \begin{pmatrix} -k_2(c_{eox} - x_2 - x_2^0) - q \\ k_2(c_{eox} - x_2 - x_2^0) \\ 0 \end{pmatrix}. \end{aligned} \quad (7.30)$$

So

$$[\varphi^1, \varphi^2](x_0) = r_{an}(u^1 - u^2) \begin{pmatrix} -k_2(c_{eox} - x_2^0) - q \\ k_2(c_{eox} - x_2^0) \\ 0 \end{pmatrix}. \quad (7.31)$$

The first and second co-ordinates of $[\varphi^1, \varphi^2](x_0)$ are generically non-zero.

Now, taking a third piecewise constant control u^3 , we consider the Lie bracket

$$\begin{aligned} [\varphi^3, [\varphi^1, \varphi^2]](x) &= \frac{\partial [\varphi^1, \varphi^2]}{\partial x}(x) \varphi^3(x) - \frac{\partial \varphi^3}{\partial x}(x) [\varphi^1, \varphi^2](x) \\ &= r_{an} k_2 (u^1 - u^2) \left\{ \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & -1 & 0 \end{pmatrix} (f(x) + u^3 g) - \frac{\partial f}{\partial x}(x) \begin{pmatrix} x_2 + x_2^0 - c_{eox} - \frac{q}{k_2} \\ c_{eox} - x_2 - x_2^0 \\ 0 \end{pmatrix} \right\}. \end{aligned} \quad (7.32)$$

Substituting initial conditions and expanding gives

$$\begin{aligned} &[\varphi^3, [\varphi^1, \varphi^2]](x_0) \\ &= r_{an} k_2 (u^1 - u^2) \left(\begin{array}{c} -k_2 \left((c_{eox} - x_2^0) + \frac{q}{k_2} \right)^2 + r_{ox} + k_{1b} (c_{eox} - x_2^0) - k_{1f} x_2^0 x_3^0 (N + x_2^0 - c_{eox}) \\ c_{eox}^2 (k_2 - k_{1f} x_3^0) + c_{eox} (q - 2k_2 x_2^0 + k_{1f} x_3^0 (N + 2x_2^0)) - x_2^0 (q - x_2^0 x_3^0 (k_2 - k_{1f} x_3^0)) \\ (c_{eox} - x_2^0) (k_{1b} + k_{1f} x_3^0 (N - c_{eox} + 2x_2^0)) \end{array} \right) \end{aligned} \quad (7.33)$$

Note that $[\varphi^3, [\varphi^1, \varphi^2]](x_0)$ has (generically) non-zero components. Thus $\{g, [\varphi^1, \varphi^2], [\varphi^3, [\varphi^1, \varphi^2]]\}$ is a linearly independent set, and so $\dim(f(x_0))$ is 3. Thus the CRC is satisfied, and the system is controllable.

7.7.2 Observability rank criterion

Again, the following definitions and results are taken from Chappell *et al.*, 1990. Assuming the output function $h(x)$ is continuously differentiable in x , the Lie derivative of h along the vector field φ^i is defined by:

$$L_{\varphi^i}(h)(x) = dh(x) \cdot \varphi^i(x), \quad (7.34)$$

where $dh(x)$ denotes the gradient vector field:

$$dh(x) = \left(\frac{\partial h(x)}{\partial x_1}, \dots, \frac{\partial h(x)}{\partial x_n} \right). \quad (7.35)$$

Consider the space of 1-forms $d\mathbf{g}$ whose elements are finite linear combinations of elements of the form:

$$d(L_{\varphi^i}(L_{\varphi^{i-1}}(\dots(L_{\varphi^1}(h)(x))\dots))). \quad (7.36)$$

Note that $dh(x)$ is in $d\mathbf{g}$. Let $d\mathbf{g}(x)$ denote the space of vectors obtained by evaluating the elements of $d\mathbf{g}$ at x . The system of equations (7.24) is said to satisfy the observability rank criterion (ORC) at x_0 if the dimension of $d\mathbf{g}(x_0)$ is n (Vajda *et al.* 1989).

To check that the ORC is satisfied for the system in question, first consider $dh(x)$:

$$dh(x) = \left(0, -\frac{\beta}{(\alpha + \beta x_2)^2}, 0 \right). \quad (7.37)$$

Note that $dh(x_0) \neq 0$. Now, let $u^1 = 0$, so that $\varphi^1 = f(x)$, then

$$L_{\varphi^1}(h)(x) = dh(x) \cdot f(x) = -\frac{\beta}{(\alpha + \beta x_2)^2} \left\{ k_{1b} (c_{eox} - x_2 - x_2^0) - k_{1f} (x_2 + x_2^0) (x_3 + x_3^0) (N - c_{eox} + x_2 + x_2^0) + k_2 x_1 (c_{eox} - x_2 - x_2^0) \right\}. \quad (7.38)$$

Thus

$$d(L_{\varphi^1}(h)(x_0)) = \frac{\beta}{\alpha^2} \begin{pmatrix} -k_2 (c_{eox} - x_2^0) \\ \frac{1}{\alpha} [k_{1b} (\alpha + 2\beta (c_{eox} - x_2^0)) + k_{1f} x_3^0 ((\alpha - 2\beta x_2^0) (N - c_{eox} + 2x_2^0) + 2\beta x_2^{0^2})] \\ k_{1f} x_2^0 (N - c_{eox} + x_2^0) \end{pmatrix} \quad (7.39)$$

Now, all three components of $d(L_{\varphi^1}(h)(x_0))$ are generically non-zero, so $dh(x_0)$ and $d(L_{\varphi^1}(h)(x_0))$ are linearly independent.

Let $u^2 \neq 0$, then $\varphi^2 = f(x) + u^2 g$. Expressions for $L_{\varphi^2}(L_{\varphi^1}(h)(x))$, $d(L_{\varphi^2}(L_{\varphi^1}(h)(x)))$ and $d(L_{\varphi^2}(L_{\varphi^1}(h)(x_0)))$ are then readily obtained using *Mathematica* (see Appendix 2). The expressions are not displayed here since they are far too large to be informative.

In order to check that the ORC is satisfied, it must be shown that $dh(x_0)$, $d(L_{\varphi^1}(h)(x_0))$ and $d(L_{\varphi^2}(L_{\varphi^1}(h)(x_0)))$ are linearly independent. To do this it is only necessary to show that the ratios of the first to third components of $d(L_{\varphi^1}(h)(x_0))$ and $d(L_{\varphi^2}(L_{\varphi^1}(h)(x_0)))$ are not the same. It is easily checked using *Mathematica* that this is indeed the case (see Appendix 2) (in fact, the ratio of the components of $d(L_{\varphi^2}(L_{\varphi^1}(h)(x_0)))$ depends upon parameters which do not even appear in $d(L_{\varphi^1}(h)(x_0))$). Thus the three vectors are linearly independent, so their span has dimension three, and the ORC is satisfied. Thus the system is observable.

7.7.3 Identifiability analysis

Now that the minimality of the system has been established, the identifiability analysis can proceed. Let $u \in U[0, T]$, the set of bounded measurable functions defined on a time interval $[0, T]$. Let $p \sim \tilde{p}$ denote that the parameter values p and \tilde{p} in Ω are indistinguishable in the experiments $(x_0(p), U[0, T])$. The system is globally identifiable at $p \in \Omega$ if $p \sim \tilde{p}$ implies that $p = \tilde{p}$, and it is locally identifiable if there exists an open neighbourhood W of p in Ω such that $p \sim \tilde{p}$ for p in W implies that $p = \tilde{p}$. Now, in order to establish the identifiability of the system, use is made of the following theorem (statement taken from Chappell *et al.* 1990, proof to be found in Vajda *et al.* 1989).

Theorem 1

Assume that the system of equations (7.20) is locally reduced at $x_0(p)$ for almost all p in Ω . Consider the parameter values of p, \tilde{p} in Ω , an open neighbourhood V of $x_0(p)$ in \mathbf{R}^n , and an analytical mapping $\lambda: V \rightarrow \mathbf{R}^n$ defined on $V \subseteq \mathbf{R}^n$ such that

$$(i) \text{ Rank } \frac{\partial \lambda(\tilde{x})}{\partial \tilde{x}} = n \text{ for all } \tilde{x} \in V \quad (7.40)$$

$$(ii) \lambda(x_0(\tilde{p})) = x_0(p) \quad (7.41)$$

$$(iii) f(\lambda(\tilde{x}), p) = \frac{\partial \lambda(\tilde{x})}{\partial \tilde{x}} f(\tilde{x}, \tilde{p}) \quad (7.42a)$$

$$g(\lambda(\tilde{x}), p) = \frac{\partial \lambda(\tilde{x})}{\partial \tilde{x}} g(\tilde{x}, \tilde{p}) \quad (7.42b)$$

$$h(\lambda(\tilde{x}), p) = h(\tilde{x}, \tilde{p}) \quad (7.42c)$$

for all $\tilde{x} \in V$. Then there exists $T > 0$ such that the system of equations (7.24) is globally identifiable at p in the experiments $(x_0(p), U[0, T])$ if and only if conditions (i), (ii) and (iii) imply that $p = \tilde{p}$.

Now, since we do not have access to all bounded measurable controls for the experiment, Corollary 1 of Vajda *et al.* 1989 must be employed, which states that

Theorem 1 extends to systems with specified controls and zero initial conditions provided that the following conditions hold:

- (i) $g(x, p) = b(p)$; that is, the input is multiplied by a constant vector b .
- (ii) $f(0, p) = 0$; that is, with no input, the system stays at rest.
- (iii) $\lambda(\tilde{x}) = T\tilde{x}$, where T is an $n \times n$ constant non-singular matrix; that is, the equivalence transformations are linear and time-invariant.

The first condition does indeed hold, since in this case g depends only on the parameter r_{an} (see equations (7.25)). Similarly the second condition holds simply by the definitions of x_2^0 and x_3^0 , and the (translated from the original) variables x_2 and x_3 . The third condition requires somewhat more work.

Let $\lambda(\tilde{x}) = (\lambda_1(\tilde{x}), \lambda_2(\tilde{x}), \lambda_3(\tilde{x}))$, and notice that $\lambda(0) = 0$ (from equation (7.41)). Now, equation (7.42) immediately gives information about λ_2 :

$$\frac{1}{\alpha + \beta\lambda_2(\tilde{x})} = \frac{1}{\tilde{\alpha} + \tilde{\beta}\tilde{x}_2}. \quad (7.43)$$

Which rearranges to yield:

$$\lambda_2(\tilde{x}) = \frac{1}{\beta} [(\tilde{\alpha} - \alpha) + \tilde{\beta}\tilde{x}_2] \quad (7.44)$$

Now, since $\lambda_2(0) = 0$, equation (7.44) implies that $\tilde{\alpha} = \alpha$ and hence

$$\lambda_2(\tilde{x}) = \frac{\tilde{\beta}}{\beta} \tilde{x}_2 \quad (7.45)$$

So the second row of the Jacobian $\frac{\partial\lambda(\tilde{x})}{\partial\tilde{x}}$ (which will eventually be the T matrix mentioned in the conditions for Corollary 1, once the linearity of λ is established) is

known. To find further elements of this Jacobian, consider equation (7.42b), which gives that:

$$\begin{pmatrix} r_{an} \\ 0 \\ 0 \end{pmatrix} = \frac{\partial \lambda(\tilde{x})}{\partial \tilde{x}} \begin{pmatrix} \tilde{r}_{an} \\ 0 \\ 0 \end{pmatrix}. \quad (7.46)$$

The first component of this equation shows that

$$\frac{\partial \lambda_1}{\partial \tilde{x}_1} = \frac{r_{an}}{\tilde{r}_{an}}. \quad (7.47)$$

The second component of equation (7.46) yields no further information. However the third component gives that:

$$\frac{\partial \lambda_3}{\partial \tilde{x}_1} = 0, \quad (7.48)$$

which completes the first column of the Jacobian.

Now consider the second component of equation (7.42):

$$\begin{aligned} & k_{1b}(c_{eox} - \lambda_2 - x_2^0) - k_{1f}(\lambda_2 + x_2^0)(\lambda_3 + x_3^0)(N - c_{eox} + \lambda_2 + x_2^0) + k_2(c_{eox} - \lambda_2 - x_2^0)\lambda_1 \\ &= \frac{\tilde{\beta}}{\beta}(\tilde{k}_{1b}(\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) - \tilde{k}_{1f}(\tilde{x}_2 + \tilde{x}_2^0)(\tilde{x}_3 + \tilde{x}_3^0)(\tilde{N} - \tilde{c}_{eox} + \tilde{x}_2 + \tilde{x}_2^0) + \tilde{k}_2\tilde{x}_1(\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0)) \end{aligned} \quad (7.49)$$

On substitution for λ_2 (from equation (7.45)), this becomes

$$\begin{aligned} & k_{1b}\left(c_{eox} - \frac{\tilde{\beta}}{\beta}\tilde{x}_2 - x_2^0\right) - k_{1f}\left(\frac{\tilde{\beta}}{\beta}\tilde{x}_2 + x_2^0\right)(\lambda_3 + x_3^0)\left(N - c_{eox} + \frac{\tilde{\beta}}{\beta}\tilde{x}_2 + x_2^0\right) + k_2\lambda_1\left(c_{eox} - \frac{\tilde{\beta}}{\beta}\tilde{x}_2 - x_2^0\right) \\ &= \frac{\tilde{\beta}}{\beta}\left(\tilde{k}_{1b}(\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) - \tilde{k}_{1f}(\tilde{x}_2 + \tilde{x}_2^0)(\tilde{x}_3 + \tilde{x}_3^0)(\tilde{N} - \tilde{c}_{eox} + \tilde{x}_2 + \tilde{x}_2^0) + \tilde{k}_2(\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0)\tilde{x}_1\right) \end{aligned} \quad (7.50)$$

Now, differentiating (7.50) with respect to \tilde{x}_1 gives

$$\begin{aligned} & -k_{1f} \left(\frac{\tilde{\beta}}{\beta} \tilde{x}_2 + x_2^0 \right) \left(N - c_{eox} + \frac{\tilde{\beta}}{\beta} \tilde{x}_2 + x_2^0 \right) \frac{\partial \lambda_3}{\partial \tilde{x}_1} + k_2 \left(c_{eox} - \frac{\tilde{\beta}}{\beta} \tilde{x}_2 - x_2^0 \right) \frac{\partial \lambda_1}{\partial \tilde{x}_1} \\ & = \frac{\tilde{\beta}}{\beta} \tilde{k}_2 (\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) \end{aligned} \quad (7.51)$$

However, from equation (7.47) we know that $\frac{\partial \lambda_1}{\partial \tilde{x}_1} = \frac{r_{an}}{\tilde{r}_{an}}$, and from equation (7.48) we

know that $\frac{\partial \lambda_3}{\partial \tilde{x}_1} = 0$, so (7.51) becomes

$$k_2 \frac{r_{an}}{\tilde{r}_{an}} \left(c_{eox} - \frac{\tilde{\beta}}{\beta} \tilde{x}_2 - x_2^0 \right) = \frac{\tilde{\beta}}{\beta} \tilde{k}_2 (\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) \quad (7.52)$$

Now equating coefficients of \tilde{x}_2 gives

$$k_2 r_{an} = \tilde{k}_2 \tilde{r}_{an}. \quad (7.53)$$

Also equating constant terms in (7.52) gives

$$(c_{eox} - x_2^0) = \frac{\tilde{\beta}}{\beta} (\tilde{c}_{eox} - \tilde{x}_2^0). \quad (7.54)$$

Since $f(0) = 0$ (from the construction of the state variables), we have that

$$k_{1b} (c_{eox} - x_2^0) - k_{1f} x_2^0 x_3^0 (N - c_{eox} + x_2^0) = 0, \quad (7.55)$$

which rearranges to give

$$(c_{eox} - x_2^0) (k_{1b} + k_{1f} x_2^0 x_3^0) = k_{1f} x_2^0 x_3^0 N. \quad (7.56)$$

Now, $k_{1f}x_2^0x_3^0N > 0$, and $(k_{1b} + k_{1f}x_2^0x_3^0) > 0$, so it follows that

$$(c_{eox} - x_2^0) > 0, \quad (7.57)$$

and hence (from equation (7.55))

$$(N - c_{eox} + x_2^0) > 0. \quad (7.58)$$

Now, substituting $\lambda(0) = 0$ into equation (7.50) gives

$$k_{1b}(c_{eox} - x_2^0) - k_{1f}x_2^0x_3^0(N - c_{eox} + x_2^0) = \frac{\tilde{\beta}}{\beta} \left(\tilde{k}_{1b}(\tilde{c}_{eox} - \tilde{x}_2^0) - \tilde{k}_{1f}\tilde{x}_2^0\tilde{x}_3^0(\tilde{N} - \tilde{c}_{eox} + \tilde{x}_2^0) \right), \quad (7.59)$$

which implies that (using equation (7.55))

$$\tilde{k}_{1b}(\tilde{c}_{eox} - \tilde{x}_2^0) - \tilde{k}_{1f}\tilde{x}_2^0\tilde{x}_3^0(\tilde{N} - \tilde{c}_{eox} + \tilde{x}_2^0) = 0 \quad (7.60)$$

since $\frac{\tilde{\beta}}{\beta} \neq 0$. Hence

$$(\tilde{c}_{eox} - \tilde{x}_2^0) > 0, \text{ and} \quad (7.61)$$

$$(\tilde{N} - \tilde{c}_{eox} + \tilde{x}_2^0) > 0. \quad (7.62)$$

Consider the third component of equation (7.42):

$$\begin{aligned} & k_{1b} \left(c_{eox} - \frac{\tilde{\beta}}{\beta} \tilde{x}_2 - x_2^0 \right) - k_{1f} \left(\frac{\tilde{\beta}}{\beta} \tilde{x}_2 + x_2^0 \right) (\lambda_3 + x_3^0) \left(N - c_{eox} + \frac{\tilde{\beta}}{\beta} \tilde{x}_2 + x_2^0 \right) - q\lambda_3 \\ &= \frac{\partial \lambda_3}{\partial \tilde{x}_2} \left(\tilde{k}_{1b}(\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) - \tilde{k}_{1f}(\tilde{x}_2 + \tilde{x}_2^0)(\tilde{x}_3 + \tilde{x}_3^0)(\tilde{N} - \tilde{c}_{eox} + \tilde{x}_2 + \tilde{x}_2^0) + \tilde{k}_2\tilde{x}_1(\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) \right) \\ &+ \frac{\partial \lambda_3}{\partial \tilde{x}_3} \left(\tilde{k}_{1b}(\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) - \tilde{k}_{1f}(\tilde{x}_2 + \tilde{x}_2^0)(\tilde{x}_3 + \tilde{x}_3^0)(\tilde{N} - \tilde{c}_{eox} + \tilde{x}_2 + \tilde{x}_2^0) - q\tilde{x}_3 \right) \end{aligned} \quad (7.63)$$

Differentiating (7.63) with respect to \tilde{x}_1 gives

$$0 = \frac{\partial \lambda_3}{\partial \tilde{x}_2} \tilde{k}_2 (\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0). \quad (7.64)$$

Thus either $\tilde{x}_2 = \tilde{c}_{eox} - \tilde{x}_2^0$ for all time or $\frac{\partial \lambda_3}{\partial \tilde{x}_2} = 0$. The former is impossible, since

$\tilde{x}_2(0) = 0$, and $(\tilde{c}_{eox} - \tilde{x}_2^0) > 0$ (from (7.61)). Hence

$$\frac{\partial \lambda_3}{\partial \tilde{x}_2} = 0. \quad (7.65)$$

Now consider the first component of equation (7.42a). Again expanding this using the information we have about λ , we obtain the following:

$$\begin{aligned} & \left\{ -k_2 \left(c_{eox} - \frac{\tilde{\beta}}{\beta} \tilde{x}_2 - x_2^0 \right) - q \right\} \lambda_1 \\ &= \frac{r_{an}}{\tilde{r}_{an}} \left(-\tilde{k}_2 \tilde{x}_1 (\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) - \tilde{q} \tilde{x}_1 \right) \\ &+ \frac{\partial \lambda_1}{\partial \tilde{x}_2} \left(\tilde{k}_{1b} (\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) - \tilde{k}_{1f} (\tilde{x}_2 + \tilde{x}_2^0) (\tilde{x}_3 + \tilde{x}_3^0) (\tilde{N} - \tilde{c}_{eox} + \tilde{x}_2 + \tilde{x}_2^0) + \tilde{k}_2 \tilde{x}_1 (\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) \right) \\ &+ \frac{\partial \lambda_1}{\partial \tilde{x}_3} \left(\tilde{k}_{1b} (\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) - \tilde{k}_{1f} (\tilde{x}_2 + \tilde{x}_2^0) (\tilde{x}_3 + \tilde{x}_3^0) (\tilde{N} - \tilde{c}_{eox} + \tilde{x}_2 + \tilde{x}_2^0) - \tilde{q} \tilde{x}_3 \right) \end{aligned} \quad (7.66)$$

