

IEEE Sensors 2004 Conference

Final Report

For the period

Starting October 24, 2004 Ending October 27, 2004

Work Performed Under Contract:

N00014-04-M-0424

Submitted By:

**Sensor Research and Development Corporation
17 Godfrey Drive
Orono, Maine 04473**

Greg Harmer, Ph.D.

Phone: 207/866-0100 x 253

Fax: 207/866-2055

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

Submitted To:

**ATTN: JOEL DAVIS
OFFICE OF NAVAL RESEARCH
BALLSTON TOWER ONE
800 NORTH QUINCY STREET
ARLINGTON, VA 22217-5660**

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188	
Public Reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comment regarding this burden estimates or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188,) Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE 16 DEC 2004		3. REPORT TYPE AND DATES COVERED FINAL REPORT, 24 OCT 04 - 27 OCT 04
4. TITLE AND SUBTITLE IEEE SENSORS 2004 CONFERENCE			5. FUNDING NUMBERS C -N00014-04-M-0424	
6. AUTHOR(S) GREG HARMER				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) SENSOR RESEARCH AND DEVELOPMENT CORPORATION 17 GODFREY DRIVE, ORONO, ME 04473			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES Published in the conference proceedings, to be submitted to IEEE Sensors Journal.				
12 a. DISTRIBUTION / AVAILABILITY STATEMENT			12 b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This work describes the use of an array of SAWs (surface acoustic wave devices) to detect and identify chemical warfare agents (nerve and mustard simulants) in a background of interferents (methanol, diesel and hexane). Even in a varying concentration and realistic gas delivery profiles, an array of SAWs is able to identify and predict the concentrations of simulants in real-time. This application differs from standard electronic noses (e-noses) as it is analytic. Much of the previous work in the field of sensor arrays and mixtures (usually binary or tertiary) tries to learn the whole sample space, due to the use of non-selective sensors. The SAWs used in this array, although not completely selective have a very high sensitivity to only the simulants. Thus, using the appropriate classification algorithms a real-time chemical detector can be realized. The system was tested with wide range of gas profiles as was able to identify and quantitate the simulants to a high degree of accuracy.				
14. SUBJECT TERMS Chemical warfare agents, detection, SAWs, classification, quantitation			15. NUMBER OF PAGES 22	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OR REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION ON THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

NSN 7540-01-280-5500

Standard Form 298 (Rev.2-89)
Prescribed by ANSI Std. Z39-18
298-102

Enclosure 1

IEEE Sensors 2004 Conference

Greg Harmer

24-27 October 2004

Contents

1	Introduction	2
2	Algorithm Preparation	2
2.1	Verification	4
3	Data Preparation	4
3.1	Sensor Validation	4
3.2	Training Data	6
3.3	Testing Data	6
4	Conference Seminars	10
5	Summary	15
	Appendix	15
A	Conference Paper	16
B	Presentation Slides	20

1 Introduction

Detecting chemical vapors requires a sensing system combined with signal processing algorithms. The sensing systems can range from a sophisticated GC-MS to an array of simple chemical sensors that are found in typical e-nose applications. The work described in this paper uses surface acoustic wave sensors (SAWs) that are designed primarily for the detection of chemical warfare agents (CWAs). This differs from most enose applications [1] slightly as they only try to remember the signature of a particular mix of gases (i.e. smell) with a limited ability to isolate the particular gases or determine their concentration.

The motivation for this work is to develop a handheld chemical detector specifically targeted at CWAs for use by personnel such as the military and homeland security. The requirements of a handheld device rule out well known expensive and bulky solutions such as GC-MS or laser spectroscopy systems. Thus, an array of inexpensive, simple sensors are employed with the use of pattern recognition processing.

Currently a laboratory based system is used. However, to better simulate real world conditions it is open to the environment where laboratory air is pulled through a large PVC pipe in which the CWA simulants (DCP (1,5-dichloropentane) - mustard, DMMP (dimethyl methylphosphonate) - nerve) and interferents (methanol, hexane and diesel oil fumes) are injected, shown in Fig. 1. A flow of about 5 L/min is allowed through the 5 cm diameter pipe and 200 ccm is drawn off to a small 22 cm³ sensor chamber.




Figure 1. The experimental system.

2 Algorithm Preparation

In a typical enose system the preprocessed data is converted into meaningful results via the use of one or more classifiers such as neural networks, support vector machines or probabilistic algorithms for example. The details of the classifier will not be discussed here, but instead will be treated as a black box system that outputs a list of classes and their corresponding confidence levels. The selected class for an unknown input pattern is the class with the highest confidence level, $k = \arg \max_j \eta_j$, where η_j is the confidence of class j being the unknown pattern.

Three strategies can be used for identification and quantification shown in Fig. 2. Brief details of these three strategies follow.




Figure 2. Methods used for classifying and quantitating analytes.

- **One Step Method**

This method is one of the easiest approaches, specially if you have an 'off-the-shelf classifier' available. Each analyte class is divided into several subclasses of different concentrations so the classifier actually trains to gas/conc pairs. This requires the concentrations to be discretized otherwise there may be reasonable overlapping of the classes, which may degrade the performance of some classifiers.

This method is best suited for analytes that only occur at specific concentrations. To increase the concentration resolution the number of classes need to be increased, which may cause the classification problem to quickly get very large.

Notwithstanding, probably due to the setup simplicity it has been used for a number of applications [1], though it is acknowledged that this method is not appropriate if the problem size is likely to expand.

- **Two Step Method**

The analyte is first determined through the use of a classifier, then the concentration is found using this knowledge. Once the analyte is known, the task of concentration determination is reasonably easy, and a number of methods would suffice.

It is important to choose the correct feature types to use at each stage. In the first stage we wish to eliminate any concentration information, leaving only the analyte information. Since the concentration predominately effects the magnitude of the response normalizing the points seems a logical choice. For the second stage, the use of magnitude information with simple prediction algorithm would do the job. Dividing the problem into two parts vastly reduces the computational demand on the classifier and the concentrations can be from a continuous range. This method has also been used by a number of people [2].

