

To the Graduate Council:

I am submitting herewith a thesis written by Barbara Hoffheins entitled "Using Sensor Arrays and Pattern Recognition to Identify Organic Compounds." I have examined the final copy of this thesis for form and content and recommend that it be accepted in partial fulfillment of the requirements for the degree of Master of Science, with a major in Electrical Engineering.

Donald W. Bouldin

Donald W. Bouldin, Major Professor

We have read this thesis
and recommend its acceptance:

Charles F. Moore
Mohan M. Trivedi

Accepted for the Council:

C. W. Mink

Vice Provost
and Dean of the Graduate School

STATEMENT OF PERMISSION TO USE

In presenting this thesis in partial fulfillment of the requirements for a Master's degree at The University of Tennessee, Knoxville, I agree that the Library shall make it available to borrowers under rules of the Library. Brief quotations from this thesis are allowable without special permission, provided that accurate acknowledgment of the source is made.

Permission for extensive quotation from or reproduction of this thesis may be granted by my major professor, or in his absence, by the Head of Interlibrary Services when, in the opinion of either, the proposed use of the material is for scholarly purposes. Any copying or use of the material in this thesis for financial gain shall not be allowed without my written permission.

Signature

Date

Barbara A. Allen
7/25/89

**USING SENSOR ARRAYS AND PATTERN RECOGNITION
TO IDENTIFY ORGANIC COMPOUNDS**

A Thesis

Presented for the

Master of Science

Degree

The University of Tennessee, Knoxville

Barbara Hoffheins

August 1989

ACKNOWLEDGMENTS

I would like to thank all of the many people who have generously assisted me with guidance and support. In particular, I am forever indebted to my friend and colleague, Dr. Robert Lauf, who has patiently and tirelessly performed the roles of mentor and teacher, help set up and perform the experiments, and assisted me with assembly of this document. I am also grateful for the part that my thesis "coach," Dr. Richard Todd, played in encouraging me and keeping me on schedule. I wish to thank the members of my committee, Dr. Mohan Trivedi, Dr. Charles Moore and Dr. Emil Muly and Dr. Donald Bouldin who also served as committee chairman and helped with final suggestions and completion of administrative requirements. Thanks also go to Joy Kerley, Norma Hensley, Shirley Hendrix, and Patsy Thornton for help with document printing, editing, and photo preparation; Claudia Walls and Clyde Hamby for technical assistance with some of the sensor experiments; Glenn Allgood for the use of his cluster analysis programs; Dr. William Dress and Laird Williams for introduction to and use of the neural network simulator work station; Law McCabe, of Cabot Corporation, who helped initiate the study of chemical sensor arrays at Oak Ridge National Laboratory; and family, friends, and colleagues who continually encouraged me in this endeavor.

This work was supported in part by Cabot Corporation under a contract with Martin Marietta Energy Systems, Inc. Other equipment and supplies were provided by Oak Ridge National Laboratory (ORNL).

Research was performed at ORNL, which is operated by Martin Marietta Energy Systems, Inc., for the U.S. Department of Energy under Contract No. DE-AC05-84OR21400.

ABSTRACT

The recent development of chemical sensor arrays promises to solve some of the industrial, military, and domestic problems of gas detection and monitoring, but there are many problem areas to be addressed before these types of devices become readily available. The work presented here represents initial expeditions into a new, hybrid discipline for chemical analysis, which combines materials science, chemical sensing techniques, and the application of pattern recognition for automatic information extraction. Specifically, two kinds of chemical sensor array design and construction are discussed. The nature of the outputs from these sensor arrays is examined for qualities such as information content, stability, reliability, and accuracy. Several methods of pattern recognition are explored for their ability to classify sensor array information. Preliminary results indicate much promise in the use of neural networks for the analysis of mixtures, which is a vexing problem. It is found that the most appropriate pattern recognition technique depends to a large degree on the complexity of the sensing problem.

TABLE OF CONTENTS

CHAPTER	PAGE
1. INTRODUCTION	1
The Problem of Olfactory Sensing	1
Design and Construction of Gas Sensor Arrays.....	5
Previous Work in Gas Sensor Array Analysis	10
2. EXPERIMENTAL METHODS	16
An Array of Conventional Sensors: Design and Fabrication.....	16
IGAS Chip: Design and Fabrication.....	21
Selection of Compounds for Olfactory Testing	38
Testing Conditions: Discrete Sensor Array.....	39
Testing Conditions: IGAS Chip.....	40
3. RESULTS AND DISCUSSION.....	45
Organization of Data	46
Cross-sensitivity.....	47
Uniqueness	58
Effects of Concentration	64
Reproducibility	72
Dynamic Behavior.....	74

CHAPTER	PAGE
4. CONCLUSIONS.....	78
LIST OF REFERENCES	83
Supplementary References.....	87
APPENDIXES.....	89
APPENDIX A. PATTERN RECOGNITION	
METHODOLOGY	90
Classical Methods: Cluster Analysis.....	93
Neural Network Designs	95
Data Conversions for the Neural Network Programs	98
APPENDIX B. NEURAL NETWORK	
RESPONSE DATA.....	107
VITA	122

LIST OF TABLES

TABLE	PAGE
2.1. Thick-Film Master Compositions.....	30
2.2. Electrical Characteristics of Printed Tin Oxide Inks	32
2.3. Electrical Characteristics of Printed Zinc Oxide Inks.....	33
3.1. Classification of Hexane and Ethanol Signatures by the K-means Algorithm	55
3.2. Classification of Hexane and Ethanol Signatures by the Maximin-Distance Algorithm.....	56
3.3. Classification of Hexane and Ethanol Signatures by the Hamming Network	57
3.4. Classification of Water and Ethanol Signatures by the K-means Algorithm	69
3.5. Classification of Water and Ethanol Signatures by the Maximin-Distance Algorithm	70