Differentiating (7.66) with respect to \tilde{x}_1 gives

$$\frac{r_{an}}{\tilde{r}_{an}} \left\{ -k_2 \left(c_{eox} - \frac{\tilde{\beta}}{\beta} \tilde{x}_2 - x_2^0 \right) - q \right\} = \frac{r_{an}}{\tilde{r}_{an}} \left(-\tilde{k}_2 (\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) - \tilde{q} \right) + \frac{\partial \lambda_1}{\partial \tilde{x}_2} \left(\tilde{k}_2 (\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) \right) \quad (7.67)$$

since $\frac{\partial \lambda_1}{\partial \tilde{x}_1} = \frac{r_{an}}{\tilde{r}_{an}}$. Differentiating (7.67) with respect to \tilde{x}_3 shows that

$$\frac{\partial^2 \lambda_1}{\partial \tilde{x}_2 \partial \tilde{x}_3} = 0. \quad (7.68)$$

Now, differentiating (7.50) with respect to \tilde{x}_3 gives

$$\begin{aligned} & -\frac{\partial \lambda_3}{\partial \tilde{x}_3} k_{1f} \left(\frac{\tilde{\beta}}{\beta} \tilde{x}_2 + x_2^0 \right) \left(N - c_{eox} + \frac{\tilde{\beta}}{\beta} \tilde{x}_2 + x_2^0 \right) + \frac{\partial \lambda_1}{\partial \tilde{x}_3} k_2 \left(c_{eox} - \frac{\tilde{\beta}}{\beta} \tilde{x}_2 - x_2^0 \right) \\ & = -\frac{\tilde{\beta}}{\beta} \tilde{k}_{1f} (\tilde{x}_2 + \tilde{x}_2^0) (\tilde{N} - \tilde{c}_{eox} + \tilde{x}_2 + \tilde{x}_2^0) \end{aligned} \quad (7.69)$$

Differentiating (7.69) with respect to \tilde{x}_2 (and using (7.65) and (7.68)) gives

$$\frac{\partial \lambda_3}{\partial \tilde{x}_3} k_{1f} \left(2 \left(\frac{\tilde{\beta}}{\beta} \right) \tilde{x}_2 + N - c_{eox} + 2x_2^0 \right) + k_2 \frac{\partial \lambda_1}{\partial \tilde{x}_3} = \tilde{k}_{1f} (2\tilde{x}_2 + \tilde{N} - \tilde{c}_{eox} + 2\tilde{x}_2^0) \quad (7.70)$$

Now differentiating again with respect to \tilde{x}_2 (and again using (7.65) and (7.68)) gives

$$\frac{\partial \lambda_3}{\partial \tilde{x}_3} = \frac{\tilde{k}_{1f}}{k_{1f}} \frac{\beta}{\tilde{\beta}}, \quad (7.71)$$

which implies that (from equations (7.48) and (7.65), and because $\lambda_3(0) = 0$ (see equation (7.41))

$$\lambda_3 = \frac{\tilde{k}_{1f}}{k_{1f}} \frac{\beta}{\tilde{\beta}} \tilde{x}_3. \quad (7.72)$$

Now that λ_2 and λ_3 are known, equation (7.49) becomes

$$\begin{aligned}
 & k_{1b} \left(c_{eox} - \frac{\tilde{\beta}}{\beta} \tilde{x}_2 - x_2^0 \right) - k_{1f} \left(\frac{\tilde{\beta}}{\beta} \tilde{x}_2 + x_2^0 \right) \left(\frac{\tilde{k}_{1f}}{k_{1f}} \frac{\beta}{\tilde{\beta}} \tilde{x}_3 + x_3^0 \right) \left(N - c_{eox} + \frac{\tilde{\beta}}{\beta} \tilde{x}_2 + x_2^0 \right) \\
 & + k_2 \lambda_1 \left(c_{eox} - \frac{\tilde{\beta}}{\beta} \tilde{x}_2 - x_2^0 \right) \\
 & = \frac{\tilde{\beta}}{\beta} \left(\tilde{k}_{1b} (\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) - \tilde{k}_{1f} (\tilde{x}_2 + \tilde{x}_2^0) (\tilde{x}_3 + \tilde{x}_3^0) (\tilde{N} - \tilde{c}_{eox} + \tilde{x}_2 + \tilde{x}_2^0) + \tilde{k}_2 \tilde{x}_1 (\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) \right)
 \end{aligned} \tag{7.73}$$

Differentiating (7.73) with respect to \tilde{x}_2 gives

$$\begin{aligned}
 & -k_{1b} \frac{\tilde{\beta}}{\beta} - k_{1f} \left(\frac{\tilde{k}_{1f}}{k_{1f}} \frac{\beta}{\tilde{\beta}} \tilde{x}_3 + x_3^0 \right) \left(2 \left(\frac{\tilde{\beta}}{\beta} \right) \tilde{x}_2 + \frac{\tilde{\beta}}{\beta} (N - c_{eox} + 2x_2^0) \right) \\
 & + k_2 \left(\frac{\partial \lambda_1}{\partial \tilde{x}_2} \left(c_{eox} - \frac{\tilde{\beta}}{\beta} \tilde{x}_2 - x_2^0 \right) - \frac{\tilde{\beta}}{\beta} \lambda_1 \right) \\
 & = \frac{\tilde{\beta}}{\beta} \left(-\tilde{k}_{1b} - \tilde{k}_{1f} (\tilde{x}_3 + \tilde{x}_3^0) (2\tilde{x}_2 + \tilde{N} - \tilde{c}_{eox} + 2\tilde{x}_2^0) - \tilde{k}_2 \tilde{x}_1 \right)
 \end{aligned} \tag{7.74}$$

Differentiating again, this time with respect to \tilde{x}_3 , and then rearranging, gives

$$\left(N - c_{eox} + 2x_2^0 \right) + \frac{k_2}{\tilde{k}_{1f}} \frac{\tilde{\beta}}{\beta} \frac{\partial \lambda_1}{\partial \tilde{x}_3} = \frac{\tilde{\beta}}{\beta} (\tilde{N} - \tilde{c}_{eox} + 2\tilde{x}_2^0). \tag{7.75}$$

Now, on substitution for λ_2 and λ_3 , equation (7.63) becomes

$$\begin{aligned}
 & k_{1b} \left(c_{eox} - \frac{\tilde{\beta}}{\beta} \tilde{x}_2 - x_2^0 \right) - k_{1f} \left(\frac{\tilde{\beta}}{\beta} \tilde{x}_2 + x_2^0 \right) \left(\frac{\tilde{k}_{1f}}{k_{1f}} \frac{\beta}{\tilde{\beta}} \tilde{x}_3 + x_3^0 \right) \left(N - c_{eox} + \frac{\tilde{\beta}}{\beta} \tilde{x}_2 + x_2^0 \right) - q \frac{\tilde{k}_{1f}}{k_{1f}} \frac{\beta}{\tilde{\beta}} \tilde{x}_3 \\
 & = \frac{\tilde{k}_{1f}}{k_{1f}} \frac{\beta}{\tilde{\beta}} \left(\tilde{k}_{1b} (\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) - \tilde{k}_{1f} (\tilde{x}_2 + \tilde{x}_2^0) (\tilde{x}_3 + \tilde{x}_3^0) (\tilde{N} - \tilde{c}_{eox} + \tilde{x}_2 + \tilde{x}_2^0) - \tilde{q} \tilde{x}_3 \right)
 \end{aligned} \tag{7.76}$$

Differentiating (7.76) with respect to \tilde{x}_3 and simplifying gives

$$\begin{aligned}
 & -\frac{\beta}{\tilde{\beta}} \left(\frac{\tilde{\beta}}{\beta} \tilde{x}_2 + x_2^0 \right) \left(N - c_{\text{eox}} + \frac{\tilde{\beta}}{\beta} \tilde{x}_2 + x_2^0 \right) - q \frac{1}{k_{1f}} \frac{\beta}{\tilde{\beta}} \\
 & = \frac{1}{k_{1f}} \frac{\beta}{\tilde{\beta}} \left(-\tilde{k}_{1f} (\tilde{x}_2 + \tilde{x}_2^0) (\tilde{N} - \tilde{c}_{\text{eox}} + \tilde{x}_2 + \tilde{x}_2^0) - \tilde{q} \right).
 \end{aligned} \tag{7.77}$$

Equating coefficients of \tilde{x}_2^2 yields the relationship

$$\left(\frac{\tilde{\beta}}{\beta} \right)^2 = \frac{\tilde{k}_{1f}}{k_{1f}}. \tag{7.78}$$

Now equating coefficients of \tilde{x}_2 in equation (7.77) and substituting equation (7.78) gives

$$N - c_{\text{eox}} + 2x_2^0 = \frac{\tilde{\beta}}{\beta} (\tilde{N} - \tilde{c}_{\text{eox}} + 2\tilde{x}_2^0), \tag{7.79}$$

which when substituted into equation (7.75), tells us that

$$\frac{\partial \lambda_1}{\partial \tilde{x}_3} = 0. \tag{7.80}$$

Now equation (7.66) becomes

$$\begin{aligned}
 & \left\{ -k_2 \left(c_{\text{eox}} - \frac{\tilde{\beta}}{\beta} \tilde{x}_2 - x_2^0 \right) - q \right\} \lambda_1 \\
 & = \frac{r_{\text{an}}}{\tilde{r}_{\text{an}}} \left(-\tilde{k}_2 \tilde{x}_1 (\tilde{c}_{\text{eox}} - \tilde{x}_2 - \tilde{x}_2^0) - \tilde{q} \tilde{x}_1 \right) \\
 & + \frac{\partial \lambda_1}{\partial \tilde{x}_2} \left(\tilde{k}_{1b} (\tilde{c}_{\text{eox}} - \tilde{x}_2 - \tilde{x}_2^0) - \tilde{k}_{1f} (\tilde{x}_2 + \tilde{x}_2^0) (\tilde{x}_3 + \tilde{x}_3^0) (\tilde{N} - \tilde{c}_{\text{eox}} + \tilde{x}_2 + \tilde{x}_2^0) + \tilde{k}_2 \tilde{x}_1 (\tilde{c}_{\text{eox}} - \tilde{x}_2 - \tilde{x}_2^0) \right)
 \end{aligned} \tag{7.81}$$

Differentiating equation (7.81) with respect to \tilde{x}_3 and using equations (7.68) and (7.80) gives

$$0 = \tilde{k}_{1f} \frac{\partial \lambda_1}{\partial \tilde{x}_2} (\tilde{x}_2 + \tilde{x}_2^0) (\tilde{N} - \tilde{c}_{eox} + \tilde{x}_2 + \tilde{x}_2^0) \quad (7.82)$$

Since (except for possibly two values of \tilde{x}_2) $(\tilde{x}_2 + \tilde{x}_2^0) (\tilde{N} - \tilde{c}_{eox} + \tilde{x}_2 + \tilde{x}_2^0) \neq 0$, it follows that

$$\frac{\partial \lambda_1}{\partial \tilde{x}_2} = 0. \quad (7.83)$$

Thus the Jacobian is complete, and is given by (using equation (7.78) to simplify the (3,3) element)

$$T = \begin{pmatrix} \frac{\tilde{r}_{an}}{r_{an}} & 0 & 0 \\ 0 & \frac{\tilde{\beta}}{\beta} & 0 \\ 0 & 0 & \frac{\tilde{\beta}}{\beta} \end{pmatrix}. \quad (7.84)$$

This shows that the final condition of Corollary 1 is satisfied (i.e. λ is linear), and thus Corollary 1 can be applied to our system.

Now, the first component of equation (7.42a) becomes

$$\left\{ -k_2 \left(c_{eox} - \frac{\tilde{\beta}}{\beta} \tilde{x}_2 - x_2^0 \right) - q \right\} \frac{r_{an}}{\tilde{r}_{an}} \tilde{x}_1 = \frac{r_{an}}{\tilde{r}_{an}} \left(-\tilde{k}_2 \tilde{x}_1 (\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) - \tilde{q} \tilde{x}_1 \right), \quad (7.85)$$

which reduces to

$$k_2 \left(c_{eox} - \frac{\tilde{\beta}}{\beta} \tilde{x}_2 - x_2^0 \right) + q = \tilde{k}_2 (\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) + \tilde{q}. \quad (7.86)$$

Equating coefficients of \tilde{x}_2 gives

$$k_2 \frac{\tilde{\beta}}{\beta} = \tilde{k}_2, \quad (7.87)$$

and so

$$\frac{\tilde{\beta}}{\beta} = \frac{\tilde{k}_2}{k_2}. \quad (7.88)$$

Equating constant coefficients gives

$$k_2(c_{eox} - x_2^0) + q = \tilde{k}_2(\tilde{c}_{eox} - \tilde{x}_2^0) + \tilde{q}, \quad (7.89)$$

which, together with equation (7.54), implies that

$$q = \tilde{q}, \quad (7.90)$$

and thus q is globally identifiable.

The second component of equation (7.42a) becomes

$$\begin{aligned} & k_{1b} \left(c_{eox} - \frac{\tilde{\beta}}{\beta} \tilde{x}_2 - x_2^0 \right) - k_{1f} \left(\frac{\tilde{\beta}}{\beta} \tilde{x}_2 + x_2^0 \right) \left(\frac{\tilde{\beta}}{\beta} \tilde{x}_3 + x_3^0 \right) \left(N - c_{eox} + \frac{\tilde{\beta}}{\beta} \tilde{x}_2 + x_2^0 \right) \\ & + k_2 \left(c_{eox} - \frac{\tilde{\beta}}{\beta} \tilde{x}_2 - x_2^0 \right) \frac{r_{an}}{\tilde{r}_{an}} \tilde{x}_1 \\ & = \frac{\tilde{\beta}}{\beta} \left(\tilde{k}_{1b} (\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) - \tilde{k}_{1f} (\tilde{x}_2 + \tilde{x}_2^0) (\tilde{x}_3 + \tilde{x}_3^0) (\tilde{N} - \tilde{c}_{eox} + \tilde{x}_2 + \tilde{x}_2^0) + \tilde{k}_2 \tilde{x}_1 (\tilde{c}_{eox} - \tilde{x}_2 - \tilde{x}_2^0) \right) \end{aligned} \quad (7.91)$$

Equating coefficients of \tilde{x}_3 in equation (7.91) and using equation (7.78) gives

$$x_2^0(N - c_{\text{eox}} + x_2^0) = \left(\frac{\tilde{\beta}}{\beta}\right)^2 \tilde{x}_2^0(\tilde{N} - \tilde{c}_{\text{eox}} + \tilde{x}_2^0). \quad (7.92)$$

Similarly equating coefficients of \tilde{x}_2^2 in equation (7.91) and simplifying gives

$$x_3^0 = \frac{\tilde{\beta}}{\beta} \tilde{x}_3^0, \quad (7.93)$$

however, since $x_3^0 = \frac{r_{\text{ox}}}{q}$ and $q = \tilde{q}$ (by equations (7.22) and (7.90)), it follows that

$$r_{\text{ox}} = \frac{\tilde{\beta}}{\beta} \tilde{r}_{\text{ox}}. \quad (7.94)$$

Equating coefficients of \tilde{x}_2 and simplifying using equations (7.78), (7.79) and (7.93) gives

$$k_{1b} = \tilde{k}_{1b}, \quad (7.95)$$

and thus k_{1b} is identifiable.

Now no more individual parameters can be shown to be identifiable without using the definition of x_2^0 in equation (7.22) (which is algebraically very difficult). This is a potential problem since if the model is unidentifiable then parameter estimates may be unreliable. However, there are a couple of possible courses of action to remedy the situation:

- Other information could be employed to fix the values of one or more of the parameters – for example, the parameter β corresponds to the sensitivity of the output to changes in the value of x_2 , the number of conducting electrons. Thus an estimate should be available from the physics literature.

- It may be possible to reparameterise the model to produce a slightly different model, which retains the important features of the present model, yet is structurally identifiable.

Now, in this particular case, both options are capable of yielding success. The latter method is chosen. In fact, the reparameterisation used results in a model equivalent to the current one, with the value of β set to unity. So in a sense it is equivalent to the former method, with units for some of the parameters and variables altered to reflect the fixing of β . The reparameterisation used is given by setting

$$\begin{aligned}
\hat{x}_1 &= \beta \bar{x}_1 \\
\hat{x}_2 &= \beta \bar{x}_2 \\
\hat{x}_3 &= \beta \bar{x}_3 \\
\hat{x}_4 &= \beta \bar{x}_4 \\
\hat{x}_6 &= \beta \bar{x}_6 \\
\hat{r}_{an} &= \beta \bar{r}_{an} \\
\hat{r}_{ox} &= \beta \bar{r}_{ox} \\
\hat{k}_2 &= \frac{k_2}{\beta} \\
\hat{k}_{1f} &= \frac{k_{1f}}{\beta^2}.
\end{aligned} \tag{7.96}$$

With this reparameterisation, the system of differential equations (7.7) becomes:

$$\begin{aligned}
\dot{\hat{x}}_1 &= \hat{r}_{an} - \hat{k}_2 \hat{x}_1 \hat{x}_4 - q \hat{x}_1 \\
\dot{\hat{x}}_2 &= -\hat{k}_{1f} \hat{x}_2 \hat{x}_3 \hat{x}_6 + k_{1b} \hat{x}_4 + \hat{k}_2 \hat{x}_1 \hat{x}_4 \\
\dot{\hat{x}}_3 &= \hat{r}_{ox} - \hat{k}_{1f} \hat{x}_2 \hat{x}_3 \hat{x}_6 + k_{1b} \hat{x}_4 - q \hat{x}_3 \\
\dot{\hat{x}}_4 &= \hat{k}_{1f} \hat{x}_2 \hat{x}_3 \hat{x}_6 - k_{1b} \hat{x}_4 - \hat{k}_2 \hat{x}_1 \hat{x}_4 \\
\dot{\bar{x}}_5 &= \frac{\hat{k}_2}{\beta} \hat{x}_1 \hat{x}_4 + r_{cl} \bar{x}_7 - q \bar{x}_5 \\
\dot{\hat{x}}_6 &= -\hat{k}_{1f} \hat{x}_2 \hat{x}_3 \hat{x}_6 + k_{1b} \hat{x}_4 + \hat{k}_2 \hat{x}_1 \hat{x}_4 \\
\dot{\bar{x}}_7 &= -r_{cc} \bar{x}_7 + q \bar{x}_5,
\end{aligned} \tag{7.97}$$

with the new observation function:

$$h(\hat{x}(t, p), p) = \frac{1}{\alpha + \hat{x}_2}. \quad (7.98)$$

The system of equations (7.97) and (7.98) is exactly the same (apart from labelling) as that described by (7.7) and (7.11), except with the value of β set to unity. The physical meanings of the parameters c_{eox} and N are also altered by the reparameterisation, being replaced by βc_{eox} and βN respectively. Thus all of the calculations of this chapter remain valid for this new reparameterised system. Thus the ‘hats’ are dropped for convenience, and the identifiability analysis can proceed for the new system.

Note that since β is characteristic of the sensor material and thus not related to the analyte, this reparameterisation does not affect which of the parameters are appropriate for use in an odour discrimination method.