- **Mapping Method**

The idea for this method is not new, though a recent paper where it is applied to e-noses highlights its potential [3]. It uses the fact that each sensor responds slightly different to each analyte. By using the response-concentration mapping function for each sensor/analyte combination, the inverse function can be used to back-calculate the concentration from the response.

For each gas/sensor combination the response-concentration dependence can be described by a mapping (or transfer) function,

$$r_{ij} = f_{ij}(c_j), \quad (1)$$

where $i = 1, \dots, m$ sensors, $j = 1, \dots, n$ analytes and c_j is the concentration of the j th analyte.

Typically, we expect this curve to be nonlinear, which could be modeled by cubic splines or fitted to some parametric equation. To find the required inverse, f needs to be monotonic, otherwise f^{-1} will be bistable and is of limited use when trying to back calculate the concentration.

Moreover, the use of the algorithm only requires the use of f^{-1} for the calculation of

$$c_{ij} = f_{ij}^{-1}(r_i), \quad (2)$$

which we trained a small neural network to learn. The trade-off for versatility is it does not extrapolate outside the concentration limits of the training data, though a modification would not be difficult to implement.

If the unknown candidate is from the j th analyte then the concentration calculated from all the sensors should be very close, whereas they should vary widely for the other gases. Hence, we can use $\sigma_j^2 = V[c_{ij}]$ as a "closeness" measure for each j . Therefore the best candidate will have the smallest variance (or standard deviation), $k = \arg \min_j \sigma_j^2$ with a concentration of $\mu_k = E[c_{ik}]$. Using the standard deviations we defined a confidence indicator of

$$\eta_k = \frac{1/\sigma_k}{\sum_i 1/\sigma_i}. \quad (3)$$

This is also a one step method and the analytes can have a continuous concentration range.

The main difference between the mapping and the two step method is the way the winning analyte is chosen. Either by a purpose built classifier or a minimum error approach respectively.

Arguably this could be considered as an integrated two step methods approach since we require the knowledge of gas identify before we can give a concentration. However, we will list it separately since it does not require a purpose built classifier. One may ask how does this method differ from using the steady-state fingerprints that are commonly used? An explanation is given in the paper.

The following work is based on the mapping method for classification. This method offers a slight processing gain over using standard steady-state fingerprint methods, details are given in the paper. Additionally, the method is easy to implement in terms of necessary training data and microprocessor implementation. It is also very tolerant to failed sensors (does not require retraining), which easily allows redundancy to be built in to the system.

2.1 Verification

The algorithms were first verified in Matlab with artificially generated data. Once verified, selective sensors to DCP and DMMP were selected and exposed to the training protocol (described next section). the data gained was processed off-line in Matlab.

The testing part of the classifier coded in LabView to work real-time in one of SRD's test systems. A number of tests were run in real-time to verify the operation of the classifier, in particular its robustness. The algorithm was still able to perform when sensors were removed at random. Sensors were also remove automatically when there was a large spike in the output due to interference.

3 Data Preparation

3.1 Sensor Validation

In order for the mapping method to work, the magnitude of response must be consistent independent of the history of the sensor. That is, the sensor needs to be reproducible with respect to the input concentration. Note that when the concentration changes sharply, some time is required for equilibrium to be reached, which is what needs to be reproducible. To verify our sensors are reproducible five gas profiles that have vastly different starting profiles but that end with the same concentration were tested, shown in Fig. 3. Three minutes were given at the final concentration to allow the sensor to reach equilibrium.




Figure 3. The left plot shows the five profiles for DCP, with the right plot showing the two possible background profiles using diesel. Other gases have same shape profiles just with different flow rates.

Figure 4 shows the corresponding responses to the profiles of Fig. 3 of a SAW towards DCP, DMMP, diesel and methanol. For DCP there is excellent consistency at the end of the 10 minutes exposure. The error bars show one standard deviation about the mean, thus the small error bars denote good repeatability. The consistency of DMMP and diesel is not as good as DCP, but sufficient.

For this work methanol is an interferent and SAWs have been designed to not react with it. Thus the large error bars coupled with the small response magnitudes indicate the responses are not much more than noise, hence this is a good result. The results are similar for hexane (plots not shown). The other interferent, diesel on the other hand shows a quite measurable response due to the chemical similarity of some its components to DCP. However, the magnitude of response to diesel is much lower then those for DCP and DMMP.




Figure 4. SAW responses to the validation profiles shown in Fig. 3. The mean of five responses is shown with the error bars showing $\pm\sigma$ about the mean.

Figure 5 shows the corresponding responses to the profiles for one of the SMOs (WO_3). The consistency at the end of the 10 minute exposure for DCP, diesel and methanol is very good. For DMMP the SMO is very slow at tracking the concentrations, which leads to its poor performance. Additionally, the SMOs do not show any selectivity towards the analytes or immunity towards the interferents.

In comparison of these SMO responses to the SAW responses (Fig. 4), the SAWs show more 'crispness', which indicates they reach equilibrium quicker, as evident by shorter recovery times.

When dealing with individual gases the sensors do not necessarily have to be selective, they just need to have different mapping functions. However, when dealing with interferents the sensor needs to be reasonably selective, which is shown by the magnitudes of the responses to each of the gases in Fig. 4. Figure 6 shows the responses for DMMP and DCP with diesel in the background carrier air; the top two plots are SAWs and the bottom two plots are SMOs. Diesel was chosen because it is one of the strongest influencing interferents for both SAWs and SMOs. That is, it is the interferent they have immunity to.

The SAW responses (top plot in Fig. 6) show very good selectivity towards DCP and DMMP. The responses to the analyte alone and with a background interferent are very repeatable, especially for DCP. The variance in the DMMP response may be due to the use of the open air test system (shown in Fig. 1). Towards the end of the exposures where the analyte concentration is constant there is a very small influence of the drifting background.