LIST OF FIGURES

FIGURE	PAGE
1.1. Prototype sensor array containing six different gas sensing elements	6
1.2. Two examples of miniature sensor arrays, IGAS chips, which were designed and built at ORNL	8
1.3. Prototype of a thick-film sensor array developed at Hitachi and its response to four distinctive compounds.....	9
1.4. Schematic of a sensor array consisting of four electrochemical sensors operated in four different catalytic modes	11
1.5. Hitachi algorithm for comparing sample pattern to signatures of known organics.....	13
1.6. Response data from the electrochemical gas sensor array.....	15
2.1. Design of the Figaro Engineering TGS gas sensor.....	17
2.2. Characteristic responses of two Figaro TGS gas sensors versus concentration of various gases in air.....	18
2.3. Prototype sensor array containing nine different gas sensing elements.....	20
2.4. Electrode patterns developed for the IGAS chip.....	25
2.5. IGAS substrates showing printed thick-film heaters.....	26
2.6. Scanning electron micrograph of IGAS chip TOS-3	28
2.7. Scanning electron micrographs of three thick-film sensor materials, screen printed and fired in air at 925°C.....	35

FIGURE**PAGE**

2.8. Equilibrium values in air over time for the large sensor array.....	41
2.9. Test cell for evaluating steady-state response of IGAS sensors to particular gases	42
3.1. Signatures of flavor extracts are masked by that of the ethanol present as a diluent	49
3.2. The presence of water vapor has little effect on the signature of lemon extract	50
3.3. Analysis of mixtures is inherently difficult, but can be tractable if the possibilities are limited.....	52
3.4. Hexane and ethanol are more reactive at different temperatures	53
3.5. Gasohol detector circuit.....	59
3.6. Responses of sensor TOS-2 exposed to seven different gases.....	60
3.7. Responses of sensor TOS-7 to eight different gases.....	61
3.8. Comparison of heptane and hexane signatures in the large sensor array.....	62
3.9. Two aromatic compounds show interesting similarities.....	65
3.10. Signatures of naphthalene vapor from whole and crushed moth balls.....	66
3.11. Signatures of several water/ethanol mixtures	67
3.12. Signatures of ethanol over a 9-month period.....	73
3.13. Time response of the methanol signature in the large sensor array.....	76

FIGURE	PAGE
3.14. Time response of the isopropanol signature in the large sensor array.....	77
A.1. Magnitude data plots for four ethanol signatures and one ethanol signature	100
A.2. Derivative data plots for an ethanol signature	101
A.3. Training set for magnitude data (9 x 16 grid)	102
A.4. Training set for magnitude data (9 x 10 grid)	103
A.5. Training set for derivative data (8 x 10 grid).....	104
A.6. Training set for derivative data (8 x 14 grid).....	105
B.1. Classification by Hopfield network trained with five magnitude data patterns.....	110
B.2. Classification by Boltzmann network trained with five magnitude data patterns.....	111
B.3. Classification by Hamming network trained with five using magnitude data.....	112
B.4. Classification by Hopfield network trained with six magnitude data patterns.....	113
B.5. Classification by Boltzmann network trained with six magnitude data patterns.....	114
B.6. Classification by Hamming network trained with six magnitude data patterns.....	115
B.7. Classification by Hopfield network trained with five derivative data patterns	116
B.8. Classification by Boltzmann network trained with five derivative data patterns	117

FIGURE	PAGE
B.9. Classification by Hamming network trained with five derivative data patterns	118
B.10. Classification by Hopfield network trained with six derivative data patterns	119
B.11. Classification by Boltzmann network trained with six derivative data patterns	120
B.12. Classification by Hamming network trained with six derivative data patterns	121

CHAPTER 1

INTRODUCTION

The Problem of Olfactory Sensing

The human sense of smell works by the identification of scent "fingerprints." Blindfolded, we can know a rose from a lily, Juicy-Fruit™ gum from a peppermint variety, and some of us can distinguish our mother's beef stew from someone else's; we do not need an organic chemist as a lifelong interpreter to guide us through life's maze of olfactory pleasures and hazards. Few people lack the sense of smell, but the sense of smell does deteriorate with age, and there are many hazardous compounds that we cannot detect by our unaided senses. There may not be a compelling reason to equip the anosmic individual with an artificial nose, but a greater understanding of this human sense and the ability to enhance this capability or to model it, even crudely, are desirable goals. There are many real-time (rapid response time) and/or remote olfactory sensing problems in domestic, industrial, and military settings that could be solved with a detector that could rapidly identify hazardous gases or by-products of interest. We can analyze a number of compounds with mass spectrometry and gas chromatography, but the associated instruments are cumbersome for many applications, require preparation of samples, and cannot provide information in the time

required. Many applications demand portable or remotely operated devices that are specifically tailored to the problem at hand to enable rapid gas analysis.

There are many kinds of olfactory sensors, but in general they all have the disadvantage of not being able to identify explicitly an unknown gas or mixture of gases because they are not inherently selective. They may be calibrated to give an alarm or quantitative measurement only when the gaseous species is known at the outset. This problem of cross-sensitivity is particularly acute when dealing with hazardous industrial chemicals, CBW (chemical and biological warfare) agents, explosives, military fuels, and so forth, because such mixtures are chemically complex. Rapid identification of many of these compounds is desirable or even critical for making decisions affecting the safety of personnel and equipment.

One theoretical method of synthesizing selectivity or minimizing cross-sensitivity was proposed by Clifford.¹ His idea was to construct an array of sensors having differential sensitivities to different gases. Such an array would, in principle, yield signatures for different gases that would have varying degrees of uniqueness. The integrated gas analysis and sensing (IGAS) chip developed at ORNL has effectively demonstrated the creation of such a sensor array on an IC-sized (integrated circuit-sized) ceramic chip.² The signature or output from the IGAS chip consists of a histogram of resistance changes in the chip's multisensor array caused by a particular gas or gas mixture that cannot be analyzed from first principles. This miniature sensor array provides unique signatures for many types of simple organic compounds. Several

other approaches to the design and construction of gas sensor arrays have also successfully displayed selectivity to gases, but few have become commercially available.