Now, since in the new system β is replaced by one, equations (7.94), (7.78), (7.88) and (7.53) give us that

$$r_{\text{ox}} = \tilde{r}_{\text{ox}}, \quad (7.99)$$

$$k_{1f} = \tilde{k}_{1f}, \quad (7.100)$$

$$k_2 = \tilde{k}_2, \quad (7.101)$$

and

$$r_{\text{an}} = \tilde{r}_{\text{an}}. \quad (7.102)$$

Equations (7.54), (7.79) and (7.92) have now reduced to:

$$c_{\text{eox}} - x_2^0 = \tilde{c}_{\text{eox}} - \tilde{x}_2^0 \quad (7.103)$$

$$N - c_{\text{eox}} + 2x_2^0 = \tilde{N} - \tilde{c}_{\text{eox}} + 2\tilde{x}_2^0 \quad (7.104)$$

$$x_2^0(N - c_{\text{eox}} + x_2^0) = \tilde{x}_2^0(\tilde{N} - \tilde{c}_{\text{eox}} + \tilde{x}_2^0). \quad (7.105)$$

Adding together equations (7.103) and (7.104) gives

$$N + x_2^0 = \tilde{N} + \tilde{x}_2^0 \quad (7.106)$$

and so

$$N = \tilde{N} + \tilde{x}_2^0 - x_2^0. \quad (7.107)$$

Rearranging equation (7.103) gives

$$c_{\text{eox}} = \tilde{c}_{\text{eox}} - \tilde{x}_2^0 + x_2^0. \quad (7.108)$$

Now, using equations (7.107) and (7.108), we have that

$$N - c_{\text{eox}} + x_2^0 = \tilde{N} + \tilde{x}_2^0 - x_2^0 - (\tilde{c}_{\text{eox}} - \tilde{x}_2^0 + x_2^0) + x_2^0, \quad (7.109)$$

which rearranges to give

$$N - c_{\text{eox}} + x_2^0 = \tilde{N} - \tilde{c}_{\text{eox}} + \tilde{x}_2^0 + (\tilde{x}_2^0 - x_2^0). \quad (7.110)$$

Now, substituting equation (7.110) into (7.105) and rearranging yields

$$x_2^0(\tilde{x}_2^0 - x_2^0) = (\tilde{N} - \tilde{c}_{\text{eox}} + \tilde{x}_2^0)(\tilde{x}_2^0 - x_2^0). \quad (7.111)$$

Thus either $\tilde{x}_2^0 = x_2^0$, in which case all of the remaining parameters are identified, or

$$x_2^0 = \tilde{N} - \tilde{c}_{\text{eox}} + \tilde{x}_2^0, \quad (7.112)$$

in which case we have a second (distinct) solution to the system of equations (7.103), (7.104) and (7.105):

$$\tilde{x}_2^0 = N - c_{eox} + x_2^0, \quad (7.113)$$

$$\tilde{c}_{eox} = N, \text{ and} \quad (7.114)$$

$$\tilde{N} = c_{eox}. \quad (7.115)$$

Thus the (reparameterised) model is shown to be at least locally identifiable, with two possible parameter sets.

7.8 Conclusions

A physical model for the interaction between gaseous analytes and an MOS gas sensor has been taken from the literature and converted into a mathematical model, described via a system of nonlinear ordinary differential equations.

Assumptions were made regarding the gas delivery system and electronic interfaces, as discussed in Chapter 6, in order to produce a model for the measured outputs of the electronic nose system. The model was reduced via decoupling and subsequently analysed for steady states. Steady states were also calculated for the system in the absence of the analyte, to be used as initial conditions for the full system, representing typical experimental operating procedures.

The system was shown to be controllable and observable, using the CRC and ORC respectively. The identifiability of the system was investigated, using the similarity transformation approach. The system was found to be at least locally identifiable, with at most two solutions to the set of equations generated by the identifiability analysis.

It remains to be seen whether use of the expression for x_2^0 (7.22) can discount the second solution set and thus guarantee global identifiability for all of the parameters.

Even without such a result, given the physical meanings of N and c_{ox} it is possible that bounds on estimates of the two parameters may be produced, which might permit positive identification of each (for example if one were expected to be many orders of magnitude greater than the other, then selecting which was which would be possible). In any case neither parameter relates directly to the analyte being sensed, so (to a point) it is not important for odour discrimination that the correct choice of solutions be made.

Chapter 8

Parameter estimation and validation of nose system model

In this chapter, the usefulness of the model of the previous chapter is investigated using data collected from the experimental test rig detailed in Chapter 5. Having established the identifiability of the model (at least locally), it is reasonable to proceed with parameter estimations and simulations to compare with experimental data in order to establish how well the model describes an actual electronic nose system. The practical usefulness of the model is investigated, in terms of using the model for discriminating between unknown analytes. The versatility and extensibility of the model is also examined as the effects of operating a discrimination system based on the model under less controlled conditions are investigated (i.e. with varying ambient temperature, strength of odour, flow rates).

8.1 Experimental data

The experimental data used for the model validation in this chapter were collected using the test rig described in Chapter 5. Three sample pots were used, the first containing nothing (except air), the second containing approximately 0.01ml of acetone diluted in approximately 6.25ml of water and the third containing a similar

dilution of isopropyl alcohol. The system was set up so that the sensors were exposed to the following repeated cycle:

- 20 min - air only
- 1 min - acetone diluted in water
- 20 min - air only
- 1 min - isopropyl alcohol diluted in water.

The outputs from the sensors were sampled once per second for a period of approximately 13 days, producing 1,134,000 data vectors corresponding to 450 exposure cycles. Over the period of the experiment, the responses of the gas sensors varied considerably. This is illustrated with a plot of the output of sensor 1 in Figure 8.1. As mentioned previously, the responses of metal oxide sensors depend heavily on factors such as their operating temperature, the temperature of the samples, the humidity of the odour carrying air and the flow rate through the gas delivery system. The long term responses of the sensors are also potentially affected by sensor poisoning effects and possible variations in the headspace concentrations of the target gases in the sample containers (Gardner and Bartlett 1999). The long duration of the experiment and fairly low concentrations of the diluted samples mean that the reductions in the concentrations of the target gases presented to the sensors over the course of the experiment may be a highly significant factor in the 'drift' of the response. A typical response of one of the gas sensors to an exposure cycle is plotted in Figure 8.2.




Figure 8.1 A plot of the output of sensor 1 over the entire duration of the experiment, showing the long term variation in the responses of the sensor.




Figure 8.2 A typical response of sensor 2 to a single exposure cycle. The first response is to an input of acetone, the second to isopropyl alcohol. The dotted line shows when the valves were operated on the gas delivery system.

The flow rate and exposure duration were selected to provide a compromise between the desire to record the full response of each sensor (i.e. for the sensor output to level off before the gas input is switched back to just air), the desire to maintain a static headspace in the sample container (see Section 6.2.1 for explanation of this), and the desire to reduce the risk of substantial poisoning of the sensors by over-exposure.

It is clear from the flatness of the saturation level of the sensor response to an input of acetone in Figure 8.3 that there was no significant headspace degradation over the period of the exposure. This was an unusually fast sensor / analyte combination – typically the sensors took longer to reach equilibrium levels, so the exposure period chosen was deemed appropriate.




Figure 8.3 A response of sensor 2 to an input of acetone. The flatness of the saturation level shows the concentration of the analyte in the headspace of the sample container was not significantly degraded over the sixty seconds for which it was sampled.

8.2 Initial model fitting and parameter estimations

There are a number of methods available for numerical simulation and parameter estimation for systems of nonlinear ordinary differential equations. Preliminary investigations showed that the system considered here is very stiff (order of stiffness approximately 10^{12}) – this rules out a number of software packages which do not provide algorithms that are capable of coping with this level of stiffness. *Facsimile* (<http://www.mcpa-software.com>) and *Berkeley Madonna* (<http://www.berkeleymadonna.com>) are two software packages which are able to simulate very stiff systems. After initial investigations with both packages, it was decided to use *Berkeley Madonna* for this model, since it seemed to cope much better than *Facsimile* for this particular system, producing simulations for parameter sets where *Facsimile* could not.

Parameter estimation with *Berkeley Madonna* was found to be relatively straightforward and flexible. The curve fitting procedure requires the minimisation of the root mean square error between real data points and corresponding simulated data points from numerical solutions to the system of differential equations. As with most procedures requiring the minimisation of a complicated cost function over a large number of parameters, the success of the minimisation routine depends heavily upon the initial guesses for the unknown parameter set. The *Berkeley Madonna* package uses the ‘downhill simplex’ minimisation algorithm (for details of this algorithm see Press *et al.* 1993) to produce curve fits.

A further restriction on the application of the software was imposed by the difficulty / impossibility of automating the process of parameter estimation for a great many exposure cycles. This meant that the parameter estimation for each exposure cycle for each sensor had to be performed manually, thus restricting greatly the amount of data that can realistically be produced.

It was decided that all parameter estimates would be made based only on the ‘on’ portion of the sensor responses (i.e. the part of the response curve in Figure 8.3 where the sensor output is dropping, and not the part after the valve has switched over to flush the sensor chamber with air and the sensor recovers to (or at least near to) its original state). This was decided for two main reasons:

- It is clear from the electrochemical reaction scheme described in equations (7.1) and (7.2) that the ‘off’ portion of the response (where the mechanism of (7.2) is no longer occurring) does not directly involve the analyte gas, so this portion of the response should not be useful for odour classification purposes. The only useful information carried in this portion relates to the final value that the sensor outputs take (which of course can also be obtained from the ‘on’ portion).
- Whilst it was possible to obtain very good simulations for the ‘on’ period of the sensor response (see the next section), it was found to be very difficult to obtain accurate simulations for both the ‘on’ and ‘off’ portions simultaneously. Even after extensive investigations using different initial guesses and parameter bounds, acceptable fits were never produced. This is possibly due to the simplicity of the

model for the gas dynamics in the chamber. The concentration of the analyte in the main chamber will not in fact drop instantaneously to zero, and it would appear that more complex modelling of the concentration profile in the vicinity of the sensor would be required to produce simulations that match the real data.

If the model had described the ‘off’ portion of the sensor response well, then this part of the response could have been useful for estimating some of the analyte-independent parameters, however, since this is not the case, this portion of the response was discarded.

8.2.1 Parameter estimation for a single sensor response

The first step was to produce a simulation using the model of Chapter 7 that matched a single sensor response. A portion of a response of sensor 1 to an exposure of diluted acetone (only the ‘on’ phase of the response – so not including the recovery of the sensor after the target gas input is removed) was isolated and *Berkeley Madonna* was used to produce parameter estimates and fit the model to the data. A step change in the value of the parameter r_{an} was used to model the introduction of the target gas to the sensor chamber. This corresponds to a step change in the concentration C_{an} of the target gas in the main chamber (see Section 7.3 and Section 7.7.3). The results are shown in Figure 8.4.




Figure 8.4 The ‘on’ portion of a response of sensor 1 to diluted acetone; x’s are the real data points, the solid line relates to simulated data using the model of Chapter 7. The r.m.s. error of the curve fit was approximately 4.58×10^{-3} V.

The curve fit shown in Figure 8.4 is very encouraging, however, the structural identifiability of the model established in the previous chapter was not empirically observed in the practicalities of the parameter estimations for this experiment. It was found that several parameters could be fixed at almost any desired value, and then acceptable model fits could be obtained by varying only the remaining parameters. It should also be noted that the curve fit in Figure 8.4 was only obtained after a great deal of experimentation with different initial guesses for the parameter values. In fact, it was necessary to manually experiment with different parameter values to get a response curve of the right qualitative form, prior to using the minimisation algorithms to home in on the minima produced. Earlier curve fits did not reproduce the shape of the response curve at the start of the odour introduction and instead produced sharp drops. This apparent failure of the software, the model, and perhaps the experiment, to produce an empirically identifiable system leads to a lack of confidence in the

accuracy of any parameter estimates produced on a single sensor response. As discussed in Chapter 6, the aim of the modelling is to produce a classification system using estimates of the parameters relating to the target species to distinguish between different analytes. However, not all of the parameters in the model depend upon the analyte in question. Thus some parameters should be constant or at least bounded over multiple responses.

8.2.2 Model fitting to multiple analyte responses

The parameters in the mathematical model all have a physical meaning. Although the precise meaning of them is slightly modified by the reparameterisation in Section 7.7.3, it is still clear that some of the parameters relate only to the e-nose system, and not at all to the analyte being detected. Thus these parameters should be constant for a given sensor operating under fixed operating conditions (i.e. same temperature, humidity, gas flow rate settings, background odours etc.), irrespective of which analyte the sensor is exposed to. This is potentially a great help in reducing the choice of parameters which can give rise to equally good model fits, and thus ensuring the consistency and reliability of parameter estimates.

A model fit was performed using *Berkeley Madonna*, to two sections of data simultaneously, both from sensor 1, from successive exposures to the analyte gases; one exposure to acetone and one to isopropyl alcohol. The following parameters were fixed to be equal over the two model fits since they relate to the nose system hardware and not the choice of analyte: k_{1f} , k_{1b} , c_{eox} , r_{ox} , N , q and α . The parameters k_2 , r_{an} and a delay parameter (included to allow for possible slight inconsistencies in the splitting up of the data – not because significant differences in the physical delay introduced by the gas dynamics of the system for different analytes were expected or observed) were allowed to take different values for each analyte. The results are plotted in Figure 8.5.




Figure 8.5 Simulations produced by simultaneous parameter estimations to the response of sensor 1 to (a) an input of acetone and (b) an input of isopropyl alcohol. The parameters corresponding to the nose system and not the analyte were kept equal for the estimation procedure. The combined r.m.s. error of the two curve fits was approximately 0.175 V.

It is clear from Figure 8.5 that the model fits to the data are significantly poorer than those obtained when the supposedly analyte-independent parameters are allowed to vary for the different analyte inputs. It was found that curve fits to multiple responses to the same analyte were also noticeably poorer than individual fits when these parameters were held constant for the different responses. This indicates that the cause of the poor fits was a combination of experimental variations and possibly the simplifications inherent in the model used. To rectify this, the parameters which relate to the nose system and not to the analyte sensed were constrained to lie within certain intervals, rather than held constant, for different response fits. This produced significantly better results. An example of a pair of curve fits produced using the same narrow bounds for the sensor / system parameters, and much more freedom for the

analyte parameters, is plotted in Figure 8.6. The fit to the isopropyl alcohol response in particular, is extremely good – having an r.m.s error of only approximately 5.45×10^{-3} V.




Figure 8.6 Simulations produced by parameter estimations to the response of sensor 1 to (a) an input of acetone and (b) an input of isopropyl alcohol. The parameters corresponding to the e-nose system and not the analyte were constrained to lie within the same narrow bounds for each curve fit. The combined r.m.s. error of the two curve fits was approximately 0.0174 V.

If the model had not been reparameterised then it may have been possible to use the fact that some of the parameters (namely r_{an} and r_{ox}) should remain constant for different sensors, to more accurately pin down their values. However, the reparameterisation of the model had the effect of making r_{an} and r_{ox} also depend upon the characteristics of the sensor and the interface electronics. Thus it is not reasonable to expect these parameters to remain constant across different sensors.

8.3 Parameter estimates over the range of the entire experiment

The method of constraining sensor / system dynamics to lie within certain bounds and allowing the analyte-dependent parameters to vary to reflect the differences in sensor responses was employed for a number of responses taken from the whole range of the experiment. Each parameter estimation run had to be performed manually. Due to the very long duration of the experiment and hence large number of response cycles (450), it was not possible to estimate the parameters for all of the responses. Instead, every tenth response cycle was extracted from the whole data set, a curve fit produced, and the parameters recorded. It was found by trial and error that the parameters k_{1f} and k_{1b} could be fixed without loss of fitting accuracy. The parameter α could be determined directly from the steady state sensor output prior to the odour input and the other parameters. Thus the remaining four system parameters were constrained within certain limits (which were found by trial and error), and the two analyte-specific parameters and a delay parameter were allowed to vary widely. The parameter ranges chosen are detailed in Table 8.1.

Parameter:	Minimum allowed value:	Maximum allowed value:
k_2	3×10^{-13}	1×10^{-11}
r_{an}	1×10^5	8×10^5
k_{1f}	6.35607×10^{-12}	6.35607×10^{-12}
k_{1b}	1.16128×10^{-6}	1.16128×10^{-6}
c_{eox}	6×10^3	8×10^3
r_{ox}	6×10^5	7×10^5
N	1×10^4	1.4×10^4
Q	0.15	0.2

Table 8.1 The ranges selected for the various parameters for curve fits to sensor 1 responses throughout the entire duration of the experiment.

The values of the analyte-dependent parameters fitted to the response curves for the 45 extracted response cycles (thus 45 responses to each of the two tested analytes) are

plotted in Figure 8.7, along with the r.m.s errors produced (a measure of goodness of fit).




Figure 8.7 The variation in the fitted values for k_2 (top), r_{an} (middle), and the associated r.m.s. errors (bottom) for curve fits to responses of sensor 1 over the duration of the experiment. The solid lines represent data corresponding to acetone responses, the dashed lines represent data corresponding to isopropyl alcohol responses.

8.3.1 Investigation of correlations between fitted parameters as a partial check for reliable parameter estimation

As shown in Figure 8.7, although the fitted values for k_2 are fairly stable, they do vary slightly over the course of the experiment, more notably so for the responses to acetone than for those to isopropyl alcohol. The fitted values for r_{an} also generally decline. These variations may be attributed to a number of possible factors. It is important for confidence in the fitted values that the fitted values for k_2 and r_{an} are not too closely

correlated with each other. If there were a clear relationship between the fitted values for the two parameters, then it might be the case that estimates for them are not reliable. A significant correlation could indicate that only some function of the two parameters (e.g. their product) might be empirically identifiable. In order to check this, Figure 8.8 plots the fitted values of k_2 against those of r_{an} for each analyte.




Figure 8.8 Plot of the fitted values of k_2 against those for r_{an} for (a) acetone response, and (b) isopropyl alcohol responses, all from sensor 1.

It is evident from Figure 8.8 that the fitted value of k_2 and r_{an} are not strongly correlated for each analyte. The correlation coefficient between k_2 and r_{an} for acetone is approximately 0.438, and that for isopropyl alcohol is approximately -0.0111 . Although the value for acetone is a little higher than might be expected for totally uncorrelated variables, this may not be significant. It may also be the case that both k_2 and r_{an} genuinely did vary over the course of the experiment due to other factors, for example, the temperature of the sample and the temperature of the sensor chamber may have affected each of the parameters, and there may have been a decrease in the

concentration of the analyte in the sample headspace caused by evaporation over the course of the experiment, affecting the 'true' value of r_{an} .

8.3.2 Investigation of the effects of variations in the temperature and humidity on the fitted parameters

The two temperature sensors (one positioned in the sensor chamber itself, and another mounted externally to monitor ambient temperature) were highly correlated with each other (with a correlation coefficient of approximately 0.9865 for the 90 responses selected). Therefore only the ambient temperature sensor data was used. The outputs of the humidity sensor (which measures relative humidity) in the chamber were strongly correlated (with a correlation coefficient of approximately 0.8866) with the ambient temperature, and more weakly correlated with the fitted values of k_2 and r_{an} , suggesting that perhaps the humidity was varying as a result of the variations in temperature. If there was a significant portion of data where the temperature remained fairly constant, yet the humidity varied, then the effects of variations in humidity on the values of k_2 and r_{an} might be determined. Since this was not the case, no firm connection between humidity and the parameter values can be established.

The matrix of correlation coefficients for the fitted values of k_2 and r_{an} and the temperature sensor output, is given in Table 8.2.

r	k_2	r_{an}	Temperature
k_2	1	0.4380	0.7408
r_{an}	0.4380	1	0.7106
Temperature	0.7408	0.7106	1

Table 8.2 Matrix of correlation coefficients (given to four significant figures) for the fitted parameters to acetone responses of sensor 1, and the recorded temperature and humidity values for the experiment.

The values of the fitted parameters k_2 and r_{an} are plotted against the ambient temperature in Figure 8.1.




Figure 8.9 Plot of ambient temperature versus (a) k_2 , and (b) r_{an} for responses of sensor 1 to acetone.

The points to the right of Figure 8.9(b) showing great variation of estimates for r_{an} around 23.5°C correspond to the fits near to the start of the experiment, where the temperature was stable, but the fits for r_{an} varied considerably. This could be due to evaporation of the acetone from the sample (see also Figure 8.7 to see the fall in the estimates for r_{an} near to the start of the experiment). The variation in the fitted values of r_{an} might be attributable to changes in the sensitivity of the sensor material with temperature (this is unlikely, since the sensor itself contains a microheater which heats the sensing material to a few hundred degrees Celsius) or to changes in the gas dynamics in the sensor chamber. However it seems most likely that the dominant mechanism by which this parameter is affected by temperature is through the evaporation rate and saturation concentration of the headspace of the sample itself. This parameter would not generally be considered useful for identification of a sample in non-laboratory conditions, since the strength of the odour (to which the parameter corresponds) could be expected to vary widely from sample to sample. For this reason there is little to be gained from attempting to remove or compensate for the effects of temperature on this parameter.

The data plotted in Figure 8.9(a) suggest a fairly good linear relationship between the fitted values of k_2 and the ambient temperature. A similar (albeit slightly weaker) relationship is found for the responses of sensor 1 to isopropyl alcohol. A linear function of the form

$$k_2(T) = k_T T + k_{2,0} \quad (8.1)$$

was fitted to the k_2 data sets, where T is the temperature (in degrees Celsius), k_T is a temperature coefficient and $k_{2,0}$ is a temperature compensated rate constant (effectively an extrapolated estimate of the value that k_2 would take at zero degrees Celsius). Note that since the linear fit is made over a small temperature range, a negative value for $k_{2,0}$ is not cause for concern.