The corresponding responses for SMOs (bottom plots in Fig. 6) do not show the same selectivity as the SAWs. Whilst the individual and constant background exposures are repeatable to responses to the drifting background are quite attenuated.

Primarily for this reason only SAWs were used in the construction of the sensor array. Thus, the results shown in Fig. 4 and Fig. 6 show that the SAWs have the characteristics that are required by the mapping function. That is they have good selectivity towards the gases to be detected and they are very invariant to the exposure history of the gases.




Figure 5. SMO responses (mean of five) to the validation profiles shown in Fig. 3.

3.2 Training Data

One of the advantages of the mapping technique is the procedure of training is relatively simple. We simply require the concentration/magnitude relations for each gas sensor pair. This is obtained by measuring the response at different concentrations and using some technique to learn this relationship. A sequence of hits with increasing concentration was delivered to the sensors and the response magnitude measured at the end of the exposure. Examples for a SAW and SMO are shown in Fig. 7. Three minute hits were used to approximate steady-state operation. As seen from Fig. 7 this is ok for the SAWs, but the SMOs would require more time to reach the same level of approximation.

To learn the mapping function (or the inverse of it more precisely) a small neural network was trained for each gas-sensors pair. A neural network was chosen because it is versatile and can learn a wide range of nonlinear relationships without having to specify parameters for a fixed relationship. The mapping functions for two SAWs are shown in Fig. 8. The interpolated curve is from the neural network that was trained from the test points. The plots clearly show the mapping functions have been accurately learnt.

3.3 Testing Data

To evaluate the two versions of the mapping method and their performance a series of 18 test protocols were developed, a selection is shown in Fig. 9. They start with simple conditions of individual hits with different hit times, and increase in difficulty where the hits and the background interferences both have time varying concentration profiles. In the last few protocols the hits are just smooth increases in concentrations from zero with no well defined instant where the gas is turned on, which is to model a plume like scenario.

The results for three versions of the mapping methods are shown in Fig. 10, along with the gas delivery profile and raw SAWs responses. Figure 10a classifies to the analytes and interferences, Fig. 10b only classifies to the analytes, and Fig. 10c is the same as (b) except a threshold on the variance of the back calculated concentrations is in place. These




Figure 6. SAW (top plots) and SMO (bottom plots) responses to the validation profiles with a constant and drifting diesel background.




Figure 7. Training data for a SAW (left) and SMO (right) to DCP. The magnitude is taken at the end of the exposure.




Figure 8. Plots of the response magnitude for a SAW (left) and SMO (right) to DCP, DMMP and diesel. The large markers are the data points from training and the interpolated points are from the neural networks.




Figure 9. A selection of gas delivery protocols used to test the classifier.

results show the sensors are selective against diesel, however when methanol is delivered in the background of diesel a noticeable response is produced. Interestingly, when delivering methanol alone there is a negligible response. In general, the best results over all the profiles as highlighted above was achieved using the method of Fig. 10c. Note how the threshold caused the omission of the transient parts of the response, which should filter out false readings while sensors stabilize.




Figure 10. Comparison of the mapping classification and quantitation methods. DMMP is shown at 10x actual concentration.




Figure 11. Comparison of the four methods to process the variances of the concentrations. DMMP is shown at 10x actual concentration.

Another problem when dealing with interferences, or any wide variety of gases, are the different concentration ranges each gas can be present at. For example, assume three sensors determine a given response is gas A with concentrations 8, 9, and 10, or gas B with concentrations 207, 209, and 214. Using the standard rule of $k = \arg \min_j \sigma_j$ would give gas A as the best choice. However, with respect to the magnitude of the concentrations gas B may be the better choice. Four options for the argument of the $\arg \min_j$ expression are listed below and evaluated for certain conditions.

- a) raw-straight: $k = \arg \min_j \sigma_j$
- b) raw-normalized: $k = \arg \min_j \sigma_j / \mu_j$
- c) log-straight: $k = \arg \min_j \hat{\sigma}_j$
- d) log-normalized: $k = \arg \min_j \hat{\sigma}_j / \hat{\mu}_j$

The overhat denotes the standard deviations or means are taken from the log concentrations. When there is a clear cut winner the choice of variance calculation does not make much difference. It is when there are two or more gases that are at about the same variance measure that the choice of methods becomes important. For this reason sequences were focused on (where diesel was included in the classifier) where there were misclassifications or low confidence levels.

Figure 11 shows the results of the four variance calculation methods for one of the test sequences that has individual hits of DCP, DMMP and diesel. It is clear that using unmodified concentrations tends to bias the gas choice to those with lower concentrations. Since DMMP has the lowest concentration it often tends to be the likely candidate, as shown in Fig. 11c. The results of methods a) and b) are very similar and give the best results. The method b) does a satisfactory job, but has a slight tendency to select DCP over DMMP. It should be noted that the performance of each method may vary with the types of sensors used and the application they are used in.

4 Conference Seminars

There were 625 participants. 445 papers were accepted, which was 55 % of the total submitted papers. There were 7 parallel tracks (6 oral and 1 poster) for the three days. The poster sessions had the same weighting of papers. i.e. will be in the proceedings with a full write-up.

The following subsections are very brief summaries of some selected seminars. Regarding the code: the first letter is the weekday, the second number is the session number, the third letter is [L]ecture or [P]oster, the next letter A-F indicates the parallel stream, and the last number the seminar number in the session.

The proceedings will be sent out on CD within a couple of weeks of the conference.

M2P-P2: A portable gas recognition system based on metal oxide gas sensor array

This was just commercial SMO sensors packaged into some type of portable unit used to distinguish between SO₂ (20ppm), NO₂ (1-5ppm), H₂S (1-5ppm). Don't recall seeing mixtures, but need to check the proceedings. They used 50 min hit, 50 min recovery cycles, which is very long; and they extract feature from the transient response for the features.

