Meas. Sci. Technol. **9** (1998) 120–127. Printed in the UK PII: S0957-0233(98)86520-X

The prediction of bacteria type and culture growth phase by an electronic nose with a multi-layer perceptron network

J W Gardner†, M Craven†, C Dow‡ and E L Hines†

† Department of Engineering, Warwick University, Coventry CV4 7AL, UK
‡ Department of Biological Sciences, Warwick University, Coventry CV4 7AL, UK

Received 16 August 1997, accepted for publication 24 September 1997

Abstract. An investigation into the use of an electronic nose to predict the class and growth phase of two potentially pathogenic micro-organisms, *Eschericha coli* (*E. coli*) and *Staphylococcus aureus* (*S. aureus*), has been performed. In order to do this we have developed an automated system to sample, with a high degree of reproducibility, the head space of bacterial cultures grown in a standard nutrient medium. Head spaces have been examined by using an array of six different metal oxide semiconducting gas sensors and classified by a multi-layer perceptron (MLP) with a back-propagation (BP) learning algorithm. The performance of 36 different pre-processing algorithms has been studied on the basis of nine different sensor parameters and four different normalization techniques. The best MLP was found to classify successfully 100% of the unknown *S. aureus* samples and 92% of the unknown *E. coli* samples, on the basis of a set of 360 training vectors and 360 test vectors taken from the lag, log and stationary growth phases. The real growth phase of the bacteria was determined from optical cell counts and was predicted from the head space samples with an accuracy of 81%. We conclude that these results show considerable promise in that the correct prediction of the type and growth phase of pathogenic bacteria may help both in the more rapid treatment of bacterial infections and in the more efficient testing of new anti-biotic drugs.

1. Introduction

The human body is constantly in contact with a variety of micro-organisms, some of which are bacteria. Normal microbial flora reside on the surface of the skin and internally from a few days to weeks (transient flora) or for years (resident flora) without normally causing a disease. Bacterial infections are caused by the growth of certain micro-organisms known as pathogens and the most common infectious diseases tend to afflict our upper respiratory system. Such diseases may be caused by a micro-organism that was originally part of our normal flora, such as *Staphylococcus aureus*, or a micro-organism that is alien and parasitic, such as *Legionella pneumophilia*. Clearly, there is a clinical need to identify the nature of the pathogenic micro-organisms within a bacterial infection as quickly as possible and then to apply the appropriate anti-biotic. Current practice involves the wiping of the infected area (for example throat or ears) with a sterile cotton swab which is then sent to a laboratory for analysis. The swab is then usually seeded onto a blood agar plate which acts as a growth medium, which is incubated at 37 °C, visually examined, stained and examined under a microscope. This laborious procedure can take many days, during which the infected patient may be receiving the wrong or no treatment and thus the disease may progress unchecked. An earlier identification of the pathogen would not only produce a more rapid treatment of the infection but also reduce the misuse of antibiotics. The latter is, unfortunately, more serious since the overuse of anti-biotics is relatively common and leads to the early appearance of resistant strains of bacteria.

Electronic nose instrumentation has advanced rapidly during the past 10 years, the majority of applications being within the foods and drinks industry (Gardner and Bartlett 1992, Kress-Rogers 1997). Indeed, reports have been written on the ability of an electronic nose to detect the freshness of foods such as fish (Schweizer-Berberich *et al* 1994) and meat (Winquist *et al* 1993). A deterioration in food freshness is often associated with microbial spoilage and this work suggests that there is potential for the clinical application of electronic noses. We believe that the first attempt to identify micro-organisms with an electronic nose was made by Craven *et al* (1994). An array of four different commercial metal oxide gas sensors (Figaro Engineering, Japan) was used to sample

0957-0233/98/010120+08$19.50

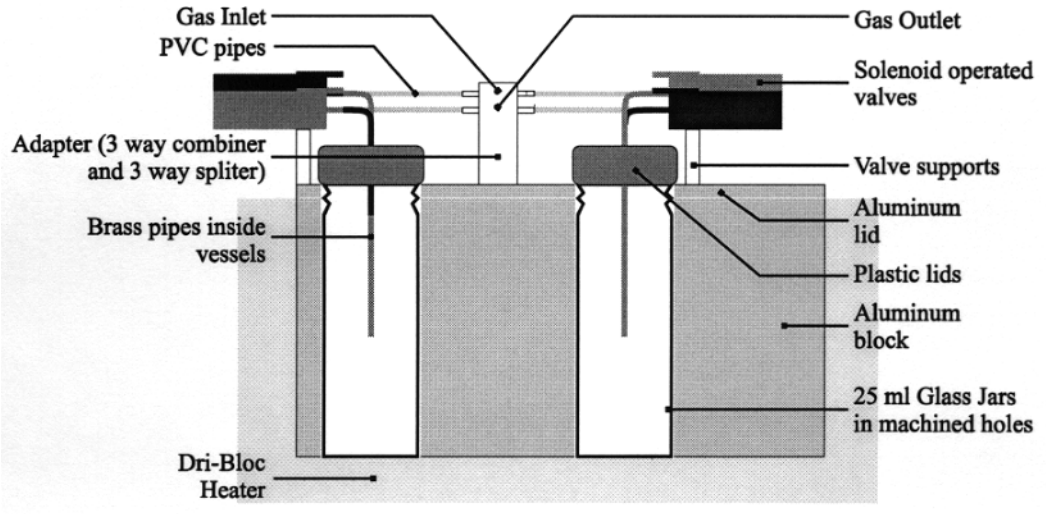

Figure 1. A cross section showing two of the three glass sample vials within the head-space sampler for bacterial specimens.