Linear fits to the acetone and isopropyl alcohol responses separately yield markedly different values for k_T (approximately 3.8×10^{-13} and 8.7×10^{-14} respectively), suggesting that the dependence of the rate constant is analyte-specific. However, using

the same temperature coefficient for all fits to compensate for temperature variations still yields positive results (a correlation coefficient of approximately 0.6029 is calculated between the temperature and the combined mean-centred acetone and isopropyl alcohol k_2 estimates). A linear fit to the fitted k_2 values for both analytes together yields a coefficient of temperature of approximately 2.3×10^{-13} . Temperature compensated rate constant estimates were produced by simply taking each estimate for k_2 and subtracting from it $k_T T$, where k_T is an estimate for the temperature coefficient obtained from linear regression using the whole data set, and T is the recorded temperature for that sample exposure. A plot of temperature compensated fitted values of k_2 for all of the responses of sensor 1 is given in Figure 8.10, showing the reduced dependence of the parameter upon the temperature.

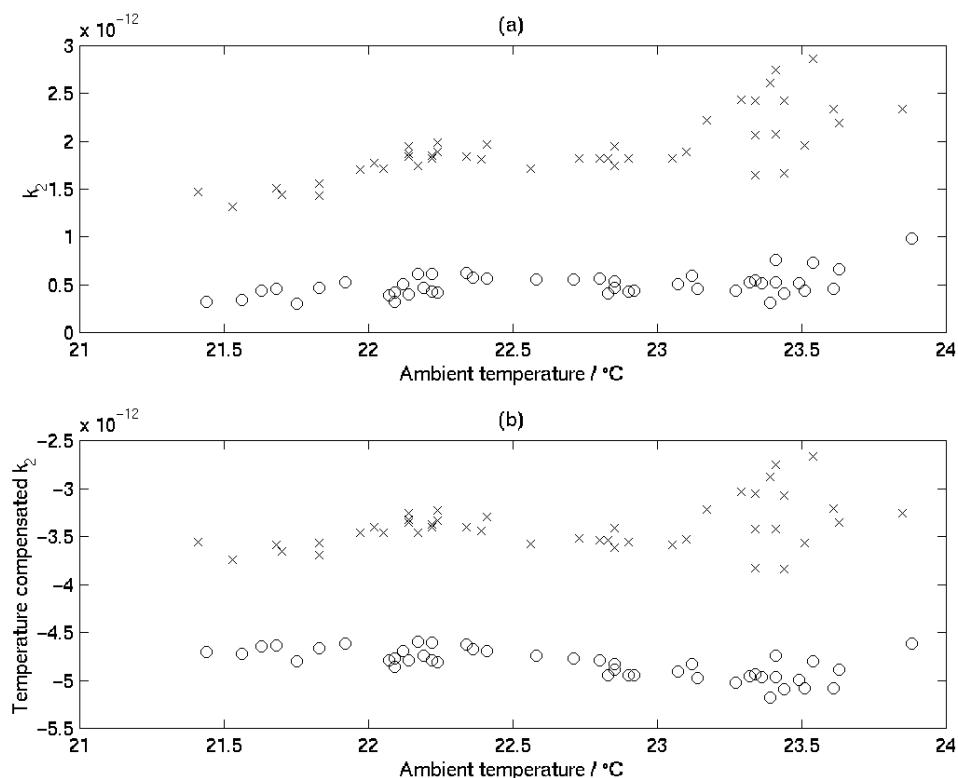


Figure 8.10 Plots of (a) fitted values of k_2 , and (b) temperature compensated fitted values for k_2 ($k_{2,0}$ as described in equation (8.1)). The x's represent values extracted from acetone responses and the o's represent values extracted from isopropyl alcohol responses, all from sensor 1. Note the reduced within-group variation in the values of k_2 for plot (b) compared with plot (a).




Figure 8.11 Plots of (a) k_2 and (b) temperature compensated k_2 , against response number, showing the slightly improved separation between the two classes after temperature compensation. The solid lines represent data corresponding to acetone responses, the dashed lines represent data corresponding to isopropyl alcohol responses, all from sensor 1.

It is evident from Figure 8.11 that the temperature compensation improves the separation of the two classes (acetone and isopropyl alcohol) slightly. However, there were possibly other factors which might affect the fitted parameter values, which were not controlled (most notably possible variations in the composition of the carrier gas (air from the room) or drift / poisoning of the sensors). Thus it would be foolhardy to rule out the possibility that the variations in the fitted parameters were not actually caused by the variations in temperature, but rather just happened to appear correlated with the temperature by chance.

8.4 Evaluation of the discrimination performance of parameter extraction versus steady state methods

It is clear from Figure 8.10 that the two analytes are clearly distinguishable from the values of k_2 , with or without temperature compensation. A simple rule-based algorithm based on the size of the estimate for k_2 from a curve fit to an unknown analyte response would produce a 100% successful classification rate (since the highest fitted value for an isopropyl alcohol response is lower than the lowest fitted value for an acetone response). However, using a classification test policy of training on the first half of the data, and testing on the second half (50% cross validation), there could be one misclassification towards the end of the testing set, depending on how the cut-off value was chosen. If instead the temperature-compensated values of k_2 are used, perfect classification rates can easily be obtained using 50% cross validation. Notice that the responses of only one sensor are required to achieve this perfect classification rate.

In order to compare this result with existing techniques, data from the same 45 response cycles (for all four sensors) were pre-processed as described in Section 3.3. Examples of the data which resulted are plotted in Figure 8.12.

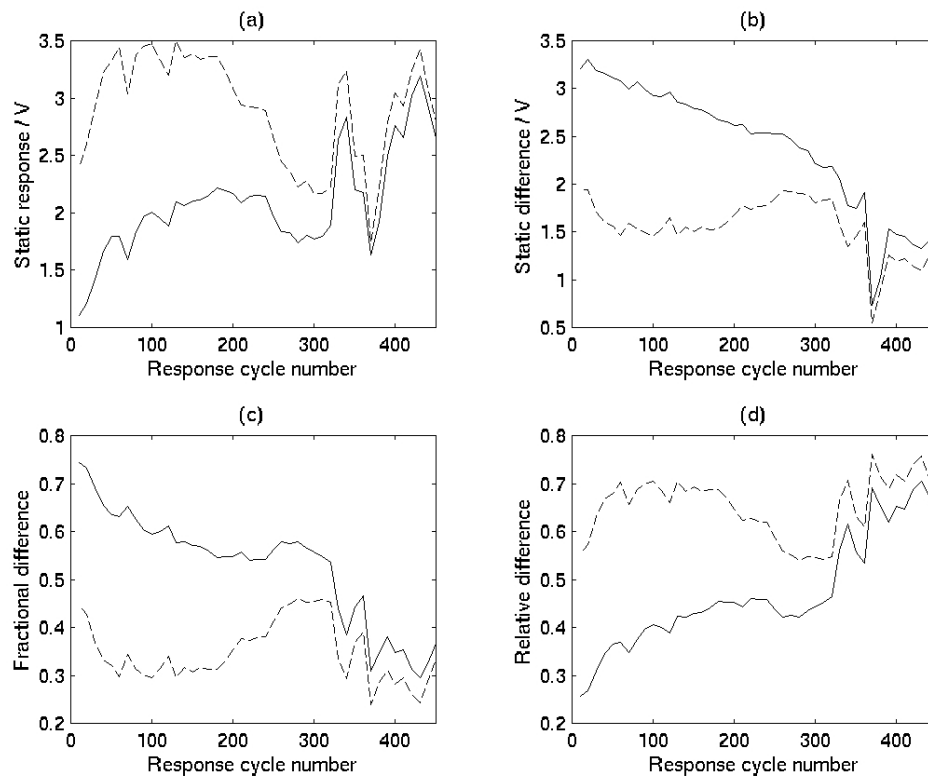


Figure 8.12 Plots of pre-processed static data from sensor 2: (a) static response, (b) static difference, (c) fractional difference, and (d) relative difference. The solid lines represent data corresponding to acetone responses, the dashed lines represent data corresponding to isopropyl alcohol responses.

It is clear from Figure 8.12 that the simple pre-processing applied to data from a single sensor does not produce the separation of the two classes that the physical parameter extraction technique does (compare the plots in Figure 8.12 with those of Figure 8.11). To test the success of the static techniques discussed in Chapters 3 and 4 when applied to the data from all four sensors, multi-input single-output linear time invariant black-box models were constructed using *Matlab*. As in the earlier chapters, various structures and orders were tested. It was found that black-box models using only information from single sensors were not able to produce 100% successful classification rates. However, when the information from all four sensors was used as inputs to the black-box models, low order (order one or two for each of the A and B polynomials) ARX models (see Chapter 3 for explanation of what this means) were

able to produce perfect classification rates (again these were tested using 50% cross validation).

8.5 Discussion and conclusions

In this chapter, it has been established that the model described and analysed in Chapter 7 is capable of producing excellent fits to real data (with r.m.s. errors as low as 4.58×10^{-3} V) from sensor exposures to simple odours (see Figure 8.4). However, despite the theoretical identifiability of the parameters in the model (which was established in Chapter 7), the actual numerical estimation of the parameters was far from straightforward.

Initial investigations showed that it was in fact very difficult to reproduce the qualitative shape of the start of the response curves (specifically the gradual rather than sharp drop of the sensor output) using most initial estimates for the parameters. Because of this, it was initially thought that the model might be unable to reproduce this feature, whilst using the simple (step change) form for the gas input to the main chamber. A model for a diffused input to the chamber (as discussed in Chapter 6) was incorporated in the software, and the shape of the start of the sensor response was easily reproduced. It was only after extensive investigations with both forms of input that it was found that with appropriate initial parameter estimates, the shape of the response can be reproduced using a square odour input function. The use of a square input function is preferable to a diffused input. This is because the inclusion of a diffused input in the model introduces at least one more parameter which must be estimated. Also, experiments with different gas flow rates show that the slow switch-on is present in the response even at relatively fast gas flow rates – suggesting that the qualitative shape at the start is (at least partly) a result of the dynamics of the sensor response mechanism rather than just a consequence of diffusion in the pipes / chamber. It is of course possible that diffusion in the pipework and / or chamber is a significant factor in the shape of the sensor response near to the start of the response. However, from the equipment and experiments used, it is impossible to isolate the two possible effects and model them both accurately.

Ideally, the sensors could be tested in a system whereby the flow of the gaseous analyte over the sensors could be controlled and known precisely, and where the volume of the chamber surrounding the sensors was minimal, so that the odour input to the model would be very close to the square profile that would be the input signal to the valve system. This would enable the dynamics of the electrochemical sensor responses to be isolated from the gas flow dynamics and thus modelled truly independently. Then the responses of the same sensors in a more usual (and practical) system (such as that considered here) could be modelled more accurately, including the gas flow dynamics.

Once initial parameter estimates had been found which enabled the accurate reproduction of the ‘on’ portion of the sensor response (using a square input function), restricting the parameter search to those parameter sets narrowed the search slightly, and slightly improved the empirical ‘identifiability’ of the system. However, it was still the case that choosing different initial estimates and bounds for various parameters resulted in the optimisation routines used by *Berkeley Madonna* finding significantly different parameter estimates, many of which produced visually (and numerically) acceptable fits to the same data set. This undesired flexibility in the choice of parameters was largely overcome via the use of the fact that the majority of the parameters in the model depend only on the sensor and system, and not the analyte that the sensor is exposed to. This means that these parameters should remain the same over multiple responses (for the same sensor, but different samples). Fitting parameters to multiple sensor responses to different analytes, whereby the sensor / system parameters were fixed to be the same for each response, was not very successful. This is possibly attributable to un-monitored external influences or experimental variations (including variations in the composition of the carrier gas, which was air from the laboratory). It is also possible that the reason for this is simply that the model, based as it is on assumptions and simplifications of the actual system, does not perfectly describe the system, and so in practice, the estimates for the supposedly analyte-independent parameters are actually affected by the analyte being sensed. Realistically the latter reason is perhaps the most likely.

The most practical compromise to enable parameter estimates to be used for odour classification, was to constrain the sensor / system parameters to lie within certain

bounds. This allowed sufficient flexibility for the model to accurately fit data from sensor responses to each analyte, but constrained the estimates enough so that repeatable and reproducible estimates for the (less constrained) analyte parameters could be extracted from the sensor responses. Using this technique, the key (for analyte identification purposes) model parameter k_2 was extracted from sensor responses over the course of a long-term experiment. It was found that the fitted values for this parameter were sufficiently well separated (between the two different analytes tested) so that perfect classification rates were easily obtainable using the data from just a single sensor.

It was found that the fitted values of the analyte parameter k_2 were correlated with the measured ambient temperature. However, using this correlation to produce temperature-compensated k_2 values produced only a slight improvement in the stability of the estimates for the parameter.

The same data were tested using the linear time-invariant black-box models investigated in detail in Chapters 3 and 4. It was found that the black-box models for static data were not able to match the discrimination performance of the parameter extraction technique when applied to only a single sensor, but when all four sensors were used for the inputs to the inverse black-box models, they were also able to produce perfect classification results. As a further test of its potential, in the following chapter the parameter extraction technique is applied to data from complex biological odours.

Chapter 9

Parameter extraction technique applied to cyanobacteria data

In this chapter, the technique for the classification of odours based upon estimates of parameters in a physical model describing the electronic nose system, developed in the previous chapter, is applied to a more realistic classification problem. The data from cyanobacteria experiments, used previously in Chapters 3 and 4, are again utilised to provide a comparison between the black-box models of the earlier chapters, this new parameter extraction technique, and the artificial neural networks used elsewhere.

This provides a challenging test for the technique, and thus a measure of its potential, as well as an opportunity to investigate how best to use the extracted parameters from an array of sensors (rather than just the required single sensor for the data in Chapter 8) to classify odours.

The application of this technique to complex odours requires slightly different assumptions than in the previous chapter. In the experiments considered in Chapter 8, the odours to which the sensors were exposed consisted only of single chemical species (with possibly some water in addition). In the cyanobacteria experiments considered in this chapter, the headspaces of the samples consist of many different chemical species. Attempting to expand the model of the previous chapter to explicitly

model the responses of the sensors to each of the many chemical species present would be futile. Even increasing the number of analytes modelled (by a few) would render the model almost useless, since there would be no realistic chance of obtaining accurate or consistent estimates of the many parameter values in the model. The experimental set-up and the data collected from it simply do not support the use of such a complex model. Instead the model is applied in exactly the same way as in Chapter 8, with the assumption that the response of a given sensor to the complex odour containing a mixture of many chemical species is approximately the same as the response to a simple odour. Because of this, the physical parameters k_2 and r_{an} extracted from the sensor responses will now no longer correspond directly to a (scaled) rate constant and a (scaled) rate of introduction of the target analyte molecules to the vicinity of the sensor. Instead they will correspond to ‘average’ or ‘effective’ values for these physical parameters for the odour as a whole (being made up, as it is, from many different compounds).

9.1 Cyanobacteria strain identification

The first test for the classification technique is the identification of the strain of a cyanobacteria colony. The data set is exactly the same as that used in Chapter 3, as is the task. The bacteria must simply be classed as ‘toxic’ or ‘non-toxic’. The linear time invariant black-box models used in Chapter 3 were very successful, obtaining a 100% successful classification rate using steady state data, as did the artificial neural networks used elsewhere (Shin *et al.* 2000). Thus if this new technique is to be regarded as successful, perfect (or at least near-perfect) classification rates are required on these data.

9.1.1 Experimental data

The experiment was designed to evaluate the ability of an electronic nose system to discriminate between two strains of cyanobacteria (blue-green algae), one toxic and the other non-toxic.

The headspaces of separate cultures of the two strains of cyanobacteria, grown in a nutrient medium (BG11), were sampled periodically by an electronic nose system over 40 days (Shin *et al.* 2000). The nose system used consisted of six commercial metal oxide resistive odour sensors (Alpha MOS, France), and two other sensors to monitor ambient temperature (LM35CZ, National Instruments) and humidity (MiniCap 2, Panametrics) (Shin *et al.* 2000). The repeated exposure cycle was as follows:

- 23 min 20 s - medium only
- 2 min - medium with toxic *microcystis aeruginosa* PCC 7806 strain
- 23 min 20 s - medium only
- 2 min - medium with non-toxic *microcystis aeruginosa* PCC 7941 strain.

The outputs from the sensors were sampled every 10 seconds, producing 350,358 data vectors corresponding to 1,150 exposure cycles.

9.1.2 Initial model fitting and parameter estimations

The experiment considered here was performed using a modified commercial instrument, in slightly less controlled conditions than the experiment studied in the previous chapter. The experiment took place in an active laboratory in the Department of Biological Sciences at the University of Warwick, where the ambient temperature and gas composition of the room varied considerably over the course of the experiment. For these reasons, the data produced from the experiment fluctuate more and are somewhat more noisy than those considered earlier. These fluctuations are shown in Figure 9.1. Note also that the sensor outputs were also sampled only once every 10 seconds, instead of once per second. A typical sensor response to an exposure to the headspace of the toxic bacteria is plotted in Figure 9.2.




Figure 9.1 The output of sensor 2 over the whole experiment, showing the diurnal and longer term variations observed.




Figure 9.2 A response of sensor 3 to exposure to the headspace gas from toxic bacteria sample.

Parameter fits to the model described in Chapter 7 and used in Chapter 8 were performed on a single response of each sensor to each sample of the two sample headspaces. As in Chapter 8, for each sensor, bounds for the parameters k_{1f} , k_{1b} , c_{ox} , r_{ox} , N and q were established, within which acceptable model fits were produced for both classes of odour and over the range of the data set. An example curve fit to a response of sensor 3 to an exposure to the headspace of the toxic sample is plotted in Figure 9.3.




Figure 9.3 The response of sensor 3 to an exposure to the headspace gases from a sample of toxic bacteria sample, together with a model fit to the data. The data points are plotted as x's, the model fit as a solid line. The r.m.s. error for the fit was 8.32453×10^{-3} V.

With the data analysed in Chapter 8, the values of the parameters k_2 and r_{an} (which were expected to be strongly analyte-dependent) were allowed to vary over a wide range during the curve-fitting process. It was found that when the sensor-dependent parameters were restricted to fairly narrow ranges, the values of k_2 and r_{an} were fairly well determined by the response curves. However, with the cyanobacteria data considered here, these crucial parameters seemed less empirically identifiable. It was found that, for a given sensor response curve, the value of r_{an} could be fixed almost arbitrarily with little or no difference to the success of the resulting curve fit obtained from varying only k_2 (and the other parameters within their bounds). This contrasts with the results of Chapter 8. This apparent lack of empirical identifiability could perhaps be caused by poor choices for the bounds on the sensor / system parameters, or by a lack of information in the data set itself. The former reason is plausible, though

extensive investigations failed to find a set of parameter bounds that produced more favourable results. There are differences between the cyanobacteria data considered here, and the data from experiments involving simpler analytes, considered in Chapter 8. The use of the model for a simple (i.e. single chemical species) odour for this experiment (which involves a complex odour containing many chemical species) involves, as mentioned previously, some possibly large assumptions. If these assumptions produce a poor description of the actual process, it is feasible that the success of parameter estimates might be affected. Another difference between this data set and that considered in Chapter 8 is the data sampling rate. In the experiment of the previous chapter, the sensor outputs were recorded once per second. For this cyanobacteria experiment the outputs were recorded only once every ten seconds. This is because the experiment was actually designed to test the ability of neural networks to classify the odours based only on static data (Shin *et al.* 2000). This had the result that the faster changes in the sensor responses (typically at the start of the response curve) were poorly recorded. It could be the case that this portion of the response is important for accurate determination of the appropriate values for the parameters k_2 and r_{an} .

To partially test this hypothesis, a section of data from the experiment described and analysed in the previous chapter was resampled at a lower sampling rate. Every fifth data point was kept, and the rest discarded. Parameter estimation runs were performed with this new data set. It was found (as with the cyanobacteria data) that the values of k_2 and r_{an} were no longer empirically identifiable, even when the remaining parameters were quite tightly bound (as before). The value of r_{an} was fixed to lie within five different bounded intervals (200,000 to 205,000, 300,000 to 305,000 ... 600,000 to 605,000), and the value of k_2 allowed to vary freely to fit the data. For each of the five parameter estimates, visually and numerically, acceptable fits were obtained. The r.m.s. errors for the five fits varied only between 4.29136×10^{-3} V and 4.68943×10^{-3} V. This would seem to support the suggestion that the reduced sampling rate is at least partially responsible for the observed lack of empirical identifiability of these two parameters.