M2P-P4: The characterization of nanostructured copper-doped tin oxide films for gas sensor microarrays

→ Just to note they used RF magnetron sputtered technique to produce SnO₂-CuO sensors that were sensitive towards NH₃ (this is a gas of interest to SRD).

M3L-A-2: Metal oxide semiconductor gas sensors utilizing modified zeolite catalysis to improve selectivity

Uses molecular sieve materials to increase selectivity by targeting particular gases. The sensors are SMOs (2 mm²) at 400 C, the sieve materials are zeolites. These are placed on the sensor with thickness of 0, 50 μm and 100 μm. Used several different types of zeolites. They showed the sensitivity decreased as the zeolite was made thicker. Several sensors were integrated into a nose system.

→ They showed zeolites changed the sensitivity but no golden bullet for the selectivity as you might hope from a sieve material. There was nothing in their results that suggested anything promising that we didn't see.

M3L-A-3: Gas sensors based on nanoparticle WO₃ thick films

They used commercial WO₃ nanopowder from Aldrich, this was mixed with InCl₃ or BiCl₃ and screen printed onto the platforms. It was annealed at 600 °C, which removes the chlorides, just leaving the doped WO₃ I suspect. The dopants were at 1.5, 3.0 and 5.0 percent weight; run at temps 250, 300, 350; exposed to NO₂ (1-10ppm), CO (10-100ppm), NH₃ (10-100ppm); no mixtures.

They ran some standard characterization tests: PCA using sensitivities, SEM showed average grain size 60-80 nm, X-ray spectroscopy, XRD analysis. Provided two mechanisms of sensing for Bi and In dopants.

→ Probably a reasonable piece of work, but I didn't see anything that was particularly novel or new.

M3P-P8: ENT bacteria classification using a neural network based Cyranose 320 electronic nose

I was interested in this due to the use of the Cyranose 320. They used the Cyranose to get the raw responses and custom processed the signals. Three unsupervised methods (PCA/fuzzy k means/SOM) were used to extract the features, which were then classified using supervised classifiers.

→ An application of several classification algorithms with the Cyranose 320.

T2L-B-1: Improving the classification accuracy in electronic nose using multi-dimensional combining (MDC)

Again, the Cyranose 320 was used to gather the raw data. A number of features were extracted from the transient response, these were reduced in dimension by either PCA, ICA or LDA. These reduced features were then passed into a number of classifiers which were combined (I think this is referred to as the MDC). They said the LDA transform performed the best.

→ I am still a little unsure on exactly the implementation of the MDC, but will check out the paper when the proceedings arrive.

T2L-B-2: Principal discriminant analysis of gas sensor transients

When we wish to reduce the dimensionality of raw data, the method used depends on whether the information is contained in the means (use LDA) or in the variances (use PCA) of the data. However, what if the information is a combination of both? Combine the PCA and LDA.

Introduce a regularization parameter,

$$T = (1 - \gamma)S_w^{-1}S_b + \gamma C_x, \quad (4)$$

where the first part is the component of LDA and second part is the component of PCA. We then just project the n eigenvectors of T in the normal way. The results shown, generally indicated that as the number of training data increased the higher the proportion of LDA that was used. See Fig. 12 for some sketches of the results. The downside is that to find the optimal γ an iterative search (using some variant of the hold-out method for example) needs to be performed, so it is dependent on the type of classifier you use I suspect.

→ This was an informative seminar that contained a novel approach I haven't seen elsewhere, something that I will integrate into SRDs processing software.

T2L-B-3: Feature evaluation for an electronic nose

Uses six SMOs with five features extracted from the transient response, which gives a total of 30 elements in the feature vector. The question asked is what subset (if any) of these features gives the same or better performance. That is, what features give the most information. Benchmarking against using all the features an exhaustive search was done with up to five features (using 3-nearest neighbor classifier for validation). They showed best results with 5 features, which performs better than when using all the features.

→ This agrees with what SRD has found with classification, a sensor must provide new information to the array otherwise it just dilutes the effort of the others. I would like to see how a PCA/LDA of the full vector performs in comparison for each number of features – would be a more *apples to apples* benchmark. I'll check the proceedings when it comes out.




Figure 12. (a-b) Showing the how the number of training samples affects the optimal regularization parameter. Comparisons of PCA, LDA and T were shown for colas (c) and juices (d).

T2L-B-4: Principle component analysis and pattern recognition combiner with visible spectroscopy in the classification of food quality

Uses a laser system (via an optical fibre into an oven) to monitor the color of the cooking process of certain products (like chickens for examples). Feature dimension of 945 is reduced to 3 via PCA. They showed a PCA diagram for different cooking times. Basically a snake-like shape, similar to PCAs of an SMO with different concs of the same gas.

→ This is just a color mapping process, nothing particularly new here. Not sure why a cheap CMOS chip couldn't be used.

T2P-P20: Wafer level packaging of micro-machined gas sensors

This grabbed my attention due to its use of micro-hotplates and small sampling chamber (something like we've discussed as SRD). However, upon questioning him (D. Braid) on the source of their hotplates, he could not say – proprietary information.

T2P-P13: QCM sensors for ammonia gas using clay/polyelectrolyte layer-by-layer self-assembly

Basically described making highly sensitive NH_3 QCM sensor. He claimed (verbally) it was selective against ethanol ('ols') and H_2S ('sulfides').

T3L-A-3: Millisecond response time measurements of high temperature gas sensors

Use of SiC sensors to deal with harsh environments for the detection of H_2 based species. Measurement uses the capacitance, they adjust the voltage to maintain a constant capacitance, then we can tell what gas is present due to the applied voltage, see Fig. 13c.




Figure 13. (a-b) Schematics of the gas delivery method. (c) The relationship between C and v_g for O_2 and H_2 . (d) Sensors response.