the head space of six pathogenic bacteria grown in blood agar; namely *Clostridia perfringens*, *Proteus*, *Haemophilus influenzae*, *Bacteroides fragilis*, *Oxford staphylococcus* and *Pseudomonas aeruginosa*. A relatively small set of 49 response vectors was obtained with six being assigned as training vectors and one as a test vector per micro-organism (and blood agar). A multi-layer perceptron neural network with a topology of $4 \times 4 \times 7$ was trained on the data using a back-propagation algorithm. In this first test the four-element electronic nose was able to classify correctly only 62% of the pathogens. Although the leave-one-out method of cross validation is an almost unbiased estimator of the true error rate, its variance is often high when it is applied to a small sample set such as that reported here. A bootstrap estimator, such as e0, uses replacement and may be preferable for small samples, but, ideally, larger sample sets should be generated. More recently, Gibson *et al* (1996) reported on the use of an array of 16 conducting polymer resistive gas sensors to detect 12 different bacteria (and yeast), again from cultures grown on agar plates. The number of replicate samples varied from six to 48 samples, resulting in 244 input vectors. Their decision to define seven parameters per sensor led to a network architecture of $112 \times 90 \times 13$. The network was then trained using a back-propagation algorithm on the entire data-set. Classification rates varied from 63.6% to 100% per micro-organism which, at first sight, appears encouraging, but finding over 10 000 adaptable weights of more than 100 active neurones, with only 244 training vectors and not using cross validation, makes the accuracy debatable.

In a clinical application it is desirable both to identify the micro-organism as rapidly as possible and to automate the process. Current identification methods based upon ribosomal RNA sequencing have revolutionized microbial taxonomy and phylogeny; however, they simply show the presence of the nucleic acid sequence rather than an intact, or indeed *viable* microbial cell. So the use of nucleic acid probes for identification purposes is valuable but does not describe the physiological *activity* of the microbe. Here we report on an instrument designed to sample and analyse automatically the head space of a micro-organism at different stages of its growth cycle, namely, lag, log, stationary and death. The greatest benefit will be obtained if an electronic nose can identify its lag or log phase and so we report on an instrument developed to recognize *both* the type of bacterium and its growth phase during its early life. We have chosen to study two model micro-organisms, *Eschericha coli* and *Staphylococcus aureus*, as examples of gram positive and gram negative pathogens. These microbial organisms have a short incubation period of about 24 h at body temperature and are often causes of ear infection and sinusitis.

2. The experimental procedure

2.1. The automated head space sampler

An autosampler and sampling methodology have been developed for the measurement of the head space of the bacterial specimens. First, 20 ml of a multipurpose nutrient broth was placed into three sterile 25 ml glass vials, one representing a reference sample; secondly, a small number of cells (0.1 ml of inoculum) was introduced into the other two vials from master cultures. These vials were connected to the autosampler via 2 mm internal diameter tubing and maintained at $37 \pm 0.1\,^{\circ}\mathrm{C}$ by a Dri-Bloc^{TM} heater, see figure 1.

The head space was sampled using air that had first been drawn through a charcoal filter to remove organic pollutants, then bubbled through water to equilibrate the water vapour pressure and finally passed through a bio-filter to remove air-borne micro-organisms, see figure 2. The apparatus was cleaned using an ethanol–water solution and 5% sodium hypochlorite solution between experimental runs. The sampling procedure was controlled via a set of miniature solenoid valves (Lee Products Ltd) interfaced to the digital output lines from a data acquisition card (National Instruments) operated under a virtual instrument (VI) written in Labview (National Instruments). Full experimental details have been published by Craven (1997).