9.1.3 Parameter estimations for responses throughout the whole experiment

In Chapter 3 the modelling techniques were implemented in *Matlab*, and thus the testing was carried out using *Matlab* programs to automate the process. This enabled the use of the whole data set (all 1,150 exposure cycles). The technique used in this chapter requires parameter estimations to be run manually for each sensor response, so to use all of the data would require 13,800 manual parameter estimation runs. This is not feasible, so instead a subset of the data set was used. Every fiftieth response was extracted from the data set, producing 46 response curves for each of the six sensors; 276 response curves in total. As mentioned in the previous section, the parameter values k_2 and r_{an} were not empirically identifiable from the response curves produced, so the value of r_{an} was constrained to lie within fairly tight bounds (the value was fixed to lie between 2×10^5 and 2.05×10^5). This represents an assumption corresponding approximately to the strength of the odour remaining fairly constant throughout the experiment. The result of this is that the fitted values of the remaining analyte-dependent parameter, k_2 , would be expected to vary with changes in the strength of the odour signal. This is undesirable when k_2 is to be used for odour discrimination but unfortunately unavoidable, since the alternative would be to constrain the rate constant k_2 more tightly, and to allow the value of r_{an} to vary for different sensor responses. That would constitute an assumption that the electrochemical rate constants for different samples remain similar, and consequently to attempt discrimination based purely (up to the accuracy of the model) upon a parameter which corresponds to the ‘strength’ of the odour. This may yield success for simple experiments but would be unlikely to be useful for more complex experiments or real-world applications. With r_{an} constrained, the value of k_2 became well-determined for each given response curve.

Curve fits were carried out to each of the 276 response curves extracted from the data set. The fitted values for k_2 were noticeably less stable than those for the simpler experiment of the previous chapter, as shown in Figure 9.4.

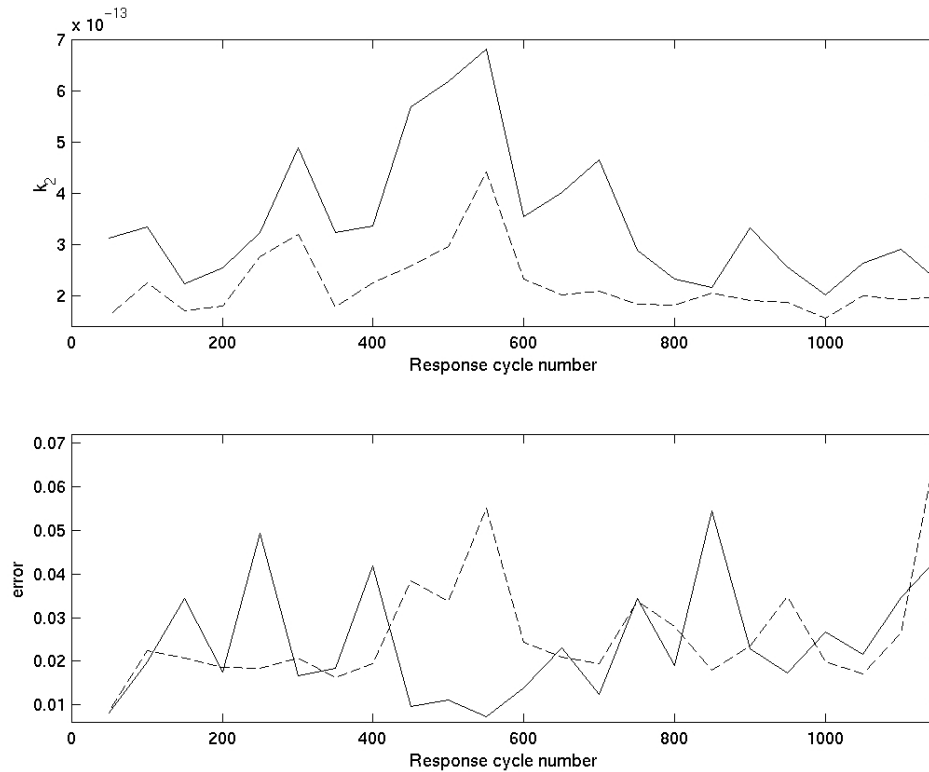


Figure 9.4 Plot of fitted values for k_2 (top) and the associated r.m.s. errors for curve fits (bottom) to responses of sensor 3 over the duration of the experiment. The solid lines represent data corresponding to responses to the headspace of toxic bacteria, the dashed lines to non-toxic bacteria responses.

The instability of the estimates for the parameter k_2 produced from the model curve fits means that a simple rule-based algorithm for discrimination based only on the estimate of the parameter for a single sensor will not yield acceptable discrimination performance. Instead, a method for discrimination based on the values extracted from the responses of all six sensors in the array must be used. As discussed previously, (see Chapter 2) there are many approaches which might be applied (since effectively the data extracted from the responses is of a similar nature to the traditionally-used pre-processed steady state information). Since the pre-processed data have already been analysed using linear time invariant black box models (see Chapter 3), this technique is applied to the extracted dynamic parameter values, for easy comparison of the discrimination power of the information within the data.

In Chapter 3 all 2300 response curves were used in the training and testing of the black box models. For fair comparison of discrimination performance, the pre-processed data corresponding to the 276 sensor responses which were used for curve fitting were extracted from the full pre-processed data set. These were then used to test black box models as in previous chapters, using a variety of structures and orders (see Chapter 3 for details of the technique used) and again 50% cross validation. It was found that after being trained on the first 23 pre-processed data vectors, simple low order (order two) FIR models were able to produce the correct classifications on the remaining 23 test vectors using the absolute response and relative difference pre-processing algorithms. The static difference and fractional difference algorithms successfully classified 22 and 21 of the test vectors respectively.

Using exactly the same technique applied to the fitted k_2 parameter values produced 22 out of 23 correct classifications using an FIR model of order two. Including an A polynomial in the model structure to form an ARX model produced perfect classification results on the testing set (using order one for the B polynomial, and order 3 for the A polynomial).

Investigations of the relationships between the fitted values of k_2 and the chamber temperature showed that the parameter estimates for the sensor responses to the headspace of the toxic bacteria were correlated with the chamber temperature with correlation coefficients of up to 0.8031 (varying slightly between the different sensors). The k_2 values for the sensor responses to the headspace of the non-toxic bacteria sample were correlated with the temperature with correlation coefficients of up to 0.7759. Taking all of the responses (to both toxic and non-toxic) together produced correlation coefficients of up to 0.6653. As in Section 8.3.2, linear regression was used to produce temperature compensated k_2 values. These compensated values are plotted in Figure 9.5. Note that since the linear fit is made over a small temperature range, a negative value for the temperature compensated k_2 is not cause for concern – the compensated k_2 does not have the same physical interpretation as k_2 itself. Using these temperature compensated k_2 values for the black box model inputs produced perfect classification success using an FIR model of order two. Interestingly, this perfect success rate was repeated when using only the temperature compensated k_2 estimate from a single sensor.

The correlation coefficients between the fitted values of k_2 and the output of the humidity sensor were low – at most 0.3105. This was deemed too low for correction of the parameter values to be worthwhile.

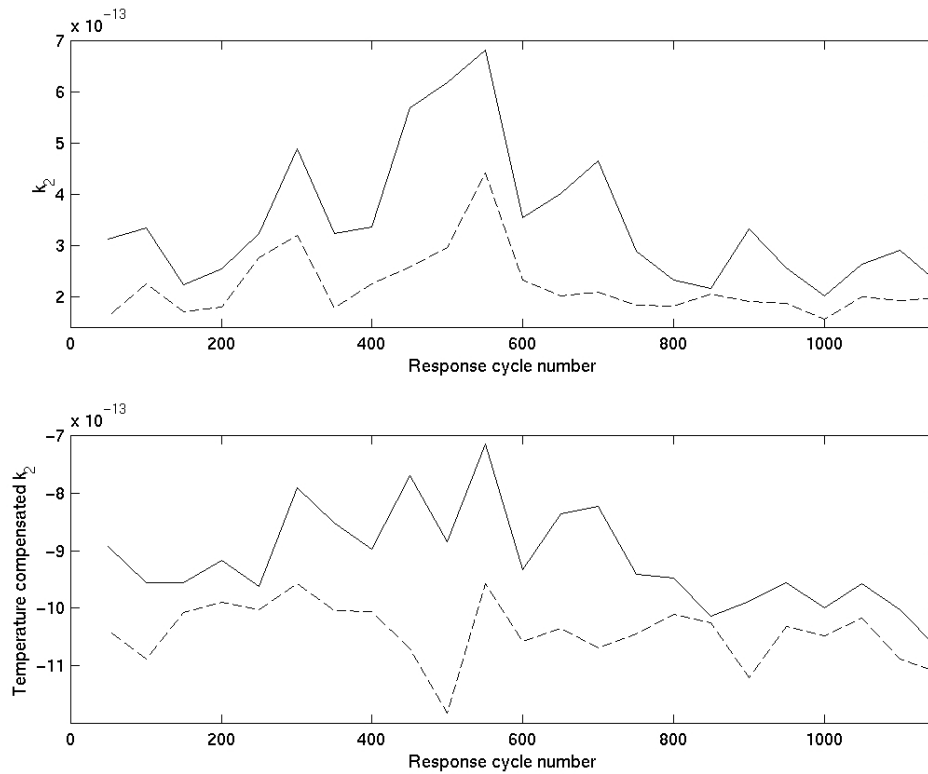


Figure 9.5 Plots of (a) k_2 and (b) temperature compensated k_2 , against response number, showing the slightly improved separation between the two classes after temperature compensation. The solid lines represent sensor responses to toxic bacteria exposure, the dashed lines represent data sensor responses to non-toxic bacteria, all from sensor 3. The temperature compensation changes the physical interpretation of the value, so negative values are not unexpected.

9.2 Cyanobacteria growth phase identification

The final test for the classification technique is the identification of the growth phase of a cyanobacteria colony. The data and task are the same as those used in Chapter 4.

Each bacteria sample must be placed into one of four growth phase classes. The linear time invariant black-box models used in Chapter 4 were only moderately successful, obtaining a maximum successful classification rate of 82.3%. The artificial neural networks used elsewhere (Shin *et al.* 2000), obtained a maximum success rate of 95.1%. The aim of the analysis of this chapter is to compare the efficacy of variations on the new parameter extraction technique with the results obtained previously.

9.2.1 Experimental data

As discussed in Section 4.2, the experiment considered here involved a single (toxic) strain of cyanobacteria, monitored over a 40 day period. As well as electronic nose data, information concerning the mean size of the bacterial cells and the biomass present in the cultures was recorded using a CellFacts instrument (Microbial Systems Ltd.). This enabled the identification of the four distinct growth phases through which the bacteria pass during their life cycle.

The experiment in question was intended to not only test the ability of an electronic nose to discriminate between the different growth phases of a cyanobacteria strain, but also to investigate the reproducibility of the measurements and success rates. For this reason the experimental system consisted of three vessels, two containing nominally identical cultures of toxic *microcystis aeruginosa* PCC 7806 in nutrient medium (BG11), and one reference vessel containing only the nutrient medium. The headspaces of these vessels were connected via a system of pipes and computer-operated valves to an electronic nose system. The nose system used consisted of six commercial metal oxide resistive odour sensors (Alpha MOS, France), and two other sensors to monitor ambient temperature (LM35CZ, National Instruments) and humidity (MiniCap 2, Panametrics) (Shin *et al.* 2000). The repeated exposure cycle was:

- 50 min - medium only
- 5 min - medium and toxic *microcystis aeruginosa* PCC 7806 strain sample 1
- 50 min - medium only
- 5 min - medium and toxic *microcystis aeruginosa* PCC 7806 strain sample 2.

The sensor outputs were again sampled every 10 seconds, producing 361,698 data vectors corresponding to 548 full exposure cycles.

The information collected using the CellFacts instrument was used to produce a 'correct' classification vector for the data, identifying the four growth phases of the bacteria using understanding of the biological processes involved. It should be noted that the boundaries between the growth phases were by no means sharp, making the manual identification of the phases subjective, and thus the noise data classification problem highly non-trivial. Some of the data from the CellFacts instrument are plotted in Figure 9.6.




Figure 9.6 Plot of data from the CellFacts instrument for the growth phase identification experiment. The upper plot shows the general increase in biomass (cell counts) with time. The lower plot shows the variation in mean size of the bacteria cells with time. The four growth phases (lag, log, stationary and late stationary), are labelled I to IV in each plot, respectively.

9.2.2 Initial model fitting and parameter estimations

The data considered in this section were gathered using the same electronic nose equipment as was used for the experiment of Section 9.1. As can be seen from Figure 9.7, there were considerable variations in the sensor outputs, both diurnally and more long-term.




Figure 9.7 The output of sensor 3 over the whole growth phase identification experiment.

As in Chapter 8 and subsequently Section 9.1 of this chapter, the model discussed in Chapter 7 was fitted to single responses of each sensor. As before, bounds for the parameters k_{1f} , k_{1b} , c_{cox} , r_{ox} , N and q were established such that acceptable model fits could be produced for sensor responses to each of the two samples over the entire course of the data set. In fact, the parameter bounds established in Section 9.1.2 for data from the previous experiment were used as these gave acceptable fitting

outcomes. An example of a model fit to a sensor response from this data set is plotted in Figure 9.8.




Figure 9.8 The response of sensor 6 to an exposure to the second toxic cyanobacteria sample, together with a model fit to the data. The data points are plotted as x's, the model fit as a solid line. The r.m.s. error for the fit was 6.91129×10^{-3} V.

9.2.3 Parameter estimations for responses throughout the whole experiment

As in Section 9.1, the data set produced by the cyanobacteria growth phase experiment was too large for manual parameter estimation runs on each of the sensor responses to be feasible. This experiment produced 548 full response cycles, so every twentieth response was extracted and used for parameter estimation. This gave a total of 336 response curves (two cyanobacteria samples, six sensors, 28 responses for each) for manual parameter estimation runs.




Figure 9.9 Plot of fitted values for k_2 (top) and the associated r.m.s. errors (bottom) for curve fits to responses of sensor 2 over the duration of the experiment.

If Figure 9.9 is compared with Figure 9.6, there is no clear correlation between the fitted k_2 values and the biological parameters used to produce the ‘correct’ growth phase classifications.

In order to test the discrimination power of the extracted rate constants, black-box models were again used. As in Section 4.3, the data were randomly reordered, and then multi-input multi-output models were trained and tested on different portions of the data set. Due to the relatively small amount of data available (only 28 response data vectors for each of the two data sets), 50% cross-validation was not considered practical – 14 data vectors spread over four classes is too few. Instead it was decided to use 18 data vectors to train the models, and ten for testing. This is still many fewer than is ideal, but time limitations prevented the extraction of more parameter values.

MIMO models of various structures (FIR, ARX, MAX and ARMAX) were trained and tested using the extracted k_2 parameter values for both the Toxic 1 and Toxic 2 data sets. Due to the small size of the data sets, each model structure was trained and tested on 100 different random reorderings of the 28 data vectors. Static pre-processed data sets were formed for the same reordered data sets and tested in exactly the same way, for comparison.

For each of the data sets, the best model structure was found to be the simplest: an FIR model of order 1. This slightly contrasts with the results of Section 4.3, where FIR models of orders between one and four were found to produce the highest success rates (using static data from the whole growth phase experiment data). This might be attributable to the greatly reduced size of the data set. The success rates obtained using the extracted parameter values, and the static (pre-processed) data are given in Table 9.1.

% successful classification using	Toxic 1:	Toxic 2:
k_2 :	55.0	60.3
Absolute response:	68.0	60.0
Difference:	73.2	70.9
Relative difference:	81.6	82.0
Fractional difference:	78.6	80.1

Table 9.1 Comparison of the average classification success rates using FIR models of order one for extracted k_2 values with those using static (pre-processed) data for the classification of cyanobacteria growth phase. The rates given are averages over 100 different random reorderings of the data set. For each reordering, the first 18 data vectors were used to train the model, and the remaining 10 to test it.

It is clear from Table 9.1 that the models for the extracted rate constant k_2 were outperformed significantly by the models for the simple pre-processed static data. The correlation coefficients between the fitted values of k_2 and the chamber temperature varied between 0.0535 and 0.5159 for the different sensors and data sets. Since these

figures are low, it is unsurprising that attempting to remove variations in the fitted parameter by linear temperature compensation was unsuccessful (producing, in fact, lower success rates than the not-compensated values).

9.3 Discussion

In this chapter, the parameter extraction techniques established in Chapter 8 were used to fit analyte-dependent rate constants to dynamic sensor responses from biological experiments involving cyanobacteria samples. In order to do this, assumptions were made regarding the application of the model described in Chapter 7 (which describes the interaction of a simple odour (i.e. an odour containing a single chemical species)) to complex odours. This was necessary, since models describing the interactions between the sensors and all potential analyte species would be too complex to be useable.

The parameter estimation techniques developed in Chapter 8 were applied to data resulting from cyanobacteria experiments. These data have been previously analysed elsewhere using artificial neural networks (Shin *et al.* 2000), and linear black-box models (in Chapters 3 and 4), and thus provided a worthwhile test for the technique of identifying complex odours via parameter estimation.

The techniques were first applied to data from a bacterial strain identification experiment, where the aim was to classify a sample of bacteria as ‘toxic’ or ‘non toxic’. As in Chapter 8, it was found that it was necessary to constrain the parameters k_{1f} , k_{1b} , c_{eox} , r_{ox} , N , and q to lie within fairly tight bounds. However, unlike in Chapter 8, with these parameters constrained the remaining (analyte-dependent) parameters k_2 and r_{an} were not found to be well-determined by the data. There are a number of possible reasons for this – the mostly likely seems to be that the cyanobacteria data were sampled only once every 10 seconds, compared with the once per second sampling rate of the experiments of Chapter 8. This reduced the dynamic information contained in the sensor response data, especially near to the start of the sensor response, where the initially gradual drop in sensor output value was almost entirely missed. This explanation was partially validated by resampling the data of Chapter 8 at a lower

frequency and performing parameter estimations with these new data. A similar lack of empirical identifiability was then found with the resampled data.

In order to obtain consistent parameter estimates, the value of the parameter r_{an} was then also constrained to lie within narrow bounds, leaving k_2 then well-determined from each sensor response curve, though potentially susceptible to changes in the odour signal strength.

The parameter k_2 was then estimated for 276 response curves from the cyanobacteria strain identification experiment. Plots of extracted k_2 values for the two bacteria strains showed variations in the estimates that were consistent between the two strains – suggesting that they result from variations in ambient conditions (especially temperature, which might affect the sensors responses to a given stimulus, and possibly more significantly the activity levels of the bacteria and the evaporation rates of odourant chemicals given off by the samples). The variations in the estimates for k_2 over the duration of the experiment meant that a simple rule-based discrimination algorithm (as was successful for the data in the previous chapter) would no longer yield successful results. Instead, linear black box models were applied to the extracted k_2 vectors (one value per sensor). The discrimination ability of the technique was tested using 50% cross validation – training on 23 test vectors, and testing on the remaining 23. Low order (order two) FIR models for static (pre-processed) data were able to correctly classify between 21 and 23 of the 23 test data points depending on the choice of pre-processing algorithm. An FIR model of order two for the k_2 data produced 22 out of 23 successful classifications. Extending the model to an ARX model (order one for the B polynomial, and two for the A polynomial) produced 100% successful classification on the test set.

Investigations of the relationships between the fitted k_2 values and the outputs from the ambient temperature and humidity sensors showed a significant correlation between k_2 and the temperature, but no great correlation between k_2 and the humidity sensor output. Linear temperature compensation was performed (as described in Section 8.3.2) on the fitted k_2 values. The compensated k_2 values were tested using black box models, as before, and they were found to produce perfect (23 out of 23) successful classification results on the test set using an FIR model of order two. This represents

an improvement over the results obtained using the not-compensated k_2 values, since the model complexity required to produce 100% success is less. This might be attributed to the fact that the bacterial activity levels could be expected to increase with increasing temperature, so that the ‘strength’ of the odour would be expected to vary with changes in temperature. Compensating for changes in k_2 with temperature might thus provide an indirect (and partial) compensation for the ‘strength’ of the odour – something otherwise lacking due to the constraints placed upon r_{an} . It should be noted that the improvement in the model successes effected by this temperature compensation were slight, and should not be regarded as conclusive evidence that such a practice should always be expected to be successful.

For the second test data set, the parameter k_2 was estimated for 336 response curves from the cyanobacteria growth phase identification experiment. For this experiment, the aim was to classify samples from two nominally identical toxic bacteria colonies into one of four growth phases, as described in Section 9.2.1). As with the strain identification experiment, very good model fits were obtained using a range of different parameter values for a given response, so the system parameters and the analyte-dependent parameter r_{an} were again constrained to lie within fairly tight bounds.