In order to change the gas delivery quickly, they set up a laminar flow of the two gases (adjacent and parallel) This was verified using laser induce fluorescence on one of the gases. This is subjected to a sensor that is mounted on a platform that can rotate three degrees in 2 ms, see Fig. 13a-b. A typical response is shown in Fig. 13d.

→ A novel gas delivery setup with gas responding sensor.

T3P-P3: Fast variable selection for gas sensing applications

Commonly, sensor systems have many 100's of variables (i.e. feature length), and we need a method to choose the optimal subset (cf T2L-B-3). They first use a coarse method to cull out the non-specific variable, then they use a fine tuning method to refine the number of variables.

The coarse method uses a separability measure for each dimension (variable) and those over a particular threshold are selected. The separability measure seems very similar to the t or z -statistic (though this was not mentioned). Fine tuning

was done by forward selection (additive from zero) or backward elimination (subtract from full coarse set). GA's were used as a benchmark. Further questioning about the poster revealed the most of the smaller sets of variables were subsets of larger ones even if derived by different methods.

→ The use of the t or z -statistic is one of the proposed processing methods for reducing the dimensionality of the spectra for the In-Aid project.

T4L-C-1: Challenges and opportunities in multi-sensing microsystems - a case study

This seminar was pretty much an overview of current projects. Of interest was the micro-GC (MEMs): greater than 30 components, less than 10 ppb with pre-conc, 2-10 min analysis, less than 1 cm^2 . See Fig. 14 for a schematic.




Figure 14. A schematic of the micro-GC system.

The long term version will have $1\text{-}2\text{ cm} \times 5\text{ cm}$ and less than 10 mW per analysis.

T4L-C-2: Source localization with a network of electronic noses

A two-step approach for locating an emission source based on analytical diffusion-advection model using point concentration measurements. These can be taken from an enose for example. The advection part of the transport model helps eliminate multiple solutions.

→ This sounds useful as application to the LACIS system as described in the HSARPA proposal. I will have to look into in more depth in the proceedings.

T4P-P23: A precise $1/f$ noise spectroscopy setup for sensor characterization

This describes a method for more accurately determining the $1/f$ spectra, which involved quite a complex setup.

→ The main interest here was the similarity to the noise work done for ONR. They used thyristor or something, but have done it with SMO, they are a lot more noisier. Interesting they found the f_{cutoff} at about the same frequency of 10ish Hz that we found from our noise work.

W2L-B-2: Cancellation of chemical backgrounds with regularized Fisher's linear discriminants

The problem is cross selectivity in gas sensors; how to cancel multivariate responses? They considered three gases in mixtures and neutral (air?), which gives 7 combinations. The analysis of a binary mixtures can be shifted to one component by adapting to the other one. To do this they reformulate the LDA to discriminate primary analytes rather than each mixture, see Fig. 15.




Figure 15. Consider 3 gases in uni-, bin-, and tertiary mixtures. (a) The straight out LDA treating each group separately, then in (b) drawing out all the mixtures with A in them.

They are able to assign proportions to the selectivity of each groups. Seems a bit like just relabelling, but there is a little more to it.

→ Looks interesting. To pick out one group you would need to know in advance what group to pick out as to know what LDA transform to use – maybe do all and see which one matches best (i.e. has the highest confidence). I will need to check out the proceedings to get more of the mathematical details.

W2L-B-5: Optimized multi-frequency temperature modulation of micro-hotplate gas sensors

→ The interest in this seminar was the use of hotplates. I spoke to the presenter after the seminar and he gave two sources (collaborators or theirs) of hotplates. I am in the process of following this up to get points of contacts and specifications.

5 Summary

This work has demonstrated the use of SRD's sensor technology in a real-time application, which can identify the gas and quantify its concentration. This was achieved using a relatively novel technique that utilizes the correlations between individual sensors in an array. This technique is robust in that bad sensors can be ignored without the need for retraining the algorithm. It therefore implicitly supports the use of redundant sensors in the array. The results were presented at the international IEEE Sensors 2004 conference in Vienna, Austria with a presentation and a paper that will appear in the proceedings. The proceedings will also be listed in IEEE Xplore where they are searchable via the web.

At the conference information was also gained from seminars and posters by other researchers in the field. A brief summary of selected papers is given in Sect. 4. Of particular interest were the signal processing topics dealing with principle component analysis and linear discriminant analysis, and some possible suppliers of micro-hotplates for the SMO platforms.

Detection of chemical warfare agents in the presence of interferents

Gregory P. Harmer

Sensor Research & Dev. Corp.
Orono, ME, USA
gph@ieee.org

Chuncui Yang

Sensor Research & Dev. Corp.
Orono, ME, USA
cyang@srdcorp.com

Brent T. Marquis

Sensor Research & Dev. Corp.
Orono, ME, USA
bmarquis@srdcorp.com

Abstract

This work describes the use of an array of SAWs (surface acoustic wave devices) to detect and identify chemical warfare agents (nerve and mustard simulants) in a background of interferents (methanol, diesel and hexane). Even in a varying concentration and realistic gas delivery profiles, an array of SAWs is able to identify and predict the concentrations of simulants in real-time. This application differs from standard electronic noses (e-noses) as it is analytic. Using SAWs that are highly sensitive to only the simulants with the appropriate classification algorithms a real-time chemical detector can be realized. The system was tested with wide range of gas profiles as was able to identify and quantitate the simulants accurately.

INTRODUCTION

Detecting chemical vapors requires a sensing system combined with signal processing algorithms. The sensing systems can range from a sophisticated GC-MS to an array of simple chemical sensors such as those found in typical e-nose applications. The work described in this paper uses surface acoustic wave sensors (SAWs) that are designed primarily for the detection of chemical warfare agents (CWAs). This differs from most e-nose applications [3] slightly as they only try to remember the signature of a particular mix of gases (i.e. smell) with a limited ability to isolate the particular gases or determine their concentration.