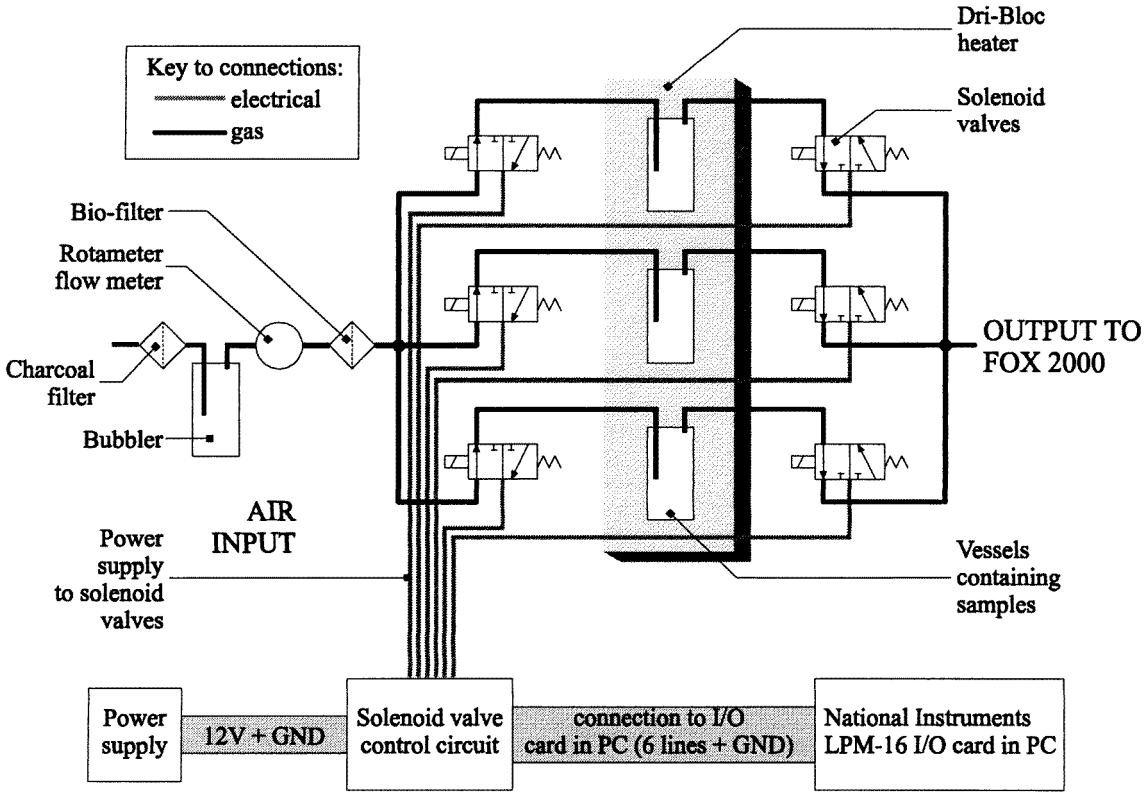

Figure 2. A schematic diagram of the automated head-space sampler showing the pipework and electrical wiring.

2.2. The electronic nose instrument

The sample from the head space was passed into a rectangular sensor chamber with a dead volume of about 100 ml which contained six commercial metal oxide odour sensors, a temperature i.c. sensor (LM35CZ) and a humidity sensor (Mini-cap, Panametrics) to measure the properties of the incoming gas, figure 3.

Although the temperature of the head-space gas entering the sensor chamber was controlled to be $37 \pm 0.1\,^{\circ}C$, there was no direct control of the temperature of the sensor chamber itself and it was found to vary by several degrees with fluctuations in the ambient room temperature. The six metal oxide odour sensors (AlphaMOS, France) were chosen on the basis of the knowledge that, as the micro-organisms grow, they convert the primary metabolites in the nutrient (carbohydrates, sugars, minerals and so on) into secondary metabolites (hydrocarbons, alcohols, aldehydes, acids, ammonia and so on). The target compounds for the six sensors were as follows: hydrocarbons (P.10.1), alcohols (T.70.2), aldehydes/heteroatoms (P.40.1), polar molecules (T.30.1, P.A.2) and non-polar compounds (P.10.2). The electrical signals from the six odour sensors were conditioned and then fed into a 12-bit ADC and recorded by a virtual instrument. The functioning of the head space sampler was reproducible to within better than 1%. A subsequent re-design of the sensor chamber to provide an axisymmetrical structure with closed-loop temperature control has reduced the temperature variation to less than $0.05\,^{\circ}C$, the dead volume to only 3.1 ml and the mixing time to less than 1 s at a flow rate of 200 ml min^{-1}. Further details of the design and performance of the revised rapid sampling system have been published elsewhere (Craven and Gardner 1997).

2.3. Cultures sampled and cell counting

A series of measurements was performed on cultures of two bacteria: *Eschericha coli* and *Staphylococcus aureus*. There were two separate experimental runs on each bacterium, making a total set of four experiments. Each experiment ran for a period of 12 h with a sampling cycle lasting 8 min: reference sample (2 min), bacterial sample A (2 min), reference sample (2 min) and bacterial sample B (2 min). Thus each experiment was comprised of 180 sample vectors making a total of 720 vectors for the entire set of experiments.