Due to the relatively small number of extracted rate constants (336 spread across two cultures and six sensors – thus 28 k_2 vectors per culture, incorporating all four growth phases) a 50% cross validation model testing regimen was deemed inappropriate. Instead, for each culture, 18 vectors were used for training and 10 for testing. Clearly the order of the samples was crucial due to the temporal progression of the cultures through their growth phases, so 100 different random reorderings were used, and the average results taken.

MIMO models of various structures and orders were used, and compared with their equivalents using corresponding static data. Simple FIR models of order one were found to produce the best success rates, of 55% for the ‘Toxic 1’ culture and 60.3% for the ‘Toxic 2’ culture using the extracted k_2 values. These results were inferior to those obtained using static data – the best of which were 81.6% for the ‘Toxic 1’ culture and 82% for the ‘Toxic 2’ culture (both using the relative difference pre-processing

algorithm). Temperature compensation was found to produce no improvement in success rates for the models for extracted k_2 values.

9.4 Conclusions

There are a number of possible causes for the relative lack of success of the models for the extracted k_2 rate constants. As mentioned before, the sampling rate of these data was too low to capture some of the features of the dynamic responses of the sensors. This may have been the cause of the noted lack of empirical identifiability of the analyte-dependent parameters in the model, and thus allowed external variations (which might have been taken care of in corresponding variations in r_{an}) to be carried through into variations in the fitted values of k_2 . The theoretical (at least local) identifiability was established in Chapter 7, suggesting that all of the parameters in the model (both system and analyte parameters) could perhaps be estimated from experiments without the need to constrain many of the parameters to lie within tight bounds. This result was not mirrored in the practical investigations of this (and the previous) chapter. It may be the case that by employing different experimental protocols (such as a higher sampling rate, different form of analyte input) this may be rectified. Alternatively, reliable estimates for some of the parameters might be obtained through separate experiments (for example experiments involving no analyte as such, but varying concentrations of oxygen in the carrier gas to estimate the oxygen reaction rate constants more accurately). Also, the experimental set-up used had significant volumes in the gas delivery system, making the dynamics of the gas delivery system difficult to isolate from the dynamics of the sensor responses. The odour input to the sensors was assumed to have a square profile, though this is almost certainly not accurate due to diffusion effects in the pipework and sensor chamber. If a system could be used with very low volumes between the sample and the sensors then the dynamics of the sensor responses could almost all be attributed to the electrochemical reactions rather than gas diffusion effects.

There is also perhaps scope for performance improvements through dynamic pre-processing of the data, to attempt to remove effects of variations in the strength of the odours (resulting from temperature changes or other influences). Array normalisation

or scaling with a reference sensor prior to parameter estimation might produce more consistent estimates for the reaction rate constant k_2 . It should also be noted that the training and test sets were smaller than desired, especially for the growth phase experiment. The development of a more automated parameter estimation procedure would allow the evaluation of the technique using much larger data sets.

To conclude, the technique of identifying odours based on extracted rate constants from a physical model for the electronic nose system which showed promise when applied to simple odours was able to produce perfect classification results for the data from the cyanobacteria strain identification experiment. However, it did not deliver improved results for the more challenging cyanobacteria growth phase experiment. It has nonetheless provided valuable insights into the experimental design and data analysis tools which might be required to produce great improvements in the results produced by this technique.

Chapter 10

Conclusions

In this thesis, a variety of techniques for odour classification using the dynamic responses of arrays of metal oxide sensors have been developed and analysed.

10.1 Black box modelling

In Chapter 3, linear time-invariant black box models were applied to the analysis of data from a cyanobacteria strain identification experiment. The data had previously been analysed using artificial neural networks, in Shin *et al.* 2000, where the best-performing neural networks (a nonlinear Fuzzy ARTMAP network) produced 100% successful classification rates. Techniques for applying black box models for both static (steady-state) pre-processed and dynamic data were developed and tested on the cyanobacteria data. It was found that very simple, low order models were able to produce perfect classification rates using 50% cross validation of static data. With dynamic data, the black box models gave successful classification rates of up to 99.3%. The success rates achieved demonstrate the suitability of linear black box modelling techniques for both static and dynamic data for simple classification problems with complex biological odours.

In Chapter 4, the techniques developed in Chapter 3 were extended and applied to a second set of cyanobacteria data. Here the classification problem was more challenging – to identify the growth phase (of which there are assumed to be four) of a culture of a single strain of cyanobacteria. The boundaries between the four growth phases were found to be not very well-defined, and were subjectively assigned (based on data from the ‘CellFacts’ instrument) for the purposes of model training. Again, this data set had previously been analysed elsewhere using nonlinear neural networks. The best networks were reported to achieve a successful classification rate of 95.1% (Shin et al. 2000). For this application it was found that the linear black box models were not able to match the artificial neural networks. Black box models for the static data produced maximum success rates of 82.3%. Models for dynamic data achieved a 76.6% success rate.

Thus it has been demonstrated that the linear black box models, for both static and dynamic data, can achieve excellent classification success rates when analysing complex odours. However, when the set of classes is not well-separated, other techniques, such as the use of nonlinear artificial neural networks, may outperform them. For the data sets considered here, the black box models for static data generally produced better results than those for dynamic data. This might suggest that the static responses of the sensors in the array carry more significant discriminatory information than the dynamic, or perhaps just that the linear black box model structures used were better suited to the static data. Since the processing power required to implement these linear black box models is minimal, this technique might be an attractive choice for use in low cost handheld devices for relatively simple classification applications.

There are a number of possible extensions to the work on linear black box models presented here. One of the potential causes of the limitations of the technique is the linearity of the models, and the nonlinearity of the sensor response processes. There are nonlinear system identification techniques that might produce improved classification performances. For both the static and dynamic models, the criteria used during the parameter estimation (training) stage were not precisely the same as those used for evaluation of the model performance, so it is plausible that a modification of the training algorithms might also yield improvements. The time-invariant nature of the models may also have had a restricting effect on their performance, so another

possible extension of the work might be to investigate the use of time-varying or self-re-calibrating analogues of these models.

10.2 Models based on the physical system

In Chapter 7, a physical model for the electrochemical mechanism governing the response of MOS sensors was presented. This mechanism was converted into a system of coupled nonlinear ordinary differential equations. These equations were analysed and reduced to a simpler, equivalent system. The observability and controllability of the system were established using the ORC and CRC respectively. The identifiability of the system was investigated, and it was found that, after reparameterisation, the system could be shown to be at least locally identifiable. This provided the theoretical basis required for confident attempts to estimate the parameters in the model from experimental data.

10.3 Classification by parameter extraction

The experimental test system described in Chapter 5 was used to conduct some long-term experiments and produce high quality data from two simple odours: acetone and isopropyl alcohol. In Chapter 8, the model presented and analysed in Chapter 7 was used to produce curve fits to sensor responses to these simple odours. The model was found to be capable of producing excellent curve fits to the data (with a typical r.m.s error of approximately 0.1%). Despite the theoretical results obtained in Chapter 7, the task of estimating the parameters of the model to fit the experimental data was found to be challenging. It was found that, for a given sensor response curve (i.e. the response of a sensor to a single odour exposure), the values of the full set of model parameters were not well-determined. Thus highly satisfactory curve fits could be obtained using a range of different parameter sets. This posed a problem since the intention was to classify odours based upon parameter estimates calculated from responses to unknown odour samples.

The majority of the parameters in the system model relate to aspects of the system that do not depend (or at least are assumed not to depend) upon the odour sensed. Thus, it was found that these parameters could be constrained to lie within fairly tight bounds, with only the analyte-dependent parameters allowed to vary widely for different response curves. With the analyte-independent parameters constrained, it was found that the remaining analyte-dependent parameters, namely r_{an} and k_2 , were well-determined by the experimental response curves, and could be repeatably extracted from the data.

The parameter k_2 corresponds to a rate constant in the electrochemical reaction mechanism, so is the primary parameter which (according to the postulated model) distinguishes analytes. In contrast, r_{an} corresponds to the rate of input, or ‘strength’, of the analyte. For a laboratory-based experiment where the signal is consistent, the information carried in this parameter may aid with the separation of classes, however this could not be assumed to be the case in real-world applications where the strength of the signal might vary widely. The parameter k_2 was extracted from responses of a single MOS sensor over a period of about 13 days, during which time the sensor response changed considerably. It was found that excellent separation of the two classes was achieved, using only this single parameter, from a single sensor. Thus a simple rule-based algorithm applied to the k_2 data was able to produce a 100% successful classification rate. By contrast, none of the pre-processed steady-state parameters extracted from the same data showed comparable separation over the course of the experiment. It was also found that linear temperature compensation of the extracted k_2 values produced slight, though not significant improvements to the consistency of the extracted parameter values over the course of the experiment.

This showed that, at least for this data set, the extracted parameter k_2 provided a more stable description of the sensor response, and produced superior classification performance. However, the classification of a simple odour into two classes was not a particularly challenging test of the newly-developed technique (indeed when pre-processed steady state data from all four MOS sensors were analysed using the black box models described in the earlier chapter, perfect classification results were also obtained), so the technique was applied to more challenging odour discrimination applications.

In Chapter 9, the data from cyanobacteria experiments analysed in Chapters 3 and 4 was used to test the suitability of the newly-developed parameter extraction technique of Chapter 8 for the analysis of complex odours. This provided a difficult test for the new technique, as well as results for comparison with the other techniques previously used (i.e. the linear black box models for static and dynamic data earlier in this thesis, and the nonlinear artificial neural networks used elsewhere (Shin *et al.* 2000)).

As mentioned, in Chapter 8, once the nose-dependent parameters had been heavily constrained, the two remaining nose-dependent parameters k_2 and r_{an} were well-determined by the experimental data. For the bacteria experiment data, in Chapter 9, this was not found to be the case. It was found to be necessary to further constrain r_{an} in order to obtain repeatable estimates for k_2 from the curve-fitting process. This was undesirable since it could be expected to remove the ability of the technique to ‘automatically’ compensate for variations in the strength of the odour signal, and perhaps to a lesser extent, variations in the ‘perceived’ strength due to drift or poisoning of the sensor. However, it was unavoidable, since otherwise estimates for k_2 would have been inconsistent and not repeatable, and so odour classification based on k_2 would have been impossible.

With r_{an} constrained, estimates for k_2 were produced from responses of the six MOS sensors used. The estimated values for k_2 for individual sensors did not show good separation between the two classes, so linear black box models were applied to the extracted k_2 values from the full array of six sensors. Using an FIR model of order 2 gave 22 successful classifications out of 23 (using 50% cross validation). Extending the model to an ARX model produced a 100% successful classification rate. A correlation was found between the extracted k_2 values and the recorded temperature. This was used to produce temperature-compensated k_2 vectors. The use of these temperature-compensated values for k_2 produced 100% successful classification using an FIR model of order 2. For comparison, using the same sensor responses, standard steady state pre-processed data vectors were produced (as in Chapter 3) and tested using the same black box model structures. These were also able to produce success rates of up to 100%.

For the cyanobacteria growth phase data set, excellent curve fits were again obtained using the model. As with the strain data, it was necessary to constrain r_{an} to lie within fairly tight bounds. The parameter k_2 was estimated from 336 response vectors related to exposure to two (nominally identical) cultures. Thus there were 28 response vectors per culture, per sensor, spread over the four growth phases. Given this small number, it was decided to use 18 vectors for training, and the remaining 10 for testing. The choice of training and testing sets was made randomly, 100 times, and the average success rates used.

Black box models for the extracted k_2 values produced successful classification rates of 55% for the ‘Toxic 1’ culture, and 60.3% for the ‘Toxic 2’ culture. This did not compare well with the results obtained on the same data sets using pre-processed static data: 81.6% and 82% for the ‘Toxic 1’ and ‘Toxic 2’ cultures respectively. Temperature compensation of the extracted k_2 values produced no improvement in the success rates obtained.

There are a number of possible reasons for the lack of success of the parameter extraction technique when applied to the growth phase identification problem. One possible reason is the low sampling rate employed for the recording of the sensor data (0.1 Hz compared with 1 Hz for the simple odour experiments analysed more successfully in Chapter 8). It may be the case that this low frequency sampling was responsible for the lack of empirical identifiability, and this in turn prevented the automatic ‘compensation’ for variations in odour signal intensity (via free estimates for the parameter r_{an}). There was significant variation in the magnitudes of the sensor responses over the duration of the experiment. Some of this may be due to changes in the bacterial activity levels as a result of their changing growth phase, but some may be due simply to changes in the ambient temperature (resulting both in differing sample headspace concentrations, and in differing sensor sensitivities). Were these effects to be isolated, one might expect the growth phase classification performance to improve. Another possible reason for the relatively poor results might be simply the lack of training data used, or the method of analysing the extracted data vectors. It may be the case that linear black box models might not be the most appropriate for classification systems based on extracted k_2 values.

The research detailed in this thesis has suggested a number of ways in which the parameter extraction technique developed here might be improved. The physical configuration of the hardware used made it difficult to isolate the dynamics of the different components of the system. The relatively high volume of the gas delivery system matched with relatively low gas flow rates meant that diffusion and mixing effects in the system between the sample headspace and the odour sensors may well have had a significant effect on the recorded sensor outputs. Faster flow rates could not have been used since this would have degraded the sample headspace over the course of the exposure, so that the concentration being sampled was not constant. One solution to this would be to use much larger samples and headspaces, though this would not necessarily be feasible in all applications. A better solution would be to employ a different hardware configuration where there was very little volume between the sample headspace and the odour sensors. Another alternative would be to use faster-responding sensors, so that faster gas flow rates could be used with less degradation of the sample headspace.

In order to capture the relevant dynamic information for odour classification, the data recording rate must be sufficiently high to record the fastest relevant changes in the sensor outputs. It seems likely that this was not the case for the cyanobacteria strain and growth phase experiments (these experiments were designed with the intention of using only static information, hence the low sampling rate). The designers of future experiments should ensure that the sampling rate used is adequate. There is also perhaps scope for the use of extra experiments (perhaps involving varying the concentration of oxygen in the carrier gas) in order to obtain high quality estimates for some of the analyte-independent parameters in the model. This might produce better estimates for these parameters than attempt to estimate all of the model parameters at once, and thus make the analyte-dependent parameters more reliably and accurately estimable from experimental data.

There is a lot of potential for progress in the analysis of the dynamic data recorded. The parameter estimation technique used here required a great deal of laborious manual intervention and data manipulation in order to extract the parameters for each sensor's response curve. This fact severely limited the amount of data available for data model building and evaluation. A software system which could automate this

process and thus produce much larger data sets would permit a much more thorough investigation and evaluation of the potential of the technique. For the cyanobacteria data sets, the extracted k_2 vectors were analysed using linear black box models. This was done in order to provide a direct comparison with the previous work using pre-processed steady state data. However, there are a great many other techniques which might be applied to the k_2 data, which might yield success, such as artificial neural networks, SIMCA, etc. (see Chapter 2 for details of these techniques). It is also feasible that a system using some combination of steady-state parameters and extracted dynamic parameters might produce superior performance to systems using just one or the other.

10.4 Final remarks

To conclude, it has been shown that linear black box models for both static and dynamic data are capable of matching the best nonlinear artificial neural networks, and producing excellent classification results (up to 100% success rates) with the cyanobacteria strain identification data set. This suggests that this essentially simple technique might be useful in a wide range of applications involving complex odours, where the number of classes is small and reasonably well separated in sensor space.

It has also been shown that for the more challenging cyanobacteria growth phase identification experiment, the simple linear techniques employed here were outperformed by nonlinear neural networks elsewhere. This suggests a limitation on the capabilities of the black box modelling technique for this type of application. For this data set, it was also demonstrated that the black box models for the pre-processed steady-state data produced superior results to those for the dynamic data.

An ODE model for the dynamic responses of MOS sensors to odour stimuli was proposed and analysed. It was established that the model was (at least locally) identifiable, and thus potentially usable for odour classification by parameter extraction.

Excellent curve-fits of the model to experimental data were produced. A method for extracting the most significant analyte-dependent parameter from experimental data was developed. This method was applied with great success to data from experiments involving simple odours – producing 100% successful classification results. The method was also applied to cyanobacteria strain and growth phase experiments. For the strain identification data, once again, 100% success rates were achieved. However, for the growth phase identification problem, the success rates were only around 55% – 60%. This was inferior to results obtained elsewhere using nonlinear artificial neural networks, demonstrating that further work is required.

It is anticipated that, with developments in sensor technology, hardware design, and data processing, this technique could evolve into a more versatile and powerful tool for odour classification than any that are currently available. The approach of mechanistic modelling rather than the more common data-driven techniques, allows the incorporation of knowledge about the mechanisms involved, and permits the creation of a much more efficient, versatile and extensible odour classification system.

References and bibliography

- Aishima, T. (1991a). Discrimination of liquor aromas by pattern recognition analysis of responses from semiconductor gas sensor array. *Anal. Chem. Acta*, **243**, 293-300.
- Aishima, T. (1991b). Aroma discrimination by pattern recognition analysis of responses from semiconductor gas sensor array. *J. Agric. Food Chem.*, **39**, 752-756.
- Al-Khalifa, S., Gardner, J. W. and Maldonado-Bascón, S. (2001). Rapid multicomponent analysis using a thermally-modulated resistive gas microsensor and a discrete wavelet transform. In *Sensors and their Applications XI*, (Eds K.T.V. Grattan & S.H. Khan), Series in Sensors, (Series Ed. B.E. Jones), pp.103-108. IOP Publishing, Bristol.
- Amoore, J. E. (1977). Specific anosmia and the concept of primary odours. *Chem. Senses Flav.*, **2**, 267-281.
- Bargagna, G., Lazzarini, B. and Partridge, A. (2000). Fuzzy logic based classification of olive oils. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 127-132. IOP Publishing, Bristol.
- Beebe, K. R., Pell, R. J. and Seasholtz, M. B. (1998). *Chemometrics: a Practical Guide*. John Wiley and Sons, New York.
- Bellman, R. and Åström, K. J. (1970). On structural identifiability. *Math. Biosci.*, **7** pp 329-339.
- Boilot, P., Hines, E. L., John, S., Mitchell, J., Lopez, F., Gardner, J. W., Llobet, E., Hero, M., Fink, C. and Gongora, M. A. (2000). Detection of bacteria causing eye infections using a neural network based electronic nose system. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 181-196. IOP Publishing, Bristol.
- Brattain, W. H. and Bardeen, J. (1953). Surface properties of germanium. *Bell Syst. Tech. J.*, **32**, 1.
- Buck, T. M., Allan, F. G. and Dalton, M. (1965). Detection of chemical species by

- surface effects on metals and semiconductors. In *Surface effects in detection* (eds. J. I. Bregman and A. Dravnieks), pp. 147-163. Macmillan, London.
- Chapman, M. J. and Godfrey, K. R. (1996). Nonlinear compartmental model indistinguishability. *Automatica*, **32**, 419-422.
- Chappell, M. J. and Gunn, R. N. (1998). A procedure for generating locally identifiable reparameterisations of unidentifiable non-linear systems by the similarity transform approach. *Math. Biosci.*, **148**, 21.
- Chappell, M. J., Godfrey, K. R. and Vajda, S. (1990). Global identifiability of the parameters of nonlinear systems with specified inputs: a comparison of methods. *Math. Biosci.*, **102**, 41-73.
- Covington, J. A., Gardner, J. W., Toh, C., Bartlett, P. N., Briand, D. and de Rooij, N. F. (2000). Characterisation of an electrodeposited conducting polymer FET array for vapour and odour sensing. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 35-42. IOP Publishing, Bristol.
- Cremoncini, A., Di Francesco, F., Lazzerini, B., Marcelloni, F., Martin, T., McCoy, S. A., Sensi, L. and Tselentis, G. (2000). Electronic noses using “intelligent” processing techniques. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 97-106. IOP Publishing, Bristol.
- Das, R. R., Shukla, K. K., Dwivedi, R. and Srivastava, A. R. (1999). Discrimination of individual gas/odor using responses of integrated thick film tin oxide sensor array and fuzzy-neuro concept. *Microelectronics Journal*, **30**, 793-800.
- Davide, F. A. M., Di Natale, C., D’Amico, A., Hierlemann, A., Mitrovics, J., Schweizer, M., Weimar, U. and Göpel, W. (1995a). Structure identification of non-linear models for QMB polymer-coated sensors. *Sens. Actuators B*. **24-25**, 830-842.
- Davide, F. A. M., Di Natale, C., D’Amico, A., Hierlemann, A., Mitrovics, J., Schweizer, M., Weimar, U., Göpel, W., Marco, M. and Pardo, A. (1995b). Dynamic calibration of QMB polymer-coated sensors by Wiener kernel estimation. *Sens. Actuators B*. **26-27**, 275-285.
- Di Natale, C., Marco, S., Davide, F. and D’Amico, A. (1995). Sensor-array calibration time reduction by dynamic modelling. *Sens. Actuators B*, **24-25**, 578-

583.