The motivation for this work is to develop a handheld chemical detector specifically targeted at CWAs for use by personnel such as the military and homeland security. The requirements of a handheld device rule out well known expensive and bulky solutions such as GC-MS or laser spectroscopy systems. Thus, an array of inexpensive, simple sensors are employed with the use of pattern recognition processing.

The sensors are based on SAW technology [2] with a 915 MHz quartz crystal oscillator used as the transduction mechanism for a polymer sensing film. Based on the mass loading phenomena, the amount of change in frequency is related to the ambient concentration of the targeted gas molecules. The 915 MHz oscillators are used because the sensitivity increases with frequency squared [2].

To better simulate real world conditions, the laboratory based system is open to the environment where room air is pulled through a large PVC pipe in which the CWA simulants (DCP (1,5-dichloropentane) - mustard, DMMP (dimethyl methylphosphonate) - nerve) and interferents (methanol, hexane and diesel oil fumes) are injected. A flow of about

5 L/min is allowed through the 5 cm diameter pipe and 200 ccm is drawn off to a small 22 cm³ sensor chamber.

The remainder of this paper describes several classification strategies that can be used for classification and quantification, and how to deal with interferents. The results show how our sensors perform in the presence of interferents using realistic, time-varying concentration profiles.

CLASSIFICATION AND QUANTIFICATION

In a typical e-nose system the preprocessed data is converted into meaningful results via the use of one or more classifiers such as neural networks, support vector machines or probabilistic algorithms for example. The details of the classifiers will not be discussed here, but instead will be treated as a black box system that outputs a list of classes and their corresponding confidence levels. The selected class for an unknown input pattern is the class with the highest confidence level. Three strategies that can be used for identification and quantification are discussed.

One Step Method

This method is one of the easiest approaches, specially if you have an 'off-the-shelf classifier' available. Each analyte class is divided into several subclasses of different concentrations so the classifier actually trains to gas/conc pairs. This requires the concentrations to be discretized otherwise there may be overlapping of the classes, which may degrade the performance of some classifiers.

This method is best suited for analytes that only occur at specific concentrations. To increase the concentration resolution the number of classes need to be increased, which may cause the classification problem to quickly get very large. Notwithstanding, probably due to the setup simplicity it has been used for a number of applications [3], though it is acknowledged that this method is not appropriate if the problem size is likely to expand.

Two Step Method

The analyte is first determined and then the concentration is found using this knowledge. Once the analyte is known, the task of concentration determination is reasonably easy, and a number of methods would suffice.

It is important to choose the correct feature types to use at each stage. In the first stage we wish to eliminate any concentration information, leaving only the analyte information. Since the concentration predominately effects the magnitude of the response normalizing the points seems a logical choice.

The second stage can be accomplished by using the magnitude information. Dividing the problem into two parts vastly reduces the computational demand on the classifier and the concentrations can be from a continuous range.

Transient and Steady-State Information

The choice of features (i.e. transient or steady-state information) to use depends on the application of the sensors; if they are to be used in an array configuration or individually.

The transient information is the 'shape' of the response, that is the magnitude as a function of time. This response information is very rich, it gives an indication of the underlying film/gas reaction that is taking place. Using this information alone, even with a single sensor, SRD is able to discriminate between gases within a limited concentration range.

However, the transient information assumes the gas delivery conforms to a specific concentration profile, usually a step-pulse, which needs to be exactly repeatable. The repeatability of the gas delivery is critical because the shape of the response is strongly influenced by the gas delivery profile. The problem with using the subtle differences between response shapes is they will be dominated by the gas delivery profiles if they are different. Hence, the delivery profile ends up being classified as opposed to the gas identity.

For certain applications extraction of the transient information may be acceptable if there is total control of the gas delivery to the sensors. However, for a handheld detector that is continuously sampling from the environment there is no control of the concentration profile of the delivered gas. In these circumstances the only available, reliable information is the instantaneous magnitude of the sensors. In this case each sensor provides a single point of information – not enough to characterize even two gases. Thus, the use of an array of sensors is imperative for successful discrimination between gases.

In the first step of the two step classification process the responses are first normalized against the baseline to calibrate to their current operating point, for the SAWs this involves a dc shift. To remove concentration information the magnitudes need to be standardized. For the transient information the response shapes are scaled to unity (see Fig. 1a) and for the steady-state information the magnitudes are scaled to zero mean and unity variance (see Fig. 1b). This results in the fingerprint for a particular gas, which can then be passed to a classifier. Once the gas is determined the unscaled magnitudes can be used to determine the concentration in the second step.

Although the gas delivery becomes less critical when using steady-state information, the monitoring of the baseline becomes vital. Typically, the sensor baseline drifts due to environmental fluctuations such as temperature, humidity and pressure. Usually these are very slow-time varying and do not present much of a problem. However, in certain instances they may change on the same time scale as a legitimate response, which presents a problem. In terms of detection, false negatives may be avoided by implementing an outlier




Figure 1. (a) Transient fingerprint of one sensor. (b) Steady-state fingerprints of eight sensors.

detection algorithm. To alleviate the extra burden this places on the baseline monitor feedback from the classifier could be incorporated.

Mapping Method

The idea for this method is not new, though a recent paper where it is applied to e-noses highlights its potential [1]. It uses the fact that each sensor responds slightly different to each analyte. By using the response-concentration mapping functions for each sensor/analyte combination, the inverse function can be used to back-calculate the concentrations from the response.

The best candidate will have the smallest variance (or standard deviation) $\sigma_j^2 = V[c_{ij}]$ with a concentration of $\mu_k = E[c_{ik}]$ where $k = \arg \min_j \sigma_j^2$. Using the standard deviations a confidence of $\eta_k = (1/\sigma_k) / \sum_i (1/\sigma_i)$ was defined. The main difference between the mapping and the two step method is the way the winning analyte is chosen; via a purpose built classifier or a minimum error approach respectively. Arguably the mapping method could be considered as an integrated two step approach since it requires the knowledge of the gas before concentration determination. However, it is considered separate since it does not require a purpose built classifier. One may ask how does this method differ from using the steady-state fingerprints that are commonly used? The explanation follows below.