The activity of the micro-organism within the nutrient was monitored by drawing a small volume (0.1 ml) into a syringe at regular 1 h intervals and performing a viable cell count using an optical densitometer. Figure 4 shows a typical plot of the viable cell count (in terms of colony forming units) against the elapsed time for *E. coli*. The first three phases of cell growth are clearly visible, the lag phase lasting about 1 h, followed by rapid cell growth during the log (or exponential) phase lasting about 2.5 h and finally a long stationary (or static) phase lasting about 8 h. The experiment was stopped before the death phase was entered because this was of little practical interest to us and can be many days in duration.

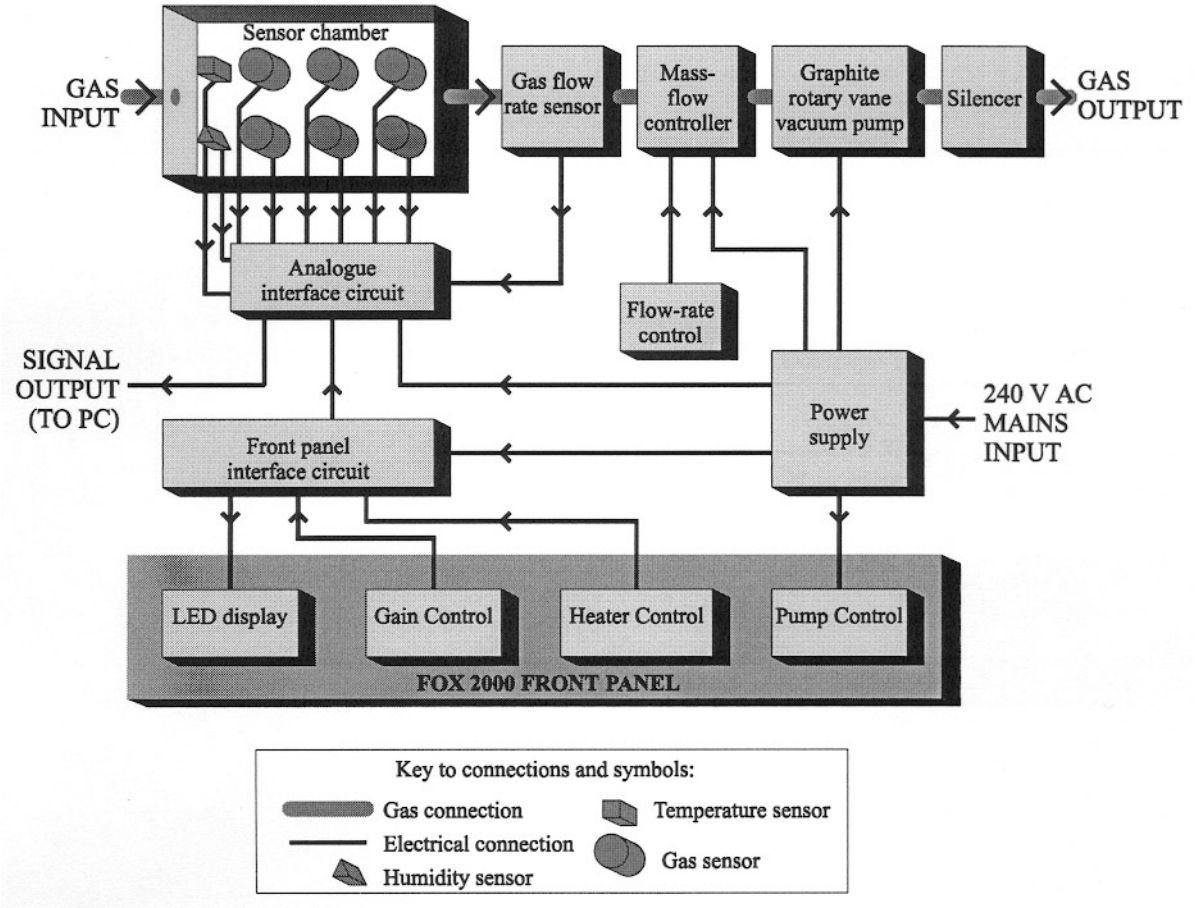

Figure 3. The design of the electronic nose system with the sensor chamber housing six commercial metal oxide odour sensors, a temperature i.c. and a capacitive humidity sensor.

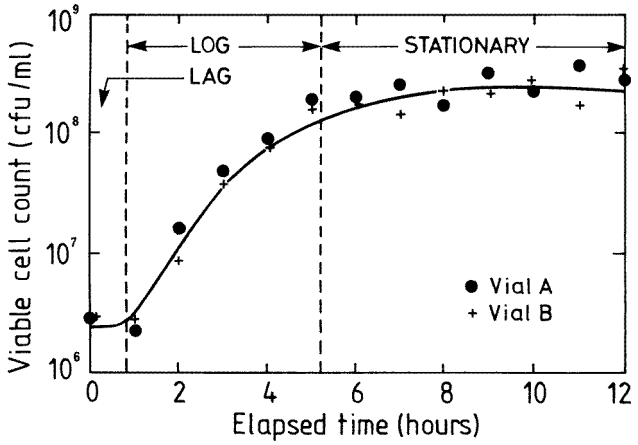

Figure 4. A typical plot showing the number of colony forming units (cfus) in 0.1 ml of inoculum for *E. coli* over a 12 h period. The different phases of growth are indicated.