Di Natale, C., Macagnano, A., Paollesse, R., Mantini, A., Tarizzo, E., D'Amico, A., Sinesio, F., Bucarelli, F. M., Moneta, E. and Quaglia, G. B. (1998). Electronic nose and sensorial analysis: comparison of performances in selected cases. *Sens. Actuators B*, **50**, 246-252.

Distante, C., Artursson, T., Siciliano, P., Holmberg, M. and Lundström, I. (2000). Odour identification under drift effect. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 89-95. IOP Publishing, Bristol.

Dravieks, A. and Trotter, P. J. (1965). Polar vapour detection based on thermal modulation of contact potentials. *J. Sci. Instrum.*, **42**, 624-627.

Endres, H., Gottler, W., Jander, H. D., Drost, S., Sberveglieri, G., Faglia, G. and Perego, C. (1995). A systematic investigation on the use of time-dependent sensor signals in signal processing techniques. *Sens. Actuators B*, **24-25**, 785-789.

Esteves de Matos, R., Mason, D. J., Dow, C. S. and Gardner, J. W. (2000). Investigation of the growth characteristics of E. Coli using headspace analysis. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 181-188. IOP Publishing, Bristol.

Evans, N. D. and Chappell, M. J. (2000). Extensions to a procedure for generating locally identifiable reparameterisations of unidentifiable systems. *Math. Biosci.*, **168**, 137-159.

Furlong, C. and Stewart, J. R. (2000). Using a portable electronic nose for identification of odorous industrial chemicals. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 285-290. IOP Publishing, Bristol.

Gardner, J. W. (1989). Electrical conduction in solid-state gas sensors. *Sens. Actuators B*, **18**, 373-387.

Gardner, J. W. (1989). A diffusion-reaction model of electrical conduction in tin oxide gas sensors. *Semicond. Sci. Technol.*, **4**, 345-350.

Gardner, J. W. (1991). Detection of vapours and odours from a multisensor array using pattern recognition Part 1. Principal component and cluster analysis. *Sens.*

Actuators B, **4**, 109-115.

Gardner, J. W. and Bartlett, P. N. (eds.) (1991). *Sensors and Sensory Systems for an Electronic Nose, NATO ASI Series E: Applied Sciences*, Vol. 212, Kluwer, Dordrecht.

Gardner, J. W. and Bartlett, P. N. (1994). A brief history of electronic noses. *Sens. Actuators B*, **18-19**, 211-220.

Gardner, J. W. and Bartlett, P. N. (1999). *Electronic noses: principles and applications*. Oxford University Press, Oxford.

Gardner, J. W., Hines, E. L. and Wilkinson, M. (1990). The application of artificial neural networks in an electronic nose. *Meas. Sci. Technol.*, **1**, 446-451.

Gardner, J. W., Pearce, T. C, Friel, S., Bartlett, P. N. and Blair, N. (1994). A multisensor system for beer flavour monitoring using an array of conducting polymers and predictive classifiers. *Sens. Actuators B*, **18**, 240-243.

Gardner, J. W., Llobet, E. and Hines, E. L. (1999). PSPICE model for resistive gas and odour sensors. *IEE Proc.-Circuits Devices Syst.*, **146**, 101-104.

Gardner, J. W., Shin, H. W., Hines, E. L. and Dow, C. S. (2000). An electronic nose system for monitoring the quality of potable water. *Sens. Actuators B*, **69**, 336-341.

Garrigues, S., Talou, T. and Nesa, D. (2000) Quartz crystal microbalance sensors based electronic nose for QC in automotive industry. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 297-302. IOP Publishing, Bristol.

Godfrey, K. R. (1983). *Compartmental models and their application*. Academic, New York.

Godfrey, K. R. and Chapman, M. J. (1990). Identifiability and indistinguishability of linear compartmental models. *Mathematics and Computers in Simulation*, **32**, 273-295.

Godfrey, K. R., Jones, R. P., Brown, R.F. and Norton, J. P. (1982). Factors affecting the identifiability of compartmental models. *Automatica*, **18**, 285-293.

Grate, J. W., Abraham, M. H., McGill, R. A. (1995). Sorbent polymer coatings for

- chemical sensors. In *Polymer films in sensor applications* (ed. G. Harsanyi), pp. 136-149. Technomic, Lancaster, PA.
- Grate, J. W. and Frye, G. C. (1996). Acoustic Wave Sensors. In *Sensors Update Volume 2 Sensor Technology – Applications – Markets* (eds. H. Baltes, W. Göpel and J. Hesse.), pp. 37-83. VCH, Weinheim.
- Gutierrez-Osuna, R. (2000). Drift reduction for metal-oxide sensor arrays using canonical correlation regression and partial least squares. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 147-152. IOP Publishing, Bristol.
- Gutierrez-Osuna, R., Nagle, H. T. and Schiffman, S. S. (1999). Transient response analysis of an electronic nose using multi-exponential model. *Sens. Actuators B*, **61**, 170-182.
- Hall, H. S. and Knight, S. R. (1950). *Higher Algebra: A Sequel to Elementary Algebra for Schools*. Macmillan, London.
- Hartman, J. D. (1954), A possible objective method for the rapid estimation of flavours in vegetables. *Proc. Am. Soc. Hort. Sci.*, **64**, 335.
- Haugen, J., Tomic, O. and Kvaal, K. (2000). A calibration method for handling the temporal drift of solid state gas-sensors. *Anal. Chim. Acta*, **407**, 23-39.
- Haykin, S. (1999). *Neural Networks: A comprehensive foundation*. Prentice Hall, New Jersey.
- Heiland, G. (1982). Homogenous semiconducting gas sensors. *Sens. Actuators B*, **2**, 343-361.
- Hermann, R. and Krener, A. J. (1977). Nonlinear Controllability and Observability. *IEEE Trans. Autom. Control*, **AC-22**, 728-740.
- Hierlemann, A.m Weimar, U., Kraus, G., Schweizer-Berberich, M. and Göpel, W. (1995). Polymer-based sensor arrays and multicomponent analysis for the detection of hazardous organic vapours in the environment. *Sens. Actuators B*, **26-27**, 126-134.
- Hierold, C. and Muller, R. (1989). Quantitative analysis of gas mixtures with non-selective gas sensors. *Sens. Actuators B*, **17**, 587-192.

- Hines, E. L., Llobet, E. and Gardner, J. W. (1999). Electronic noses: a review of signal processing techniques. *IEE Proc.-Circuits Devices Syst.*, **146**, 297-310.
- Horrillo, M. C., Sayago, I., Fernández, M. J., Gómez-Espinosa, R., Blanco, A., Otero, L., García, M., Arés, L. and Gutiérrez, J. (2000). Design and development of an electronic nose system to control the processing of dry-cured Iberian hams monitored via internet. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 75-80. IOP Publishing, Bristol.
- Hwang, B. J., Yang, J-Y. and Lin, C-W. (1999). A microscopic gas-sensing model for ethanol sensors based on conductive polymer composites from polypyrrole and poly(ethylene oxide). *J. Electrochem. Soc.*, **146**, 1231-1236.
- Jones, E. (1991). The pellistor catalytic gas detector. In *Solid state gas sensors* (eds. P. T. Moseley and B. C. Tofield, pp. 17-31, Adam Hilger, Bristol.
- Kim, J. D., Byun, H. G. and Persaud, K. C. (2000). Application of a multiplayer perceptron based on the Levenberg-Marquardt algorithm to odour pattern classification and concentration estimation using odour sensing system. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 115-120. IOP Publishing, Bristol.
- Kohler, H., Röber, J., Link, N. and Bouzid, I. (1999). New applications of tin oxide gas sensors I. Molecular identification by cyclic variation of the working temperature and numerical analysis of the signals. *Sens. Actuators B*, **61**, 163-169.
- Kreyszig, E. (1988). *Advanced engineering mathematics (6th edition)*, John Wiley and Sons, Inc., New York.
- Krzanowski, W. J. (1988). *Principles of Multivariate Analysis*, Oxford University Press, Oxford.
- Lee, D. S., Huh, J. S., Byun, H. G., and Lee, D. D. (2000). An electronic nose for recognizing combustible gases using thick film sensor array and neural network. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 107-113. IOP Publishing, Bristol.
- Legin, A., Rudnitskaya, A., Seleznev, B. and Vlasov, Yu. (2000) Taste quantification using the electronic tongue. In *Electronic noses and olfaction 2000* (eds. J. W.

- Gardner and K. C. Persaud), pp. 13-22. IOP Publishing, Bristol.
- Ljung, L. (1987). *System Identification -Theory for the User*, Prentice-Hall, New Jersey.
- Ljung, L. (1995) *System Identification Toolbox User's Guide*, The MathWorks Inc., Mass.
- Llobet, E., Hines, E. L., Gardner, J. W., Bartlett, P. N. and Mottram, T. T. (1999a). Fuzzy ARTMAP based electronic nose data analysis. *Sens. Actuators B*, **61**, 183-190.
- Llobet, E., Hines, E. L., Gardner, J. W. and Franco, S. (1999b). Non-destructive banana ripeness determination using a neural network-based electronic nose. *Meas. Sci. Technol.*, **10**, 538-548.
- Llobet, E., Ionescu, R., Al-Khalifa, S., Brezmes, J., Vilanova, X., Correig, X., Bârsan, N. and Gardner, J. W. (2001a). Multicomponent gas mixture analysis using a single tin oxide sensor and dynamic pattern recognition. *IEEE Sensors Journal*, **1**, 207-213.
- Llobet, E., Brezmes, J., Ionescu, R., Vilanova, X., Al-Khalifa, S., Gardner, J. W., Barsan, N. and Corrieg, X. (2001b). Wavelet transform and fuzzy artmap based pattern recognition for fast gas identification using a micro-hotplate gas sensor. *Technical Digest of Transducers 01*, Munich, 10-14 June.
- Lundström, I., Hedborg, E., Spetz, A., Sundgren, H. and Winqvist, F. (1991). Electronic noses based on field effect structures. In *Sensors and Sensory Systems for an Electronic Nose, NATO ASI Series E: Applied Sciences, Vol. 212* (eds. Gardner, J. W. and Bartlett, P. N.) pp. 303-319, Kluwer, Dordrecht.
- Madou, M. J. and Morrison, S. R. (1989). *Chemical sensing with solid state devices*, Academic Press, San Diego.
- Marco, S., Pardo, A., Davide, F. A. M., Di Natale, C., D'Amico, A., Hierlemann, A., Mitrovics, J., Schweizer, M., Weimar, U. and Göpel, W. (1996). 'Different strategies for the identification of gas sensing systems', *Sens. Actuators B*, **34**, 213-223
- Marco, S., Ortega, A., Pardo, A. and Samitier, J. (1998). Gas identification with tin oxide sensor array and self-organizing maps: adaptive correction of sensor drifts.

- IEEE Trans. Instrum. Meas.*, **47**, 316-320,
- Mason, J. D. (1996). Classification of networks of weightless neurons. PhD thesis, University of Warwick, Coventry, UK.
- Moncrieff, R. W. (1961). An instrument for measuring and classifying odours. *J. Appl. Physiol.*, **16**, 742-749.
- Moriizumi, T. (1988). Langmuir-Blodgett films as chemical sensors. *Thin Solid Films*, **160**, 413-429.
- Mottram, T. T. and Persaud, K. C. (2000). Automatic milking: an experiment to inspect teats using an electronic nose. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 257-263. IOP Publishing, Bristol.
- Murphy, C. (1987). Olfactory psychophysics. In *Neurobiology of taste and smell* (eds. T. E. Finger and W. L. Silver), pp. 251-273. Wiley, New York.
- Nachnani, S. (2000). Oral malodour. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 3-23. IOP Publishing, Bristol.
- Nagle, H. T., Schiffman, S. S. and Gutierrez-Osuna, R. (1998). The how and why of electronic noses. *IEEE Spectrum*, September 1998, 22-34.
- Nakamura, M., Sugimoto, I., Kuwano, H. and Lemos, R. (1994). Chemical sensing by analyzing dynamics of plasma polymer film coated sensors. *Sens. Actuators B*, **20**, 231-237.
- Olafsson, R., Martinsdottir, E., Olafsdottir, G., Sigfusson, S. I. And Gardner, J. W. (1992). Monitoring of fish freshness using tin oxide sensors. In *Sensors and sensory systems for an electronic nose* (eds. J. W. Gardner and P. N. Bartlett), pp. 257-272. Kluwer, Dordrecht.
- Pardo, A., Marco, S. and Samitier, J. (1998). Nonlinear inverse dynamic models of gas sensing systems based on chemical sensors arrays for quantitative measurements. *IEEE Trans. Instrum. Meas.*, **47**, 644-651.
- Pardo, M., Benussi, G. P., Niederjaufner, G., Faglia, G. and Sberveglieri, G. (2000a). Detection of TNT vapours with the Pico-1 Nose. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 67-74. IOP Publishing, Bristol.

- Pardo, A., Marco, S., Calaza, C., Ortega, A., Perera, A., Sundic, T. and Samitier, J. (2000b). Methods for sensors selection in pattern recognition. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 83-95. IOP Publishing, Bristol.
- Persaud, K. C. and Pelosi, P. (1991). Sensor arrays using conducting polymers for an artificial nose. In *Sensors and Sensory Systems for an Electronic Nose, NATO ASI Series E: Applied Sciences, Vol. 212* (eds. Gardner, J. W. and Bartlett, P. N.) pp. 237-256, Kluwer, Dordrecht.
- Persaud, K. C. and Travers, P.J. (1996). Arrays of broad specificity films for sensing volatile chemicals. In *Handbook of biosensors and electronic noses: medicine, food and the environment* (ed. E. Kress-Rogers), pp. 563-592. CRC Press, Ohio.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T. and Flannery, B. P. (1993). *Numerical Recipes in C*. Cambridge University Press, Cambridge.
- Reich, S., Negri, R. M., Lamagna, A. and Dori, L. (2000). Identification of pollutant gases and its concentrations with a multisensorial arrangement. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 159-163. IOP Publishing, Bristol.
- Romain, N. J., Monticelli, A. C., Maternova, D. and André, Ph. (2000). Choice of a suitable e-nose output variable for the continuous monitoring of an odour in the environment. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 141-146. IOP Publishing, Bristol.
- Samitier, J., López-Villegas, J. M., Marco, S., Cámara, L., Pardo, A. and Ruiz, O. (1994). A new method to analyse signal transients in chemical sensors. *Sens. Actuators B*, **18-19**, 308-312.
- Schiffman, S. S., Wyrick, D. W., Gutierrez-Osuna, R. and Nagle, H. T. (2000). Effectiveness of an electronic nose for monitoring bacterial and fungal growth. In *Electronic noses and olfaction 2000* (eds. J. W. Gardner and K. C. Persaud), pp. 173-180. IOP Publishing, Bristol.
- Searle, G. E., Gardner, J. W., Chappell, M. J., Godfrey, K. R. and Chapman, M. J. (2002). System identification of electronic nose data from cyanobacteria experiments. *IEEE Sensors Journal*, **2**, 218-229.

- Shin, H. W., Llobet, E., Gardner J. W., Hines, E. L. and Dow, C. S. (2000). Classification of the strain and growth phase of cyanobacteria in potable water using an electronic nose system. *IEE Proc., Sci. Meas. Technol.*, **147**, 158-164.
- Slater, J. M., Esther, J. W., Freeman, N., May, I. P. and Weir, D. J. (1992). Gas and vapour detection with poly(pyrrole) gas sensors. *Analyst*, **117**, 1265-1270.
- Strässler, S. and Reis, A. (1983). Simple models for n-type metal oxide gas sensors. *Sens. Actuators*, **4**, 465-472.
- Sundgren, H., Lundström, I., Winqvist, F., Lukkari, I., Carlsson, R. and Wold, S. (1990). Evaluation of a multiple gas mixture with a simple MOSFET gas sensor array and pattern recognition. *Sens. Actuators B*, **2**, 115-123.
- Vajda, S., Godfrey, K. R. and Rabitz, H. (1989). Similarity transformation approach to structural identifiability of nonlinear models. *Math. Biosci.*, **93**, 217-248.
- Vilanova, X., Llobet, E., Alcubilla, R., Sueiras, J. and Correig, X. (1996). Analysis of the conductance transient in thick-film tin oxide gas sensors. *Sens. Actuators B*, **31**, 175-180.
- Wilkins, W. F. and Hartman, J. D. (1964). An electronic analogue for the olfactory process. *Ann. NY Acad. Sci.*, **116**, 608-612.
- Williams, D. E. (1991). Conduction and gas response of semiconductor gas sensors. In *Solid state gas sensors* (eds. P. T. Moseley and B. C. Tofield, pp. 17-31, Adam Hilger, Bristol.
- Wilson, D. M., DeWeerth, S. P. (1995). Odor discrimination using steady-state and transient characteristics of tin-oxide sensors. *Sens. Actuators B*, **28**, 123-128.
- Windischmann, H. and Mark, P. (1979). A model for the operation of a thin-film SnO_x conductance-modulation carbon monoxide sensor. *J. Electrochem. Soc.: Solid-state science and technology*, **126**, no. 4, 627-633.
- Zwaardemaker, H. and Hogewind, F. (1920). On spray-electricity and waterfall-electricity. *Proc. Acad. Sci. Amst.*, **22**, 429-437.