The usefulness of the IGAS chip, or any alternative gas sensor array, could be dramatically expanded by giving it intelligence in the form of a pattern recognition engine. Various techniques have been proposed to handle the signal processing and pattern recognition required to convert the outputs of these arrays into useful information such as the identities and concentrations of particular chemical species. In general, the pattern recognition techniques have not been developed as fully as have the sensor arrays. Much of the work in signature recognition has taken a general approach in order to prove the principles. Gases used in the signature analysis work to date are often simple compounds and are not necessarily species that would ever occur together in a real-world situation. There is also the inherent difficulty of deconvolution of the signature from a mixture of two or more compounds; up to now, very little work has been done with complex mixtures. Though some of this early work has not addressed realistic problems, the ultimate goal of solving a specific problem should be kept in mind. For most applications, fortunately, we will have to contend with only a limited universe of possible or likely chemical species, which will reduce the information processing problem to a manageable level. The problem can possibly be further simplified by developing application specific sensing elements to enhance the signature differences for the specific compounds expected in a given situation. Consequently,

sensor and pattern recognition developments will most likely proceed in an iterative way.

Several materials and packaging problems have to be addressed before these gas detection instruments can be built. Typical power consumption requirements of many existing gas sensors are too high for realistic battery-operated designs. Some sensors are poisoned by compounds containing sulfur, and some may also be affected adversely by the environment. Other concerns involve the feasibility of incorporating the pattern recognition engine into application-specific integrated circuits so that the overall device is as small and efficient as possible. Other problems will arise as developers become involved in design and construction of an actual device.

The thrust of this thesis is to present the results of a study of chemical sensor array design and information analysis methods, with the goal of initiating a methodology or guide for designing useful gas detection and monitoring instruments. Included are discussions of the design and construction of two types of sensor arrays, the nature of their output, several approaches to data analysis, and evaluation of the different approaches in the context of some actual olfactory sensing problems.

Design and Construction of Gas Sensor Arrays

In the last several years, various types of chemical sensor arrays have been designed and constructed. Much of the work in the development of sensor arrays falls into two groups: arrays made of discrete commercial sensors and integrated sensors in which the array is fabricated as a single device. In either group the selectivity of the array can be achieved in several ways: (1) by the use of catalysts, (2) by the use of thermal gradients, and (3) by using filters. The arrays described below employ one or more of these methods. Arrays of the first two types were constructed for this work, and a more detailed description of each follows this section. These sensor arrays form useful test beds to explore the responses to various substances and to evaluate the usefulness of various signature recognition schemes.

One early example³ of an array constructed from several discrete commercial gas sensors is shown in Figure 1.1. This prototype consists of six commercial gas sensors of the Taguchi type* installed on a 10-cm-diam bakelite disk and mounted in an airtight chamber. The chamber is fitted with gas inlets and outlets for controlled gas flow. Resistance changes of the gas sensors are monitored by a computer data acquisition system. Each discrete sensor has its own internal heating element; power consumption for the complete six-element array is about 4 W.

*Figaro Engineering Co., Osaka, Japan.




Figure 1.1. Prototype sensor array containing six different gas sensing elements. The sensors' differential sensitivities to different gases could form the basis of an intelligent system. (Photo courtesy of M. W. Siegel, Carnegie-Mellon University.)

To reduce the size, power consumption, and manufacturing costs of a metal-oxide gas sensor array, the IGAS chip was designed and constructed using conventional thick-film technology.² The IGAS chip, Figure 1.2, achieves the effect of a multisensor array on a single substrate by creating a continuous metal-oxide film whose catalytic properties vary from place to place on the chip. Embedded electrodes allow us to map the responses of the different sensing areas. The catalytic activity is varied by creating a thermal gradient along the length of the chip, by distributing different catalysts in different areas along the surface, or by a combination of both techniques. The IGAS chip measures 2.5 x 1.0 cm and consumes approximately 2.5 W.

The gas sensor chip developed by Hitachi⁴ is a similar thick-film implementation with six discrete metal-oxide sensor areas on an alumina substrate, Figure 1.3. The chip is heated to around 400°C with a platinum heater printed on the bottom of the substrate. The response of the Hitachi sensor to four different compounds is shown in the figure.

Much development work has focused on the pressing need to reduce the power consumption of metal-oxide gas sensors. Because the high power consumption is a direct consequence of the high operating temperatures required, one obvious approach is to reduce the total size of the active sensing elements. Workers at General Motors Research⁵ have created a sensor array on an ultrathin, thermally isolated silicon membrane fabricated by chemical machining. The resulting device operates at 300°C, requiring about 150 mW.






Figure 1.2. Examples of two miniature sensor arrays, IGAS chips, which were designed and built at ORNL. Size of the chip substrate is 2.5×1.0 cm.




L-Menthol



Ammonia



Hydrogen sulfide



Trimethyl pentane




Figure 1.3. Prototype of a thick-film sensor array developed at Hitachi and its response to four distinctive compounds.⁴

In addition to the approaches cited above, sensor arrays with varying degrees of selectivity have been demonstrated that use sensing techniques other than the resistance change of a semiconducting metal oxide. Stetter and coworkers⁶ have described an instrument having four electrochemical sensors and two catalytic filaments that can be heated separately or together to give four operating modes (Figure 1.4). The combination of four sensors and four modes effectively creates a 16-sensor array. The response of their instrument, although achieved with a somewhat different physical approach, lends itself to the same type of pattern recognition scheme as other sensor arrays.

Other sensor array concepts that are being developed at present include the use of surface acoustic wave (SAW) devices with chemically selective coatings,⁷ metal-oxide semiconductor (MOS) sensors with molecular sieves,⁸ and palladium/silver (Pd/Ag) gate combinations with metal-insulator semiconductor (MIS) diodes.⁹

Previous Work in Gas Sensor Array Analysis

The conventional approach to signature recognition of sensor array data is through the application of chemometrics. Clifford describes an approach of this type in his patent "Selective Gas Detection and Measurement System."¹ A system of equations is developed, one equation per gas sensor. Clifford states that the number of gas sensors for each system must be equal to or greater than the number of gases in




Figure 1.4. Schematic of a sensor array consisting of four electrochemical sensors operated in four different catalytic modes.⁶

the system, but this is not a requirement in all chemometric systems.⁴ The response of each element in the array is measured for each selected gas to determine the constants used in the equations. In addition, all the sensors in the array must have a response to at least one of the gases included in that system. Other workers have evaluated several types of chemometric approaches to determine suitability and practicality for various classification problems,¹⁰ and have used this kind of pattern recognition to cull sensors from an array if they yield redundant information.¹¹

One of the few implemented chemometric classification systems has been done at Hitachi.⁴ The algorithm is presented schematically in Figure 1.5. A standard pattern is calculated for each of seven pattern classes selected for this system. The standard pattern has some defined range to accommodate ranges in concentration. For an unknown sample pattern, x , similarity values are calculated by multiplying the difference between the x pattern and the standard pattern classes, i , by a weighting factor, w_j , where j specifies the sensor number. The smallest similarity value, S_i , identifies the unknown sample, x , as belonging to the pattern class i . This algorithm has been implemented in an 8-bit microcomputer, which has been incorporated into an instrument with the sensor and associated power and signal processing circuitry. Analog signals from the sensor are converted into digital signals forming the input to the microcomputer, which calculates the similarity values and identifies and quantifies the gas.