3. Data processing techniques

3.1. Pre-processing algorithms

The choice of the parameter used to define the response of an odour sensor has been shown to influence strongly the performance both of chemometric (Gardner 1991) and of neural network (Gardner *et al* 1992, Endres *et al* 1995) data processing techniques. Therefore we have evaluated the performance of a number of different pre-processing algorithms—referred to here as sensor (feature) models. Table 1 lists nine different sensor models x_i together with the equation used to define them and their abbreviation, where V is the output voltage of the odour sensor i and is directly proportional to its electrical resistance.

In addition, it can be useful to normalize the sensor response data and so we have also studied the effect of sensor normalization (denoted by the letter s), vector array normalization (v) and autoscaling (a). The three normalized parameters y_i are also defined in table 1. Sensor normalization sets the sensor values so that they all lie in the range of $[0, +1]$ whereas autoscaling sets the mean value to 0 and the variance to 1. Both techniques give equal weighting to each sensor and thus compensate for differences in the magnitudes of the signals. Array normalization divides each sensor value by the norm of the array vector. This, in effect, removes the concentration dependence of the magnitude of the sensor output which is useful when the intensity of the odour is irrelevant, namely when the same micro-organism is sampled but there are more cells present!

3.2. The pattern analysis

A back-propagation neural network was first applied to the output of an electronic nose in 1990 (Gardner *et al*

Table 1. Definitions of sensor models and normalization methods.

Model/method	Equation	Abbreviation
Difference	$x_i = (V_{ref}^{max} - V_{odour}^{min})$	df
Relative	$x_i = \frac{V_{odour}^{min}}{V_{ref}^{max}}$	rl
Fractional difference	$x_i = \frac{V_{ref}^{max} - V_{odour}^{min}}{V_{ref}^{max}}$	fd
Absolute final output	V_{odour}^{final} and V_{ref}^{final}	af
Minimum output	V_{odour}^{min} and V_{ref}^{min}	mn
Final relative	$x_i = \frac{V_{odour}^{final}}{V_{ref}^{final}}$	fr
Modified difference	$x_i = (V_{odour}^{max} - V_{odour}^{min}) - (V_{ref}^{max} - V_{ref}^{min})$	md
Modified fractional difference	$x_i = \frac{V_{ref}^{max} - V_{odour}^{min}}{V_{ref}^{max} - V_{ref}^{min}}$	mf
Final fractional difference	$x_i = \frac{V_{odour}^{final} - V_{ref}^{final}}{V_{ref}^{final}}$	ff
No normalization		n
Sensor normalization	$y_i = \frac{x_i - x_i^{min}}{x_i^{max} - x_i^{min}}$	s
Vector array normalization	$y_i = \frac{x_i}{(x_1^2 + x_2^2 + \ldots + x_n^2)^{1/2}}$	v
Autoscaling	$y_i = \frac{x_i - \bar{x}_i}{\sigma_i}$	a

Table 2. The effect of the pre-processing algorithm (sensor model) on the ability of a back-propagation neural network to predict a bacterial class. The performance is defined by the network sum of squares error (SSE) and the percentage of test vectors assigned to each class (abbreviations defined in table 1). Confidence levels are indicated by the standard deviation, and models ranked in order.

	SSE		Correct (%)		Incorrect (%)		Unknown (%)	
Model	Mean	σ	Mean	σ	Mean	σ	Mean	σ
df	346.94	128.53	77.22	9.19	10.28	4.30	12.50	6.73
md	685.88	250.56	61.35	13.93	21.08	9.04	17.57	12.34
af	388.86	164.16	61.11	14.36	7.22	6.42	31.67	13.28
mn	417.96	177.99	60.14	24.86	9.86	6.90	30.00	24.67
ff	595.86	264.34	44.69	32.56	10.24	7.71	45.07	34.81
fd	543.13	125.22	43.54	22.97	10.17	4.07	46.28	24.75
rl	547.24	124.13	40.00	25.68	8.99	7.07	51.01	30.37
fr	558.73	223.16	36.01	39.92	6.74	8.95	57.26	45.69
mf	705.44	188.79	23.89	26.69	8.54	4.98	67.57	29.37

1990) and since then it has not only been reported widely by other researchers but also offered as a pattern analysis technique by most of the manufacturers of commercial instruments. In our experiments we have employed the commonly used multi-layer perceptron (MLP) trained by the back-propagation algorithm. The MLP analyses were carried out using a software package called SNNS (Stuttgart Neural Network Simulator), version 4.1. The MLPs were designed using a graphical interface, then trained and tested on data using an internal processing language (which also allowed automated systematic training and testing of MLPs). The performance of the MLPs was then analysed using built-in analysis software.