Appendix 1

Let C_h be the concentration of odour gases in the headspace, \dot{Q}_h be the flow rate of gas through the headspace, C_{hs} the saturated concentration of odour gases and k_e an evaporation rate constant. Then, the physical system relevant to the odour generation stage of the nose system can be described by the following first order ordinary differential equation:

$$\dot{C}_h(t) = -\dot{Q}_h C_h + k_e(C_{hs} - C_h). \quad (\text{A1.1})$$

Separating variables gives:

$$\int_{x=C_h(0)}^{C_h(t)} \frac{dx}{k_e C_{hs} - (\dot{Q}_h + k_e)x} = \int_{\tau=0}^t d\tau, \quad (\text{A1.2})$$

and so, after integrating,

$$\ln \left| \frac{k_e C_{hs} - (\dot{Q}_h + k_e)C_h(0)}{k_e C_{hs} - (\dot{Q}_h + k_e)C_h(t)} \right| = (\dot{Q}_h + k_e)t. \quad (\text{A1.3})$$

Taking exponentials and rearranging gives

$$(k_e C_{hs} - (\dot{Q}_h + k_e)C_h(0)) \exp(-(\dot{Q}_h + k_e)t) = k_e C_{hs} - (\dot{Q}_h + k_e)C_h(t). \quad (\text{A1.4})$$

Further rearrangement, and substitution of the initial condition that initially the headspace is saturated (i.e. $C_h(0) = C_{hs}$) gives the solution of the differential equation:

$$C_h(t) = C_{hs} \left\{ \frac{k_e}{\dot{Q}_h + k_e} + \left(1 - \frac{k_e}{\dot{Q}_h + k_e} \right) \exp(-(\dot{Q}_h + k_e)t) \right\}. \quad (\text{A1.5})$$

Appendix 2

Mathematica output

```

f1 = -k2 * (ceox - x2 - x20) * x1 - p * x1;

f2 = k1b * (ceox - x2 - x20) -
    k1f * (x2 + x20) * (x3 + x30) *
    (N1 - ceox + x2 + x20) +
    k2 * (ceox - x2 - x20) * x1;
f3 = k1b * (ceox - x2 - x20) -
    k1f * (x2 + x20) * (x3 + x30) *
    (N1 - ceox + x2 + x20) + rox - p * (x3 + x30);
g = {ran, 0, 0};
f = {f1, f2, f3};
dfdx =
  Simplify[
    {{D[f1, x1], D[f1, x2], D[f1, x3]},
     {D[f2, x1], D[f2, x2], D[f2, x3]},
     {D[f3, x1], D[f3, x2], D[f3, x3]}}];

phi1phi2 = (u1 - u2) * dfdx.g

dphi1phi2dx =
  Simplify[
    {{D[phi1phi2[[1]], x1], D[phi1phi2[[1]], x2],
      D[phi1phi2[[1]], x3]},
     {D[phi1phi2[[2]], x1], D[phi1phi2[[2]], x2],
      D[phi1phi2[[2]], x3]},
     {D[phi1phi2[[3]], x1], D[phi1phi2[[3]], x2],
      D[phi1phi2[[3]], x3]}}];

phi3 = f + u3 * g;
dphi3dx =
  Simplify[
    {{D[phi3[[1]], x1], D[phi3[[1]], x2],
      D[phi3[[1]], x3]},
     {D[phi3[[2]], x1], D[phi3[[2]], x2],
      D[phi3[[2]], x3]},
     {D[phi3[[3]], x1], D[phi3[[3]], x2],
      D[phi3[[3]], x3]}}];

phi3phi1phi2 = Simplify[dphi1phi2dx.phi3 - dphi3dx.phi1phi2]

phi1phi2 /. {x1 -> 0, x2 -> 0, x3 -> 0}

phi3phi1phi2 /. {x1 -> 0, x2 -> 0, x3 -> 0}

init =
  Simplify[
    phi3phi1phi2 /. {x1 -> 0, x2 -> 0, x3 -> 0,
      x20 ->
        (1/2) *
        (ceox - N1 - (k1b * p) / (k1f * rox) +
          Sqrt[
            (ceox - N1 - (k1b * p) / (k1f * rox))^2 +
            4 * (k1b * ceox * p) / (k1f * rox)]),
      x30 -> rox / p}]

```

```

h = 1 / (α + β * x2);

dh = {D[h, x1], D[h, x2], D[h, x3]}

Lφ1hx = Simplify[dh.f]

dLφ1hx =
Simplify[{D[Lφ1hx, x1], D[Lφ1hx, x2],
D[Lφ1hx, x3]}]

dLφ1hx0 =
Simplify[dLφ1hx /. {x1 → 0, x2 → 0, x3 → 0}]

φ2 = Simplify[f + u2 * g]

Lφ2φ1hx = Simplify[dLφ1hx.φ2]

$$\frac{1}{(\alpha + x_2 \beta)^3}$$


$$(\beta (-k_2 (ceox - x_2 - x_{20}) (ran u_2 + x_1 (-ceox k_2 - p + k_2 (x_2 + x_{20}))) (\alpha + x_2 \beta) + k_1 f (x_2 + x_{20}) (-ceox + N_1 + x_2 + x_{20}) (rox + k_1 b (ceox - x_2 - x_{20}) - p (x_3 + x_{30}) - k_1 f (x_2 + x_{20}) (-ceox + N_1 + x_2 + x_{20}) (x_3 + x_{30})) (\alpha + x_2 \beta) + (k_1 b (ceox - x_2 - x_{20}) + k_2 x_1 (ceox - x_2 - x_{20}) - k_1 f (x_2 + x_{20}) (-ceox + N_1 + x_2 + x_{20}) (x_3 + x_{30})) (k_1 b (\alpha + (2 ceox - x_2 - 2 x_{20}) \beta) + k_2 x_1 (\alpha + (2 ceox - x_2 - 2 x_{20}) \beta) + k_1 f (x_3 + x_{30}) (2 (x_2 + x_{20}) (\alpha - x_{20} \beta) + N_1 (\alpha - (x_2 + 2 x_{20}) \beta) + ceox (-\alpha + (x_2 + 2 x_{20}) \beta))))))$$

dLφ2φ1hx =
Simplify[{D[Lφ2φ1hx, x1], D[Lφ2φ1hx, x2],
D[Lφ2φ1hx, x3]}]

```

$$\left\{ \frac{1}{(\alpha + x_2 \beta)^3} (k_2 \beta ((\text{ceox} - x_2 - x_{20})) \right.$$

$$\begin{aligned}
& (\text{ceox} k_2 + p - k_2 (x_2 + x_{20})) (\alpha + x_2 \beta) + \\
& (k_{1b} (\text{ceox} - x_2 - x_{20}) + \\
& k_2 x_1 (\text{ceox} - x_2 - x_{20}) - k_{1f} (x_2 + x_{20}) \\
& (-\text{ceox} + N_1 + x_2 + x_{20}) (x_3 + x_{30})) \\
& (\alpha + (2 \text{ceox} - x_2 - 2 x_{20}) \beta) + \\
& (\text{ceox} - x_2 - x_{20}) \\
& (k_{1b} (\alpha + (2 \text{ceox} - x_2 - 2 x_{20}) \beta) + \\
& k_2 x_1 (\alpha + (2 \text{ceox} - x_2 - 2 x_{20}) \beta) + \\
& k_{1f} (x_3 + x_{30}) (2 (x_2 + x_{20}) (\alpha - x_{20} \beta) + \\
& N_1 (\alpha - (x_2 + 2 x_{20}) \beta) + \\
& \text{ceox} (-\alpha + (x_2 + 2 x_{20}) \beta))))),
\end{aligned}$$

$$\frac{1}{(\alpha + x_2 \beta)^4} (\beta (-3 \beta (-k_2 (\text{ceox} - x_2 - x_{20}))$$

$$\begin{aligned}
& (\text{ran} u_2 + x_1 (-\text{ceox} k_2 - p + k_2 \\
& (x_2 + x_{20}))) (\alpha + x_2 \beta) + \\
& k_{1f} (x_2 + x_{20}) (-\text{ceox} + N_1 + x_2 + x_{20}) \\
& (\text{rox} + k_{1b} (\text{ceox} - x_2 - x_{20}) - \\
& p (x_3 + x_{30}) - k_{1f} (x_2 + x_{20}) \\
& (-\text{ceox} + N_1 + x_2 + x_{20}) (x_3 + x_{30})) \\
& (\alpha + x_2 \beta) + (k_{1b} (\text{ceox} - x_2 - x_{20}) + \\
& k_2 x_1 (\text{ceox} - x_2 - x_{20}) - \\
& k_{1f} (x_2 + x_{20}) \\
& (-\text{ceox} + N_1 + x_2 + x_{20}) (x_3 + x_{30})) \\
& (k_{1b} (\alpha + (2 \text{ceox} - x_2 - 2 x_{20}) \beta) + \\
& k_2 x_1 (\alpha + (2 \text{ceox} - x_2 - 2 x_{20}) \beta) + \\
& k_{1f} (x_3 + x_{30}) (2 (x_2 + x_{20}) (\alpha - x_{20} \\
& \beta) + N_1 (\alpha - (x_2 + 2 x_{20}) \beta) + \\
& \text{ceox} (-\alpha + (x_2 + 2 x_{20}) \beta)))) + \\
& (\alpha + x_2 \beta) (-k_2 (\text{ceox} - x_2 - x_{20}) (\text{ran} u_2 + \\
& x_1 (-\text{ceox} k_2 - p + k_2 (x_2 + x_{20}))) \beta + \\
& k_{1f} (x_2 + x_{20}) (-\text{ceox} + N_1 + x_2 + x_{20}) \\
& (\text{rox} + k_{1b} (\text{ceox} - x_2 - x_{20}) - \\
& p (x_3 + x_{30}) - k_{1f} (x_2 + x_{20}) \\
& (-\text{ceox} + N_1 + x_2 + x_{20}) (x_3 + x_{30})) \\
& \beta - k_2^2 x_1 (\text{ceox} - x_2 - x_{20}) \\
& (\alpha + x_2 \beta) + k_2 (\text{ran} u_2 + \\
& x_1 (-\text{ceox} k_2 - p + k_2 (x_2 + x_{20}))) \\
& (\alpha + x_2 \beta) + k_{1f} (x_2 + x_{20}) \\
& (\text{rox} + k_{1b} (\text{ceox} - x_2 - x_{20}) - \\
& p (x_3 + x_{30}) - k_{1f} (x_2 + x_{20}) \\
& (-\text{ceox} + N_1 + x_2 + x_{20}) (x_3 + x_{30})) \\
& (\alpha + x_2 \beta) + k_{1f} (-\text{ceox} + N_1 + x_2 + x_{20}) \\
& (\text{rox} + k_{1b} (\text{ceox} - x_2 - x_{20}) - \\
& p (x_3 + x_{30}) - k_{1f} (x_2 + x_{20}) \\
& (-\text{ceox} + N_1 + x_2 + x_{20}) (x_3 + x_{30})) \\
& (\alpha + x_2 \beta) + k_{1f} (x_2 + x_{20}) \\
& (-\text{ceox} + N_1 + x_2 + x_{20}) \\
& (-k_{1b} + k_{1f} (\text{ceox} - N_1 - 2 (x_2 + x_{20})) \\
& (x_3 + x_{30})) (\alpha + x_2 \beta) + \\
& (k_{1b} (\text{ceox} - x_2 - x_{20}) + k_2 x_1 \\
& (\text{ceox} - x_2 - x_{20}) - k_{1f} (x_2 + x_{20}) \\
& (-\text{ceox} + N_1 + x_2 + x_{20}) (x_3 + x_{30})) \\
& (- (k_{1b} + k_2 x_1) \beta + k_{1f} (x_3 + x_{30}) \\
& (2 \alpha + (\text{ceox} - N_1 - 2 x_{20}) \beta)) + \\
& (-k_{1b} - k_2 x_1 + k_{1f} (\text{ceox} - N_1 - 2 \\
& (x_2 + x_{20})) (x_3 + x_{30})) \\
& (k_{1b} (\alpha + (2 \text{ceox} - x_2 - 2 x_{20}) \beta) + \\
& k_2 x_1 (\alpha + (2 \text{ceox} - x_2 - 2 x_{20}) \beta) + \\
& k_{1f} (x_3 + x_{30}) (2 (x_2 + x_{20}) (\alpha - x_{20} \\
& \beta) + N_1 (\alpha - (x_2 + 2 x_{20}) \beta) + \\
& \text{ceox} (-\alpha + (x_2 + 2 x_{20}) \beta))))),
\end{aligned}$$

$$\begin{aligned}
& \mathbf{dL}\phi 2\phi 1\mathbf{x0} = \\
& \mathbf{Simplify}[\mathbf{dL}\phi 2\phi 1\mathbf{x} /. \{\mathbf{x1} \rightarrow \mathbf{0}, \mathbf{x2} \rightarrow \mathbf{0}, \mathbf{x3} \rightarrow \mathbf{0}\}] \\
& \left\{ \frac{1}{\alpha^3} (k2 \beta ((ceox - x20) (ceox k2 + p - k2 x20) \alpha + \right. \\
& \quad (ceox (k1b + k1f x20 x30) - \\
& \quad \quad x20 (k1b + k1f (N1 + x20) x30)) \\
& \quad (\alpha + 2 (ceox - x20) \beta) + (ceox - x20) \\
& \quad (k1b (\alpha + 2 (ceox - x20) \beta) + k1f x30 \\
& \quad \quad (-ceox \alpha + N1 \alpha + 2 x20 \alpha + 2 ceox x20 \beta - \\
& \quad \quad \quad 2 N1 x20 \beta - 2 x20^2 \beta))) \right\}, \frac{1}{\alpha^4} \\
& (\beta (\alpha (k2 \text{ran u2} \alpha + k1f x20 (-ceox + N1 + x20) \\
& \quad (-k1b + k1f (ceox - N1 - 2 x20) x30) \alpha + \\
& \quad k1f x20 (\text{rox} + k1b (ceox - x20) - p x30 - \\
& \quad \quad k1f x20 (-ceox + N1 + x20) x30) \alpha + \\
& \quad k1f (-ceox + N1 + x20) \\
& \quad (\text{rox} + k1b (ceox - x20) - p x30 - \\
& \quad \quad k1f x20 (-ceox + N1 + x20) x30) \alpha - \\
& \quad k2 \text{ran u2} (ceox - x20) \beta + \\
& \quad k1f x20 (-ceox + N1 + x20) \\
& \quad (\text{rox} + k1b (ceox - x20) - p x30 - \\
& \quad \quad k1f x20 (-ceox + N1 + x20) x30) \beta + \\
& \quad (ceox (k1b + k1f x20 x30) - \\
& \quad \quad x20 (k1b + k1f (N1 + x20) x30)) \\
& \quad (-k1b \beta + k1f x30 \\
& \quad \quad (2 \alpha + (ceox - N1 - 2 x20) \beta)) + \\
& \quad (-k1b + k1f (ceox - N1 - 2 x20) x30) \\
& \quad (k1b (\alpha + 2 (ceox - x20) \beta) + k1f x30 \\
& \quad \quad (-ceox \alpha + N1 \alpha + 2 x20 \alpha + 2 ceox \\
& \quad \quad \quad x20 \beta - 2 N1 x20 \beta - 2 x20^2 \beta))) - \\
& \quad 3 \beta (-k2 \text{ran u2} (ceox - x20) \alpha + \\
& \quad k1f x20 (-ceox + N1 + x20) \\
& \quad (\text{rox} + k1b (ceox - x20) - p x30 - \\
& \quad \quad k1f x20 (-ceox + N1 + x20) x30) \alpha + \\
& \quad (ceox (k1b + k1f x20 x30) - \\
& \quad \quad x20 (k1b + k1f (N1 + x20) x30)) \\
& \quad (k1b (\alpha + 2 (ceox - x20) \beta) + \\
& \quad \quad k1f x30 (-ceox \alpha + N1 \alpha + 2 \\
& \quad \quad \quad x20 \alpha + 2 ceox x20 \beta - 2 \\
& \quad \quad \quad N1 x20 \beta - 2 x20^2 \beta)))) \right\}, \\
& \frac{1}{\alpha^3} (k1f \beta (x20 (-ceox + N1 + x20) \\
& \quad (-p - k1f x20 (-ceox + N1 + x20)) \alpha + \\
& \quad (ceox (k1b + k1f x20 x30) - \\
& \quad \quad x20 (k1b + k1f (N1 + x20) x30)) \\
& \quad (-ceox (\alpha - 2 x20 \beta) + N1 (\alpha - 2 x20 \beta) + \\
& \quad \quad 2 x20 (\alpha - x20 \beta)) - \\
& \quad x20 (-ceox + N1 + x20) \\
& \quad (k1b (\alpha + 2 (ceox - x20) \beta) + \\
& \quad \quad k1f x30 (-ceox \alpha + N1 \alpha + 2 x20 \alpha + 2 ceox \\
& \quad \quad \quad x20 \beta - 2 N1 x20 \beta - 2 x20^2 \beta)))) \}
\end{aligned}$$

$$\begin{aligned}
& \mathbf{ratiol} = \\
& \mathbf{Simplify}[\mathbf{dL}\phi 2\phi 1\mathbf{x0}[[1]] / \mathbf{dL}\phi 2\phi 1\mathbf{x0}[[3]]]
\end{aligned}$$

```
ratio11 =  
Simplify[  
  ratio1 /.  
    {x20 →  
      (1/2) *  
        (ceox - N1 - (klb * p) / (klf * rox) +  
          Sqrt[  
            (ceox - N1 - (klb * p) / (klf * rox))^2 +  
              4 * (klb * ceox * p) / (klf * rox)]),  
      x30 → rox / p}]
```

$$\begin{aligned}
& - \left(k_2 \text{ rox}^2 \left(k_1 b^2 p^2 (k_2 p - k_1 f \text{ rox}) + \right. \right. \\
& \quad k_1 b k_1 f p \text{ rox} \left(2 \text{ ceox} k_2 p + 2 k_2 N_1 p + \right. \\
& \quad \left. p^2 - 2 \text{ ceox} k_1 f \text{ rox} - 2 k_1 f N_1 \text{ rox} - \right. \\
& \quad \left. k_2 p \sqrt{\left(\frac{1}{k_1 f^2 \text{ rox}^2} (k_1 b^2 p^2 + 2 k_1 b \right. \right. \\
& \quad \left. \left. k_1 f (\text{ceox} + N_1) p \text{ rox} + k_1 f^2 \right. \right. \\
& \quad \left. \left. (\text{ceox} - N_1)^2 \text{ rox}^2 \right)} \right) + \\
& \quad k_1 f \text{ rox} \sqrt{\left(\frac{1}{k_1 f^2 \text{ rox}^2} (k_1 b^2 p^2 + \right. \\
& \quad \left. 2 k_1 b k_1 f (\text{ceox} + N_1) p \text{ rox} + \right. \\
& \quad \left. k_1 f^2 (\text{ceox} - N_1)^2 \text{ rox}^2 \right)} \left. \right) + \\
& \quad k_1 f^2 \text{ rox}^2 \left(\text{ceox}^2 (k_2 p - k_1 f \text{ rox}) + \right. \\
& \quad \left. (-k_2 N_1 p - p^2 + k_1 f N_1 \text{ rox}) \right. \\
& \quad \left. \left(-N_1 + \sqrt{\left(\frac{1}{k_1 f^2 \text{ rox}^2} (k_1 b^2 p^2 + \right. \right. \right. \\
& \quad \left. \left. 2 k_1 b k_1 f (\text{ceox} + N_1) p \text{ rox} + \right. \right. \\
& \quad \left. \left. k_1 f^2 (\text{ceox} - N_1)^2 \text{ rox}^2 \right)} \right) \left. \right) + \\
& \quad \text{ceox} \left(p^2 - k_2 p \sqrt{\left(\frac{1}{k_1 f^2 \text{ rox}^2} (k_1 b^2 p^2 + \right. \right. \\
& \quad \left. \left. 2 k_1 b k_1 f (\text{ceox} + N_1) p \text{ rox} + \right. \right. \\
& \quad \left. \left. k_1 f^2 (\text{ceox} - N_1)^2 \text{ rox}^2 \right)} \right) + \\
& \quad \left. k_1 f \text{ rox} \left(2 N_1 + \sqrt{\left(\frac{1}{k_1 f^2 \text{ rox}^2} (k_1 b^2 \right. \right. \right. \\
& \quad \left. \left. p^2 + 2 k_1 b k_1 f (\text{ceox} + N_1) \right. \right. \\
& \quad \left. \left. p \text{ rox} + k_1 f^2 (\text{ceox} - N_1)^2 \right. \right. \\
& \quad \left. \left. \text{rox}^2 \right)} \right) \left. \right) \left. \right) \left. \right) \left. \right) / \\
& \left(k_1 b p \left(k_1 b^3 p^4 - k_1 b^2 k_1 f p^2 \text{ rox} \right. \right. \\
& \quad \left. \left(-2 \text{ ceox} p - 2 N_1 p + \text{ rox} + \right. \right. \\
& \quad \left. p \sqrt{\left(\frac{1}{k_1 f^2 \text{ rox}^2} (k_1 b^2 p^2 + 2 k_1 b \right. \right. \\
& \quad \left. \left. k_1 f (\text{ceox} + N_1) p \text{ rox} + k_1 f^2 \right. \right. \\
& \quad \left. \left. (\text{ceox} - N_1)^2 \text{ rox}^2 \right)} \right) \left. \right) + \\
& \quad k_1 b k_1 f p \text{ rox}^2 \left(\text{ceox}^2 k_1 f p + k_1 f N_1^2 p + \right. \\
& \quad \left. p^2 - 2 k_1 f N_1 \text{ rox} - \right. \\
& \quad \left. k_1 f N_1 p \sqrt{\left(\frac{1}{k_1 f^2 \text{ rox}^2} (k_1 b^2 p^2 + 2 \right. \right. \\
& \quad \left. \left. k_1 b k_1 f (\text{ceox} + N_1) p \text{ rox} + \right. \right. \\
& \quad \left. \left. k_1 f^2 (\text{ceox} - N_1)^2 \text{ rox}^2 \right)} \right) + \\
& \quad \left. k_1 f \text{ rox} \sqrt{\left(\frac{1}{k_1 f^2 \text{ rox}^2} (k_1 b^2 p^2 + 2 \right. \right. \\
& \quad \left. \left. k_1 b k_1 f (\text{ceox} + N_1) p \text{ rox} + k_1 f^2 \right. \right. \\
& \quad \left. \left. (\text{ceox} - N_1)^2 \text{ rox}^2 \right)} \right) - \text{ceox} k_1 f \\
& \quad \left. \left(2 \text{ rox} + p \right) \sqrt{\left(\frac{1}{k_1 f^2 \text{ rox}^2} (k_1 b^2 p^2 + \right. \right.
\end{aligned}$$


```
ratio2 = Simplify[dLphi[x0][[1]] / dLphi[x0][[3]]]
```

```
ratio21 =  
Simplify[  
ratio2 /.  
{x20 ->  
  (1/2) *  
  (ceox - N1 - (k1b * p) / (k1f * rox) +  
  Sqrt[  
    (ceox - N1 - (k1b * p) / (k1f * rox))^2 +  
    4 * (k1b * ceox * p) / (k1f * rox)]),  
x30 -> rox / p}]  
k2 rox  
-----  
k1b p
```

Converted by [Mathematica](#) September 15, 2002