Figure 1.5. Hitachi algorithm for comparing sample pattern to signatures of known organics.⁴

J. R. Stetter and coworkers⁶ devised a recognition system noteworthy for its simple, yet workable, solution for their sensor array instrument described above. As shown in Figure 1.6, for each gas, operating four sensors in each of four modes yielded 16 responses, which are arbitrarily assigned to 16 channels. "Fingerprint" patterns for each of 19 hazardous gases were made up with the 16 data channel numbers from the sensing device, listed in order of greatest response to the selected gas. Only 2 of the 19 gases tested had the same fingerprint when the three channel numbers with the greatest responses were used as the basis for the fingerprint; however, the identities of these 2 gases were determined unambiguously when the fingerprint was expanded to the five highest response channels. The significance of this kind of identification is in its obvious economy. No time-consuming chemometric manipulations are needed, and the system can be coded easily in a microcomputer. In proving the principle, Stetter et al. have designed a general-purpose instrument, but the identification principles could hold for specific applications with smaller sets of defined gases.




Figure 1.6. Response data from the electrochemical gas sensor array.⁶

CHAPTER 2

EXPERIMENTAL METHODS

An Array of Conventional Sensors: Design and Fabrication

A typical commercial gas sensor, Figure 2.1, consists of a small ceramic tube coated with tin (IV) oxide and a suitable catalyst. A heater coil maintains the temperature at 200 to 500°C so that the tin oxide can react with combustible gases. The electrical resistance of the tin oxide is measured between two electrodes at opposite ends of the tube. The entire assembly is packaged in a protective housing with an open area (covered by wire mesh) to allow gases to enter. Because the resistance change of this sensor can be large and reasonably linear over many decades of gas partial pressure, much work has been devoted to the creation of sensitive, inexpensive gas sensors and alarms.

These conventional tin oxide gas sensors can be made more sensitive to specific gases (e.g., methane or carbon monoxide) by the choice of catalyst applied to the oxide surface.¹² For example, platinum enhances the sensitivity to light hydrocarbons such as butane, whereas palladium increases the sensitivity to hydrogen and carbon monoxide. Figure 2.2 shows the characteristic responses of two different Figaro gas sensors to various gases. Unfortunately, as can be seen from that figure, any one sensor cannot distinguish between high concentrations of a less





Figure 2.1. Design of the Figaro Engineering TGS™ gas sensor.¹²




Figure 2.2. Characteristic responses of two Figaro TGSTM gas sensors versus concentration of various gases in air. For each case, the sensor's resistance is normalized to its resistance in air containing 200 ppm of CO. It can be seen that any one sensor exposed to an unknown atmosphere gives a single reading (as an example, the dashed line shown above) that does not contain enough information to determine the identity or concentration of an unknown species.²

reactive gas and low concentrations of a more reactive gas. For many applications, this cross-sensitivity or lack of selectivity, severely limits the potential of the device as an analytical tool. It has therefore been suggested¹ that by coordinating the outputs of several different gas sensors, each of which has characteristic responses for different gases, a signature could be derived that would identify any particular gas and possibly give an indication of the gas concentration. This concept was initially tested by M. W. Siegel,* who constructed the prototype sensor array shown in Figure 1.1.

To collect data for this work, a sensor array was constructed with nine discrete commercial sensors of the Taguchi type. This sensor array can be configured into either of two modes; first, nine identical sensors can be operated at different power levels (and hence, different temperatures), and second, different models of sensors can be inserted into the different positions in the array. In this way each element of the array has different response characteristics that in general will not be identical for different gases. This arrangement allows wide control over the range of sensitivities of the different sensing elements.

The test chamber is a 1-ft³ Lucite cabinet (Figure 2.3). The sensors are arranged on a shelf in the middle of the cabinet with a baffle to diffuse the flow of the incoming gas. Dc power is supplied to the heater in each sensor through individual load resistors selected to control the power applied to each sensor. The sensor array consisted of nine Taguchi-type sensors operated at power levels ranging from 30 to 100%

*The Robotics Institute, Carnegie-Mellon University (CMU).




Figure 2.3. Prototype sensor array containing nine different gas sensing elements.

of the manufacturer's recommended operating power (900 mW). Sensor 1 was operated at 100% and sensor 9 at 30%.

A Taguchi metal-oxide sensor exhibits a large change in resistance in the presence of gases to which it is sensitive. The resistance of the model used in this work (TGS 812) will decrease by 4 to 5 orders of magnitude when exposed to alcohol vapor, for example. The response to the gas occurs in a matter of seconds to minutes, but recovery to its original resistance in air can take hours. The dynamic response is not well understood, and observations demonstrated that recovery is highly dependent on the substance to which the sensor was most recently exposed. The steady-state response of the sensor to a given substance is much less sensitive to previous exposures and therefore more reproducible. Consequently, all measurements were made after allowing the sensor array to come to equilibrium with the sample gas (~5 min).

IGAS Chip: Design and Fabrication

The author was part of a team of workers at the Oak Ridge National Laboratory that developed the Integrated Gas Analysis and Sensing (IGAS) chip.¹³ This device is an example of a sensor array fabricated on a single substrate, designed to overcome the nonselectivity of a single conventional gas sensor without being substantially larger or more expensive, or requiring more power to operate. A detailed discussion of this device follows.²

A natural outgrowth of the bulky, multisensor device described in the previous section was the idea of integrating many sensors on a single substrate by varying the catalytic properties of the tin oxide layer, either by distributing different catalysts or by varying the temperature in the different reactive regions on the chip. We achieved a dramatic increase in functional density by using thick-film circuit technology to create an array of closely spaced electrodes on a refractory substrate so that the equivalent of up to 25 different sensors could be fabricated on an area of $\sim 2 \text{ cm}^2$. The resulting device plugs into a 20-pin, dual in-line IC socket. The chip represents a significant advance in the development of a "smart" chemical sensor. Although conventional thick-film technology was used to develop this new gas sensing device, some of the design requirements dictated the formulation of nontraditional thick-film inks. We chose the thick-film technique because it is an inexpensive, reliable manufacturing method which that provides much flexibility to accommodate design modifications as development progresses.