4. Results

4.1. The classification of bacteria types

The combined data sets from experiments 1 and 3 on *E. coli* and *S. aureus* were first used to train a back-propagation network with momentum. The performances of the networks were then evaluated using split-sample validation from the data gathered from experiments 2 and 4. Clearly this is a much more challenging task than using one experimental run *both* to train and to test the network because it includes culture-to-culture variation. The nine sensor parameters, x_i, and three normalization methods, y_i, listed in table 1 were evaluated together with no normalization and so there were 36 different training sets (9×4) and 36 different test sets; each one containing 360 vectors (180 samples per micro-organism). In the first instance, the target vectors were defined to have a value of either 1 or 0 because the class membership of *E. coli* or *S. aureus* was known *a priori* and the age of the micro-organism was ignored.

A network was chosen with a single hidden layer because this reduces the number of processing neurones and thus the number of training vectors required to obtain a reliable answer. This choice was based on previous work in which it had been shown that the use of two or more hidden layers has a marginal effect on the network performance (Holmberg *et al* 1995). Our default network topology was

$6 \times 20 \times 2$ (with $12 \times 20 \times 2$ for the absolute and minimum models, af and mn, respectively). The output layer was designed to follow the standard encoding of unordered categories (that is, classes) which is one output neurone per micro-organism. A tanh function was used for the transfer function of the processing elements in the hidden and output layers because some inputs have negative values (for example in the autoscaled data).

We decided to have ten nodes per bacteria class in the hidden layer, namely 20 hidden neurones in the network used to analyse the *E. coli* and *S. aureus* samples. The large number of neurones in the hidden layer was considered necessary because an early-stopping training technique was employed. Early stopping is a technique commonly used to improve the ability of a network to generalize. (Generalization is the ability of the ANN to classify correctly input vectors which have not been used for training.) The number of hidden neurones has a significant influence on the ability of an MLP to generalize (Bishop 1995). To date no reliable method for determining the optimum number of hidden neurones has been discovered, although some attempts have been made (Hines *et al* 1993, Fekadu *et al* 1993).

Training of the output of the neural network was carried out using the '402040' technique (unlike Gibson *et al* (1996) who employed the '5050' approach). This superior technique defines the lower target output band as occurring in the range $[-1, -0.5]$ and the upper target output band was defined as occurring in the range $[+0.5, +1]$. An input vector was deemed *correctly* classified when exactly one output was in the upper band, all the other outputs (in this case just one) were in the lower band, and the highest target output corresponded to the highest output. Conversely, a vector was *incorrectly* classified when exactly one output was in the upper band, all other outputs were in the lower band and the highest target output did not correspond to the highest output. An input vector was classified as unknown when it was neither correctly classified nor incorrectly classified. Classifying output vectors in this way allowed for unknown feature vectors to be more readily detected. A commonly adopted rule of 'winner-takes-all' forces all vectors to be classified and thus artificially increases the number of true and false classifications. (Although the winner-takes-all method may give a more impressive figure for the number of correct classifications, it does not necessarily increase the difference between the numbers of correct classifications and incorrect classifications, which is a more important measure.) The network training parameters were set throughout to values of learning rate $\eta = 0.001$, momentum coefficient $\alpha = 5.0$, flat-spot elimination constant $c = 0.1$ and maximum tolerance of error per output $d_{max} = 0.1$. The relatively low value of η and high value of α helped to reduce the effect of local minima, which is necessary for a MLP with a large number of hidden neurones, such as that adopted here. The input vectors in the training set were randomized for each training cycle to improve generalization. Each set of training/testing vectors was used to train ten MLPs, so that, if an individual MLP did not train well due to the initial random weights, the overall result would not be significantly affected. The weights were assigned randomly in the region $[-0.5, +0.5]$. The sum of squared errors (SSE) parameter was computed over the entire testing feature set. Further to the initial training/testing data-set combination, it was decided to repeat the same neural network analysis but transposing the training and testing feature sets. Consequently, the MLPs were also tested with feature sets from experiments 1 and 3 and tested with data from experiments 2 and 4. MLPs trained and tested with such a combination were denoted using '24–13' instead of the '13–24' combination used above.

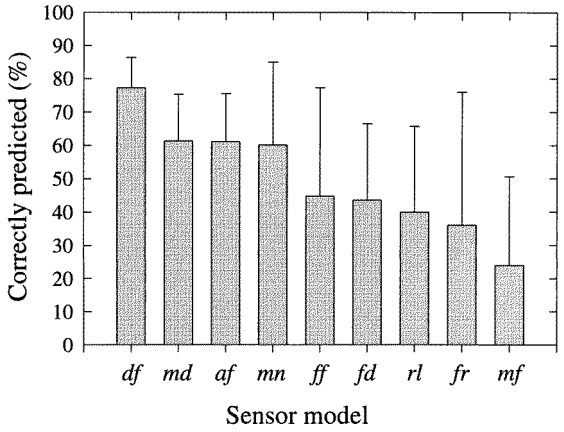

Figure 5. A bar chart showing the effect of the choice of the sensor model on the ability of the MLP to predict the head space of *E. coli* and *S. aureus* bacteria correctly. The error bars indicate the standard deviation of the result over different experimental runs.