Comparison of Mapping and Steady-State Methods

Intuitively the mapping and steady-state methods may appear similar, both employ the rescaling of the raw magnitudes to determine which gas is the most likely. The steady-state method uses a classifier to do this whereas the inverse uses a comparison of variances.

However, a very important distinction is that the mapping method implicitly caters for nonlinear concentration relations whereas the steady-state method assumes linear relations. This is evident by the scaling process, the steady-state method scales all the magnitudes by the same constant amount. In contrast, the mapping method uses the trained mapping functions for each gas/film combination. The scaling method for the steady-state could be extended to be nonlinear and gas dependent if desired, in which case the two methods would effectively be the same.

To illustrate this concept, an artificial eight sensor array was

created with arbitrary mapping functions for four gases. In the first case linear mapping functions were used. To easily view the fingerprints the first two principle components are plotted, shown in Fig. 2a. Since the sensors are linear all the points for the same gas coincide. This shows that irrespective of the concentration, the scaled fingerprints within each gas are identical. Figure 2b shows the same plot for quadratic mapping functions. The points for each gas do not coincide due to the use of a linear scaling technique with inherently nonlinear (quadratic) data. The results for the mapping method (not shown) were identical in both cases. Thus, it is clear that the mapping method is more robust against nonlinear sensor characteristics.




Figure 2. Linear (a) and quadratic (b) fingerprints.

This is one of the reasons for selecting the mapping method. Additionally, it is easy to generate training data since steady-state information is more immune to difference in gas delivery and it is very tolerant of failed sensors (does not require retraining).

Dealing with Interferents

Ideally perfectly selective sensors would be available that only respond to one gas of choice in which case interferences could be ignored. But this is almost never the case, therefore signal processing is required to compensate. Two options are:

- Train to the analytes and interferences. This may suffice for a small number of interferences, but could become too complex for many interferences. This method could have trouble with false-positives, i.e. analytes are interpreted as interferences.
- Only train to the desired gases and tighten the robustness controls, like input limits and variance thresholds for example. This would be more robust when many interferences are possible, but could have trouble with false-negatives, i.e. interferences are interpreted as analytes, which is usually more preferable than false-positives.

Another problem when dealing with interferences, or any wide variety of gases, are the different concentration ranges each gas can be present at. For example, assume three sensors determine a given response is gas A with concentrations 8, 9, and 10, or gas B with concentrations 207, 209, and 214. Using the straight variance the best choice is gas A. However, with respect to the magnitude of the concentrations gas B may

be the better choice. Four options to standardize the variance are listed below.

- a) raw-straight: $k = \arg \min_j \sigma_j$
- b) raw-normalized: $k = \arg \min_j \sigma_j / \mu_j$
- c) log-straight: $k = \arg \min_j \hat{\sigma}_j$
- d) log-normalized: $k = \arg \min_j \hat{\sigma}_j / \hat{\mu}_j$

The overhat denotes the standard deviations or means are taken from the log concentrations. Alternatively, the means response magnitudes could be used in place of concentrations. When there is a clear cut winner the choice of variance calculation does not make much difference. It is when there are two or more gases that are at about the same variance measure that the choice of methods becomes important. For this reason sequences were focused on (where diesel was included in the classifier) where there were misclassifications or low confidence levels.

RESULTS

In order for the mapping method to work well, the same magnitude of response independent of the gas exposure history of the sensor is required. Note that when the concentration changes sharply, some time is required for equilibrium to be reached, which is what needs to be reproducible. To verify our sensors are reproducible five gas protocols that have vastly different starting profiles but that end with the same concentration were employed. Three minutes were given at the final concentration to allow the sensor to reach equilibrium. As shown by Fig. 3a the response magnitudes converge at the end of the hit, which verifies there is little memory effect of the sensors.

When dealing with individual gases the sensors do not necessarily have to be selective, they just need to have different mapping functions. However, when dealing with interferences the sensor needs to be reasonably selective. Figure 3b shows




Figure 3. Sensor validation.

the responses for DMMP and DCP with diesel, methanol and hexane interferents in the background carrier air. The sensor shown is selective to DCP and DMMP but also reacts a little with diesel.

These two tests show that the sensors have the characteristics that are required by the mapping function. That is they have good selectivity towards the gases to be detected and they are very invariant to the exposure history of the gases.

To evaluate the mapping method and its performance a series of 18 test protocols were developed. They start with simple conditions of individual hits with different hit times, and increase in difficulty where the hits and the background interferents both have time varying concentration profiles. In the last few protocols the hits are gradual increases in concentration from zero to model a plume like scenario.

The results for three versions of the mapping method are shown in Fig. 4, along with the gas delivery profile and raw SAW responses. Figure 4a classifies to the analytes and interferents, Fig. 4b only classifies to the analytes, and Fig. 4c is the same as (b) except a threshold on the variance of the back calculated concentrations is in place. These results show the sensors are selective against diesel, however when methanol is delivered in the background of diesel a noticeable response is produced. Interestingly, when delivering methanol alone there is a negligible response. In general, the best results considering all the profiles as highlighted above was achieved using the method of Fig. 4c.

An important observation from all of the results of Fig. 4 is during the transient part of the response the classification is unreliable. This occurs at the start of a hit, after the end of a hit, or during a hit when there is a steep concentration gradient. The threshold helps omit parts of the transient response that occur while the sensors stabilize. Once the sensors equilibrate the results become more reliable.




Figure 4. Comparison of the mapping methods. DMMP is shown at 10x actual concentration.

From experiments of the four variance standardization methods the unmodified concentrations c) tends to bias the gas choice to those with lower concentrations. Since DMMP has the lowest concentration it often tends to be the likely candidate. The results of methods a), b) and d) are similar and give good results. It should be noted that the performance of each method may vary with the types of sensors used and the application they are used in.