Development of this device was initiated with funding by Cabot Corporation. During the period of development (~ 1.5 years), hybrid circuit methods were exploited to produce a functional, rugged, miniature, and relatively inexpensive intelligent gas sensing element. Several designs were built and tested under steady-state conditions. In a related program, workers at CMU were to study the sensor's response to transient conditions and develop the software needed to apply the sensor to specific analytical problems. The results obtained at CMU will be published elsewhere.¹⁴

Design Specifications

The IGAS chip was designed to meet the following requirements:

1. The active area must consist of a semiconducting oxide film (e.g., SnO_2 , ZnO , Fe_2O_3) whose catalytic properties vary from place to place on the surface. The response to various gases should be rapid, reversible, and cover a large dynamic range. The baseline resistance of the oxide film in air should not be too high to conveniently measure (realistically, less than $\sim 10^6 \Omega$).
2. A large number of electrodes must be arranged so that the electrical responses of all the different areas of the sensor can be quickly mapped by addressing the electrodes in pairs.
3. The interconnection between the electrodes on the chip and the measurement equipment should be rugged, inexpensive, and compatible with standard circuit components.
4. An integral heater must be able to keep the sensor at a sufficiently high temperature for reaction with the sample gases. The heater should be available in either of two configurations, depending on whether a uniform temperature or a thermal gradient is desired.
5. The entire sensor should be stable with respect to thermal cycling and thermal aging. Specifically, mechanical incompatibilities such as poor adhesion to the substrate and spalling must be avoided. Although the present sensors have evolved with all these considerations in mind, it is important to remember that the concept is quite general and alternative designs can be imagined for specific applications.

Substrate Layout

Several electrode patterns were used, as shown in Figure 2.4. The first pattern provides for measurement at nine points along the length of the chip by addressing opposite pairs of electrodes. The second pattern provides for measurement at 17 points, giving potentially greater resolution and a more complex signature. The third pattern gives 25 measurement points and is conveniently arranged so that three strips of different sensor materials (e.g., with different catalysts applied) can be laid out along the thermal gradient. In each case, the photo masks were designed so that eight sensors are printed at once on a scored alumina substrate. The conductor material was a commercial gold-palladium (Au-Pd) frit-bonded thick-film composition, printed and fired according to standard practice.

Heater configurations are shown in Figure 2.5. The heater material was a commercial RuO_2 resistor composition. Although the heaters were not originally designed to be used together on the same substrate, the masks for the large heater can be rotated by 180° , allowing power for the large heater to be applied across pins 10 and 11 while power to the small heater is applied across pins 1 and 20. The possible benefit of a two-heater configuration is more precise thermal management and less thermal stress on the small heater.




Figure 2.4. Electrode patterns developed for the IGAS chip.²
 Top: Three different patterns as described in the text. Bottom: One complete substrate, as printed, having eight individual IGAS chip substrates (approximately actual size).




Figure 2.5. IGAS substrates showing printed thick-film heaters (black strips). The large resistor (top) heats the chip uniformly, whereas the end resistor (bottom) provides a thermal gradient along the length.² In each case, power is applied to the heater through the first pair of electrodes. (The small dots are alignment marks to aid in the printing steps.)

Sensor Material

After the electrodes and heaters are fired onto the substrates, the semiconducting oxide layer can be deposited by either of two techniques. One approach involves in situ decomposition of an anhydrous tin chloride/stearic acid mixture. The other method involves the formulation of printable inks that can be applied and fired by traditional thick-film practice.

The stearic acid method, as described by Taguchi,¹⁵ involves the following steps. Anhydrous tin (IV) chloride, a clear liquid, was mixed with powdered stearic acid and warmed on a hot plate until the stearic acid dissolved. Upon further heating, the mixture became dark brown. At that point, the mixture was painted on the substrate and fired at 700°C in air. This process had to be repeated several times to build up an adequate film thickness on the substrate. The organic component was intended to be burned off, creating desired porosity in the tin oxide layer. On top of the tin oxide layer, one or more alcohol solutions of precious metal chlorides were then painted on selected areas and fired for 5 min at 700°C in air. Films with sheet resistances in the range of 5 to 30 kΩ/square were deposited by this technique.

Figure 2.6 shows the active layer of one sensor made by this method. The tin oxide has a crenulated surface with a distribution of fine ruthenium catalyst particles. It will be shown later that this type of sensor is fast and quite responsive to many organic compounds. However, the stearic acid method had several important shortcomings. First, it was difficult to control and reproduce in the laboratory




Figure 2.6. Scanning electron micrograph of IGAS chip TOS-3.¹³ A tin oxide layer was deposited by the stearic acid/tin chloride method and fired at 700°C in air. The process was repeated approximately five more times, then the surface was doped with $\text{RuCl}_3 \cdot \text{H}_2\text{O}$ /ethanol which was painted on and fired at 700°C in air (1 cm = 10 μm).

(although the technique can, of course, be done reproducibly with specialized manufacturing equipment). The tin oxide layers were not sufficiently uniform, and the stearic acid mixture tended to liquefy readily in the initial firing and partially overlap the heater strip, occasionally rendering the first measurement point on the chip unusable. Second, the strength and adhesion of the tin oxide layer was quite poor, and its initial resistance varied greatly from point to point on the chip. Third, the distribution of catalyst was difficult to control; it is likely that the sensor layer shown in Figure 2.6 contains much more precious metal than necessary.

To overcome these difficulties, it became clear that printable compositions were needed. Our proprietary sensor inks contain three major components: a metal oxide, which is the active agent; glass frit for adhesion; and organic vehicles that burn off during firing. The catalyst can be either applied after firing, as before, or incorporated directly into the ink.