Table 2 summarizes the performance of the MLP in classifying the bacteria correctly for each of the nine sensor pre-processing algorithms defined in table 1. The results have been ranked in table 2 according to the difference between the average percentages of correct, and incorrect plus unknown, classifications. Under this ranking criterion, the difference sensor parameter df does best and the modified fractional difference model mf worst of all.

Figure 5 shows a plot of only the percentage correctly classified for each of the nine sensor models, from which it appears that the difference sensor model df is better than the others at classifying the odour type. Similarly, the effect of the choice of normalization method has been investigated and figure 6 shows a plot of the percentage correctly classified for each of the three normalization methods (and none) defined in table 1, averaged over all of the sensor models. A ranking of the normalization algorithms was also made using the same criteria as before and was as follows (highest first): autoscaling a, sensor normalization s, no normalization n, and array normalization v.

The best performance of the MLP was actually found to occur when the minimum output sensor (compare with the previous best df) model was combined with sensor normalization mn/s since this resulted in 96.1% of all the test vectors being correctly classified; 2.2% were incorrectly classified and 1.7% unknown. In this case the SSE was only 74.21 and the data-set combination was 13–24.

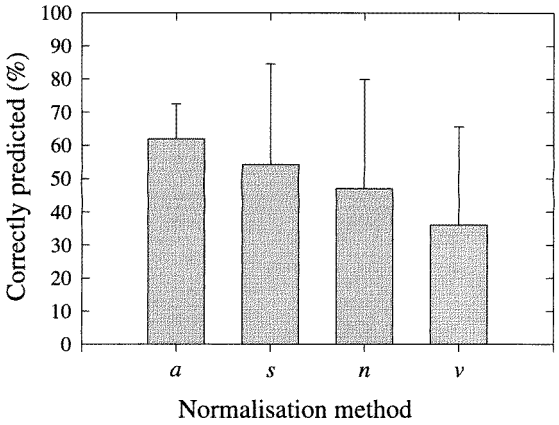

Figure 6. A bar chart showing the effect of the choice of the normalization method on the ability of the MLP to predict the head space of *E. coli* and *S. aureus* bacteria correctly. The error bars indicate the standard deviation of the result over different experimental runs.

Table 3. The confusion matrix showing the best performance of bacterial classification using a mn/s model. The accuracy of the classification is defined as $(166 + 180)/(180 + 180)$, namely 96.1%.

	True class	
Predicted class	*E. coli* (180)	*S. aureus* (180)
E. coli (166)	166	0
S. aureus (188)	8	180
Unknown (6)	6	0

The overall performance of a predictive classifier can be quickly appreciated through the use of a confusion matrix, which is shown in table 3 for the best MLP (mn/s). The overall accuracy was, as stated above, 96.1%, but more meaningful statements are that *S. aureus* was correctly classified 100% of the time whereas *E. coli* was only correctly classified 92.2% of the time. Thus, the prediction of *E. coli* is less accurate than is that of *S. aureus* due to there being a greater variance in the response vectors, but it is still a respectable figure.

4.2. The classification of the bacterial growth phase

The MLP design, topology, training methods and testing methods employed to predict the growth phase were identical to those used for the classification of the bacteria type—with the exception of the number of neurones in the output layer. Three neurones were used to represent the first three growth phases, lag, log and stationary. The true phases were obtained by inspecting the growth curves (such as figure 4) and locating the changes in the slope against time.

Another consideration was the natural occurrence of different numbers of training vectors for each growth phase. This can cause a problem because the network learns the vectors associated with the largest class more forcefully; however, the problem was avoided by selecting the training vectors for each class an identical number of times.

Table 4. The confusion matrix showing the optimal performance of bacterial phase classification using an af/a model. The accuracy of the classification is defined as $(2 + 154 + 133)/(14 + 162 + 182)$, namely 80.7%.

	True class of growth phase		
Predicted class	Lag (14)	Log (162)	Static (182)
Lag phase (4)	2	2	0
Log phase (198)	12	154	32
Static phase (134)	0	1	133
Unknown (22)	0	5	17

Once again the performance of each sensor model was investigated and ranked according to the difference in percentages stated above. However, in this case the minimum output sensor model (mn) did not perform best but was third with significantly better performances by the absolute final difference model (af) and the difference model (df). Similarly, the performance of the three different normalization techniques (and no normalization) were investigated and in this case the autoscaling (a) again ranked highest, followed by sensor normalization (s), no normalization (n) and array normalization (v). This result is not surprising because the intensity of the odour will be stronger during the log and stationary phases than it will during the lag (initial) phase and so array normalization weakens the classification.