CONCLUSION

This work has shown that there is a processing gain when using the mapping method for classification due to the implicit nonlinearity built into the system. The higher the nonlinearity of the sensors, the more gain that can be attained. In addition, several modifications were explored to determine the performance in the presence of interferents. The best results were found when using mapping functions of only the gases to be detected and normalizing the variances with respect to the back-calculated concentrations. The use of SRD's selective sensors in these experiments enhanced the detection of the CWAs, but the specifics of the sensors were not the focus of this paper. The use of selective sensors is one of the main properties that differentiate SRD's system from a typical e-nose, they allow targeting of specific gases that can then be quantified.

Since the motivation for this work is directed towards the development of an inexpensive handheld detector, it is primarily applicable to systems that only have very simple gas delivery subsystems. That is, without the use of valves or flow controllers, which means only the instantaneous magnitudes can be considered. Future development plans entail increasing the complexity of the gas delivery subsystem to include some valve switching components. This will allow the delivery of a short pulse from which the more information rich transient response can be utilized. The length of the pulse determines the trade-off between obtaining enough of the transient signal to make a reliable classification and maintaining the approximation that the concentration is constant during the pulse. The sensors that are exposed to the pulse need an uncontaminated baseline, which can be achieved by a number of known methods.

REFERENCES

- [1] Carmel, L., Sever, N., Lancet, D., and Harel, D. An eNose algorithm for identifying chemicals and determining their concentration. *Sensors and Actuators B* 93 (2003), 77-83.
- [2] Janshoff, A., Galla, H. J., and Steinem, C. The quartz-crystal microbalance in life science. *Angewandte Chemie-International Edition* 39 (2000), 4004-4032.
- [3] Pearce, T. C., Schiffman, S. S., Nagle, H. T., and Gardner, J. W. *Handbook of Machine Olfaction: Electronic Nose Technology*. Wiley-VCH, Germany, 2003.

B Presentation Slides

Detection of chemical warfare agents in the presence of interferents

Greg Harmer, Chuncai Yang and Brent Marquis
Sensor Research & Development
27-10-2004

Overview

- Introduction
- Classification Schemes
 - One-Step Method
 - Transient and Steady-State Features
 - Two-Step Method
 - Mapping Method
 - Comparison of Two-Step and Mapping Methods
- Dealing with Interferents
- Results
 - Reproducibility
 - Selectivity
 - Classifier - Training
 - Classifier - Variance Calculations
- Conclusion

Slide 2


Introduction

- Handheld detector
- Small, cheap, easy to operate
- Directed towards CWAs
 - DCP (dichloropentane - mustard)
 - DMMP (dimethyl methylphosphonate - nerve)
- Use of SAW sensors with polymer coatings

Slide 3

Test System

- Isolated gas delivery and sensor sampling
- A more realistic test system



Slide 4

Classification Schemes


- One Step: Input space quickly becomes large
- Two Step: divide problem into manageable parts



Slide 5

Transient & Steady-State Features


- Transient Information
 - e.g. Points: $R(t)$, Slopes: $R'(t)$, Curve Fitting
 - Normalize out conc info, can discriminate with 1 sensor



Slide 6

Transient & Steady-State Features

- Steady-State Information
 - e.g. Magnitudes when sensor is in equilibrium
 - Normalize out conc; scale and shift
 - Implicit use of sensor array



Slide 7

Mapping Method


- Find mapping functions: $r_{ij} = f_{ij}(c_j)$
 - $i = 1, \dots, m$ sensors; $j = 1, \dots, n$ gases
- Take a test sample for each sensor, r_i
 - Find each possible conc: $c_{ij} = f_{ij}^{-1}(r_i)$
- For each gas find variance of concs: $\sigma_j^2 = I[c_{ij}]$
 - This is a measure of 'closeness', choose $k = \arg \min_j \sigma_j^2$
 - With conc: $\mu_k = E[c_k]$

L. Carmel et al, An eNose algorithm for identifying chemicals and determining their concentrations, Sens Act B, 93, 77-83, 2003.

Slide 8

Comparison of Mapping and SS Methods

- What is the difference between the two?
 - SS: Uses linear scaling and shifting, assumes $r_{ij} \propto c_j$
 - Map: Implicit use of nonlinear mapping functions
- Eg. 8 sensors and 4 gases; range of concs
 - SS finger prints (dim = 8), plot first two PCs




Slide 9

Dealing with Interferents

- Training
 - Train to all gases; analytes and interferents
 - Only train to the gases of interest; analytes
- Variance calculation
 - eg. 3 sensors give A: 8, 9, 10; and B: 207, 209, 214.
 - Min variance gives A, but wrt absolute conc B might be the better choice. Try:
 - Raw-straight: $k = \arg \min_j \sigma_j$
 - Raw-normalized: $k = \arg \min_j \sigma_j / \mu_j$
 - Log-straight: $k = \arg \min_j \hat{\sigma}_j$
 - Log-Normalized: $k = \arg \min_j \hat{\sigma}_j / \hat{\mu}_j$
 - Where the hat denotes the use of log concs

Slide 10


Results – Reproducibility & Selectivity







Slide 11

Results – Selectivity

- Very selective against Hexane and Methanol
- Some sensitivity to Diesel



Slide 12



References

- [1] T. C. Pearce, S. S. Schiffman, H. T. Nagle, and J. W. Gardner. *Handbook of Machine Olfaction: Electronic Nose Technology*. Wiley-VCH, Germany, 2003.
- [2] R. Negri and S. Reich. Identification of pollutant gases and its concentrations with a multisensor array. *Sensors and Actuators B*, 75:172–178, 2001.
- [3] L. Carmel, N. Sever, D. Lancet, and D. Harel. An eNose algorithm for identifying chemicals and determining their concentration. *Sensors and Actuators B*, 93:77–83, 2003.