To conduct a preliminary study of the effect of ink composition and firing parameters, four master inks (SnO_2 , ZnO , Fe_2O_3 , and glass) were formulated, as shown in Table 2.1. The same binder, dispersant, and solvent were used in each, so the three batches could be blended in any desired proportion. The inks were formulated by the following procedure. Each powder was milled in isopropanol to produce a fine particle size. The oxide powders were milled in a vibratory mill in a plastic jar with zirconia medium to break up agglomerates. The glass was pulverized in a steel mortar, ball-milled 16 hours with alumina

Table 2.1. Thick-Film Master Compositions

	<u>Tin Oxide</u>	<u>Zinc Oxide</u>	<u>Iron Oxide</u>	<u>Glass</u>
Powder (g)	47	56	52.4	25
Powder analysis	99.9 SnO ₂	99.8 ZnO	99 Fe ₂ O ₃	75 SiO ₂ 12 Na ₂ O 8 CaO 5 MgO
Binder (g)	3.4	5	5	5
Solvent (g)	13.5	20	20	20
Dispersant (cm ³)	0.5	0.5	0.5	0.12

medium, and screened to -325 mesh (<44 microns). The powders were dried in air at 100°C.

To maximize the sensitivity, it is important for the oxide layer to be porous. Therefore, we formulated the inks with more organic vehicle than a normal commercial thick-film composition. To achieve a target of about 30 vol % solids in the ink, an amount of solvent equal to twice the theoretical volume of the powder was added along with a small amount of a proprietary dispersant. The mixture was initially quite dry

and was therefore blended by hand with a spatula until the dispersant was well distributed and the mixture became semifluid. The slurry was then liquefied with an ultrasonic dismembrator for 1 min at ~180 W of power.

The binder (a commercial acrylic resin) was added to the liquefied slurry in the form of small beads as supplied by the manufacturer. The mixture was stirred well, covered, and placed in an oven for several hours at 60°C to dissolve the resin completely. Each ink was thoroughly mixed again before each use.

Because of the lower volume fraction of solids, our inks were somewhat less viscous than typical commercial thick-film compositions, but each of the inks had a suitable viscosity range for screen printing onto alumina substrates. They were mixable in all proportions. A trial printing of the glass master composition fused to the substrate after 15 min at 925°C. Higher firing temperatures (950 and 975°C) gave slightly smoother films, indicating greater fusion of the glass frit, as expected.

Tin oxide and zinc oxide inks were each formulated with three different nominal glass contents (11, 20, and 33 vol %, solids basis). These inks were printed on a standard resistor test pattern that included several resistor geometries ranging from 1 x 1 mm to 1.5 x 4 mm. The test pieces were leveled, dried, and then fired at 925, 950, and 975°C. Sheet resistivities of the fired tin oxide layers were measured at 125 and 200°C (Table 2.2). Resistivities of the zinc oxide specimens were measured only at 200°C (Table 2.3). Because the test pattern contains many sizes and shapes of resistors, we believe the standard deviations of the sheet resistivities given in Tables 2.2 and 2.3 mainly reflect edge

Table 2.2. Electrical Characteristics of Printed Tin Oxide Inks

Glass Content (%) ^a	Firing Temp. (°C)	Sheet Resistivity (MΩ/square) at 125° C	Sensitivity ^b at 125°C	Sheet Resistivity (MΩ/square) at 200°C	Sensitivity ^b at 200°C
11	925	8.2 ± 4	1.5	1.7 ± 8	1.3
11	950	81 ± 56	3.6	6.5 ± 3	13
11	975	40 ± 12	2.1	13 ± 3	10
20	925	40 ± 14	1.3	5.3 ± 2.6	14
20	950	18 ± 7	1.3	1.6 ± .9	2.9
20	975	10 ± 4	1	1.7 ± .6	3.3
33	925	12 ± 5.6	1	2.4 ± 1.5	3.1
33	950	27 ± 7	1	1.2 ± .5	3.1
33	975	27 ± 7	1	1.6 ± 1	2.9

^aNominal vol % of glass master ink, balance tin oxide master ink.

^bResistivity in air divided by resistivity in acetone vapor, with no catalyst applied.

Table 2.3. Electrical Characteristics of Printed Zinc Oxide Inks

Glass Content % ^a	Firing Temperature (°C)	Sheet Resistivity (MΩ/square) at 200°C	Sensitivity ^b at 200°C
11	925	36 ± 10	4.0
11	950	0.8 ± .4	2.4
11	975	3.1 ± .8	1.8
20	925	93 ± 70	12
20	950	1.0 ± .15	3
20	975	2.5 ± 2	1.3
33	925	100 ± 49	2.0
33	950	1.7 ± .7	1.5
33	975	19 ± 7	1.7

^aNominal vol % of glass master ink, balance zinc oxide master ink.

^bResistivity in air divided by resistivity in acetone vapor, with no catalyst applied.

effects (such as interactions with the gold electrodes), rather than intrinsic variations in the printed films themselves.¹⁶ The samples were exposed to acetone vapor as a crude measure of their relative sensitivities. Resistivities changed rapidly and reversibly by as much as a factor of 14. Initial response to the gas took a few seconds and recovery took about a minute.

Figure 2.7 shows typical examples of our sensor inks printed and fired according to standard practice. Note the improved uniformity as compared to the tin oxide film deposited by the stearic acid method. Note also the fine porosity and correspondingly high surface areas of the printed compositions.

One can see from Tables 2.2 and 2.3 that the sheet resistivities of the printed sensors were quite high. The resistivities and sensitivities of the tin oxide samples were strongly temperature dependent, as expected. The tin oxide compositions were generally more responsive to acetone vapor than were the zinc oxide compositions, suggesting that SnO_2 has greater intrinsic catalytic properties than ZnO with respect to the oxidation of acetone.

Because of the high resistivities of these compositions, the soda-lime glass frit was later replaced by a more conductive vanadium oxide-based frit; however, sensors made with these newly developed inks were not studied extensively so that pattern-recognition work could focus on the large array with its greater flexibility to control individual sensing elements.






Figure 2.7. Scanning electron micrographs of three thick-film sensor materials, screen printed and fired in air at 925°C.² Top: SnO₂ + 11% glass; center; ZnO + 11% glass; bottom: Fe₂O₃ + 11% glass (1 cm = 10 μm).