The best performance overall was obtained using the absolute final output sensor model and autoscaling (af/a) when 80.7% of all test vectors were correctly classified, 10.3% were incorrectly classified and 6.2% were unknown. The value of the network error SSE was 423.4 and the data sets 24–13 were employed. The confusion matrix for the classification of the growth phase is shown in table 4. The imbalance in the class representation within the test sets is clear with only 14 test vectors for the short lag phase. Classification of the lag phase was, not surprisingly, the most difficult, only 14.3% being correctly classified, but it was more often mistaken for the log phase rather than unknown, which is encouraging. The other classes performed much better, the log phase being correctly classified 95.1% and the static phase being correctly classified 73.1% of the time. An important result to note here for clinical applications is that the lag and log phases were correctly predicted in 87% of all cases, since these earlier growth phases are of greater practical value.

5. Conclusions

An analytical instrument has been developed for the automated head-space analysis of bacteria cultures. The instrument has been used to collect large numbers of samples (360 per bacterium) for *E. coli* and *S. aureus*. Results show that the type of bacterium can be correctly predicted for 96% of all samples taken during a 12 h incubation period. In fact 100% of the *S. aureus* samples were correctly predicted, including the 14 samples taken during the first hour of incubation and in the lag phase! Moreover, the recent development of

an axisymmetrical temperature-controlled sensor chamber (Craven and Gardner 1997) should yield even better results. The rapid and fully automated identification of a bacterium could prove to be an important clinical application of an electronic nose in, for example, an ear, nose and throat department of a hospital.

The growth phase of the bacteria was correctly predicted for 81% of all unknown samples. This lower value may have been due to the way in which the growth phase was found by the visual inspection of the cell count curves. There is considerable scope for error at the interfaces between the lag and log phases and log and stationary phases which may be reduced in the future through the development of fuzzy membership functions. Nevertheless, an instrument that can automatically monitor the growth of bacteria should find application in testing the performance of new anti-biotic drugs on known pathogens. Furthermore, the accurate prediction of bacteria type within a time scale of minutes should help to optimize and therefore speed up the treatment of unknown bacterial infections.

This work demonstrates the potential application of an electronic nose in clinical medicine. (In related work by Wang *et al* (1997), it may be possible to diagnose diabetes with a five-element nose detecting elevated levels of acetone on the breath.) Field trials, testing cultures grown from swabs taken directly from patients, are now required. These studies should include the analysis of multi-cultured samples and the effect of the choice of growth medium, for we abandoned blood/brain fusion due to significant batch-to-batch variation, so further work is needed.

Acknowledgments

We are grateful to the Engineering and Physical Science Research Council and Mr David Morgan (Garfield Engineering) for the financial support of a studentship (Mark Craven) and to AlphaMOS (France) for the donation of the odour sensors.

References

Bishop C J 1995 *Neural Networks for Pattern Recognition* (Oxford: Oxford University Press)

Craven M 1997 *PhD Thesis* Department of Engineering, Warwick University

Craven M and Gardner J W 1997 *Meas. Control.* at press

Craven M A, Hines E L, Gardner J W, Horgan P, Morgan D and Ene I A 1994 *Neural Networks and Expert Systems in Medicine and Healthcare* ed E C Ifeachor and K G Rosen (Plymouth: University of Plymouth) pp 226–34

Fekadu A A, Hines E L and Gardner J W 1993 *Artificial Neural Networks and Genetic Algorithms* ed R F Albrecht *et al* (New York: Springer) pp 691–8

Gardner J W 1991 *Sensors Actuators* B **4** 109–16

Gardner J W and Bartlett P N (eds) 1992 *Sensors and Sensory Systems for an Electronic Nose* (Dordrecht: Kluwer)

Gardner J W, Hines E L and Tang H C 1992 *Sensors Actuators* B **9** 9–15

Gardner J W, Hines E L and Wilkinson M 1990 *Meas. Sci. Technol.* **1** 446–51

Gibson T D, Prosser O, Hulbert J N, Marshall R W, Corcoran P, Lowery P and Ruck-Keene E A 1996 *Proc. Eurosensors X, Leuven, Belgium, 8–11 September 1996* ed S Middelhoek (Amsterdam: Elsevier) pp 1341–4

Hines E L, Gianna C C and Gardner J W 1993 *Neural Networks: Techniques and Applications* ed P J Lisboa and M J Taylor (New York: Ellis-Harwood) ch 8

Holmberg M, Winquist F, Lundstrom I, Gardner J W and Hines E L 1995 *Sensors Actuators* B **26–27** 246–9

Kress-Rogers E (ed) 1997 *Handbook of Biosensors and Electronic Noses* (Boca Raton, FL: CRC)

Schweizer-Berberich M, Vaihinger S and Gopel W 1994 *Sensors Actuators* B **18–19** 282–90

Wang P, Tan Y, Xie H and Shen F *Biosensors Biolectron.* at press

Winquist F, Hornsten E G, Sundgren H and Lundstrom I 1993 *Meas. Sci. Technol.* **4** 1493–500