Interconnections

Connector pins for the IGAS chip are a commercial product, duPont Connector Systems* 75503-003 preplated brass stock, with a tin-lead (Sn-Pb) solder coating. The connectors are supplied on a continuous roll and are designed to clamp onto the edges of the substrate, providing a mechanical bond to supplement the adhesive or solder joint. Initially the pins were joined to the thick-film electrodes with silver-filled polymer adhesives (Ablebond† 943-1 and 71-1). We found that the strength of these adhesives was inadequate to withstand repeated handling of the sensors, and there was evidence of deterioration at the high operating temperatures. As a result, vapor phase soldering was later chosen as the preferred method of attachment.

Sensor Fabrication

The entire fabrication process is summarized as follows:

1. Substrates are prepared for printing by boiling in a cleaning solution‡ for 5 to 10 min.
2. The substrate is placed in the screen printer and aligned

*Formerly Berg Electronics.

†Ablestik Laboratories, Gardena, California

‡Crystal Clean, Aremco Products, Ossining, New York.

with the mask, which is formed on a 200-mesh stainless steel screen. Standoff is adjusted to 1.0 mm.

3. Electrodes are printed with duPont 9910 (Au) or 9985 (Pt/Au) thick-film conductor compositions at a squeegee speed of ~3 cm/s. They are allowed to level for 10 min at room temperature and then dry on a hot plate at 100 to 150°C.

4. Firing is done in air using a belt furnace. The speed is set to give ~10 minutes at a peak temperature of 925°C. Total time in the furnace is ~1 hour.

5. Heaters are printed with duPont 1711 resistor composition, leveled, dried, and fired as in step 4.

6. The liquid tin chloride/stearic acid mixture is brushed on and fired in air at 700°C for 5 min. This process is repeated several times to build up an adequate film thickness. Alcohol solutions of precious metal chloride are then brushed on, dried, and fired in air at 700°C (the palladium mixture was fired at 800°C). Alternatively, our proprietary inks (containing the catalyst as a soluble organometallic compound) are printed on and fired as in step 4.

7. Electrode pads are cleaned with a small fiberglass brush.

8. Leads are attached with conductive adhesive (silver-filled polyimide) and baked for 30 min at 150°C followed by 30 min at 275°C. Alternatively, leads are soldered using 63/37 Sn-Pb solder cream, reflowed by a fluorocarbon vapor-phase soldering unit (temperature of the vapor is 216°C).

Selection of Compounds for Olfactory Testing

Several compounds and mixtures of interest were tested, including water/alcohol; hexane/alcohol and gasoline mixtures; essential oils, flavorings, and foods; moth balls (naphthalene); and malathion, as an analog of chemical warfare agents. The mixtures chosen are representative of a wide variety of substances encountered in chemical process industries and particularly in food and fragrance production.¹⁷ Specifically, we examined mixtures of water and ethanol because of their significance to the distilling and pharmaceutical industries. Hexane/ethanol mixtures served as a paradigm for automotive fuels. Fresh whole milk and spoiled milk were used to study the possible detection of rancidity. Flavorings and essential oils were studied to determine the behavior of complex phenols and esters and the complicating effect of large quantities of alcohol as a diluent. Also, because water is such a ubiquitous substance and metal-oxide sensors are known to be influenced by humidity, we intentionally ran some measurements with water-saturated air as the carrier gas.

Of interest was discovering the characteristic effects of the many substances, how sensitive the array would be for different concentrations of the water/alcohol and hexane/alcohol mixtures, and whether the array was sensitive to the decomposition products of dairy foods and fruits. Substances avoided were those likely to contain mercaptans or some other form of sulphur, such as onions or garlic, which contain allyl disulfide.¹⁸ It is well known that sulfur is a poison to noble metal

catalysts, which are used with the Taguchi sensors and the IGAS chip and it is suspected that this characteristic makes them inherently unsuitable for environments in which these compounds are likely to occur.

All the examples to be discussed present some of the characteristics of sensor array data that must be examined in order to determine the feasibility of applying pattern recognition techniques for automatic identification of gaseous compounds in a system. It is therefore desirable to characterize the ranges of response and individuality of the signatures, to characterize the effect of water or other diluents, to define the limits of reproducibility, and to understand the complicating effects of reactive solvents and poisoning agents.

Testing Conditions: Discrete Sensor Array

Most of the substances measured are volatile liquids at room temperature. Vapors were introduced into the test chamber by passing air through a flask containing the liquid. Air flow was nominally 340 l/h. Because of the baffle and because the system was allowed to come to equilibrium, the response was insensitive to flow rate of the sample gas. In the case of particularly volatile substances like alcohols, the sensors were able to equilibrate before appreciable evaporative cooling of the vapor generator could take place (~3 min). This characteristic is important because temperature changes will affect the vapor pressure and hence the concentration of the species in the gas stream.

After most exposures, the system was allowed to equilibrate with air. For most tests this took several hours, but recovery from certain substances was incomplete after 16 hours or more. In general, tests were repeated to verify the reproducibility of the signatures. (In some practical applications, a sensor would not normally be allowed to come to equilibrium in air between measurements, and it is assumed that the sensor would still be sensitive to changes in the concentration of gases in a system.)

For the present work, we used the first configuration of the sensor array (identical sensors operated at different power levels). We used the TGS 812 general-purpose gas sensor, which is sensitive to a wide variety of organics and other reactive gases such as ammonia and hydrogen cyanide. An initial bake-out of the sensors was done with all of the sensors at full power (900 mW) for several days (Figure 2.8). This is recommended by the manufacturer and is related to baking off accumulated moisture and other contaminants.¹⁹

Testing Conditions: IGAS Chip

The test chamber for evaluating steady-state response of IGAS sensor chips to specific gases consists of a glass cell mounted to a glass-ceramic base (Figure 2.9). The base supports a 20-pin chip receptacle, which provides two power feed-throughs for the chip's heater and 18 connections to a terminal box to measure the electrical resistance




Figure 2.8. Equilibrium values in air over time for the large sensor array. Sensors were all operating at full power (900 mW) in air. Note variability among nominally identical sensors.



Figure 2.9. Test cell for evaluating steady-state response of IGAS sensors to particular gases.¹³ The electrical box, to measure the resistance between any two electrodes on the chip, was intentionally designed to be compatible with any measurement system used on the large array. Power is supplied to the heater through the coaxial cable (lower right) from a controlled dc power supply. Gases from a vapor generator enter through the tube at the top of the flask.

between any two electrodes on the chip. Note that the terminal box is identical for both the IGAS test cell and for the large array (Figure 2.3). Power is supplied to the heater from a controlled dc power supply.

The top of the glass cell is fitted with a stopper that secures the gas intake and exhaust tubes. The intake tube is bent to the side of the chamber to avoid a direct flow of gas across the chip, which would tend to cool it. This tube opens about 1 cm above the base of the test chamber and provides a slow, diffused flow of the test atmosphere around the chip to the gas exhaust near the top of the chamber.

The intake air, which constantly feeds through the chamber during testing, passes through a dry, empty flask to establish the baseline resistance of the different areas of the chip in air. This flask is removed and replaced with one containing a pure organic liquid at room temperature. The air picks up vapor and carries it to the test chamber, exposing the IGAS chip to the test gas.

To begin testing, the IGAS chip is mounted in the test chamber under an air flow rate of 1.1 l/h. A potential of 15 to 20 V is applied to the heater on the chip, and it is left overnight to reach thermal equilibrium and come to chemical equilibrium with the air. The electrical resistance of the tin oxide between opposite pairs of electrodes is then measured with a digital multimeter by addressing the appropriate points in the terminal box.

Introduction of the test gas to the IGAS chip is accomplished by placing a flask of pure organic liquid in line with the air intake, thereby causing the air to pick up the test gas vapor and transport it to the chip without changing the flow rate of the air into the test chamber. The

IGAS chip is left under flowing air with gas vapor for 5 min to allow it to come to equilibrium with the test atmosphere. The resistance between opposite electrodes along the length of the chip is then measured.

The flask containing the pure organic liquid used for the first test is removed, and an empty dry flask is inserted in its place. The IGAS chip is left under flowing air at the same constant flow rate, and the power to the heater is left on throughout the test procedure. The chip is left for ~30 minutes to return to equilibrium with air.

The resistances at the various points on the IGAS chip are measured after this 30 min recovery period before the next sample gas is introduced and the test procedure repeated.

The plot of the resistance in gas divided by the resistance in air measured between opposite points along the length of the chip is recorded for each gas tested. These "signatures" are compared to show how well each IGAS chip can differentiate between various similar gases (e.g., alcohols) and between gases of greater diversity (e.g., alcohols vs. ketones).

No attempt was made to determine the absolute gas concentrations, but it is likely that relative concentrations reflect the vapor pressure of each species at room temperature. This should be kept in mind particularly when dealing with mixtures. (For instance, the vapor over a water-ethanol mixture will be richer in ethanol than the liquid phase.) The reason for using this method of sample delivery was to better approximate head space analysis in which the unknown sample is present in bulk.

CHAPTER 3

RESULTS AND DISCUSSION

The bulk of experimental results was the generation of many signatures using a wide variety of substances and several sensor arrays. They were tested under various conditions, including complicating effects of mixtures. The signatures were evaluated in terms of their uniqueness and reproducibility, which can be verified by inspection in some cases. Various pattern recognition techniques were then applied to signatures that superficially met these criteria. The goals of this task were: (1) to evaluate these methods for their ability to satisfactorily distinguish different patterns and (2) to gain a better understanding of just how "different" the patterns need to be. A related issue is how different (or noisy) the same pattern can be and still be recognized by the algorithm. It became clear that different algorithms are more appropriate for different problems and that the complexity of the problem can be reduced if the universe of possible outcomes is limited.

A study of the current literature dealing with sensor array development²⁰ indicates what one expects to find, namely, that different sensors have different responses to a particular gas. Algorithms used by Kaneyasu et al.⁴ and Stetter et al.⁶ successfully distinguished certain pure substances, but no attempt was reported by these workers to analyze mixtures. Also, in these studies, few groups of similar compounds were

analyzed, so it is difficult to know how these algorithms might perform if they were expected to distinguish among several ketones, for example. These reports also did not discuss day-to-day variability or reproducibility of the signatures. Specific test cases are discussed below to illustrate how these complications arise and their impact on the pattern recognition problem.

Organization of Data

Sensor array data were plotted in two ways. The discrete sensor array data were usually plotted as a set of resistance measurements, one for each sensor. The resistances are plotted along the y axis and the sensor number or position is plotted along the x axis. (When looking at the graphs, it is important to remember that a low resistance corresponds with a strong sensor response.) Typically, the resistance measurement was taken after the system came to equilibrium (~3 to 5 minutes). For reference, ranges of "air" values of sensors before the gas was introduced to the system are shown on some of the plots. IGAS chip data were usually plotted as a set of ratios of resistance in gas to resistance in air. Because the response of the discrete sensors in the large array covered as much as four decades of resistance, it was appropriate to plot the data logarithmically. Because of its lower operating temperature, the response of the IGAS chip was less dramatic, and signature differences were better seen in a linear plot. (In the design

of an application-specific instrument, the particular sensor responses will largely determine the front-end signal processing scheme of a detector circuit.) Examples of actual signatures will be examined in the following sections to illustrate the important attributes of sensor array data and evaluate the applicability of particular algorithms. The following algorithms were examined: K-means and maximin-distance cluster analysis algorithms, and the Hamming, Hopfield, and Boltzmann machine paradigms on the ANSimTM* artificial neural simulator. A discussion of these algorithms and data conversion methods is located in Appendix A. A complete tabulation of the performance of each neural network is presented in Appendix B in a form designed to show at a glance the relative successes of individual algorithms.

Cross-Sensitivity

Probably the most important questions to be answered about a sensor array are whether it is selective and whether it can distinguish substances of interest in a mixture. Many sensor arrays, described in the literature, have unique signatures for various compounds, but most have no information about their abilities to identify compounds in mixtures.

*ANSimTM is a registered trademark of Science Applications International Corporation (SAIC), San Diego, California.

Both sensor arrays used in this project were subjected to mixtures, and both sensor arrays had some successes with mixtures. The large array of discrete sensors was tested more extensively with mixtures, so more of the results described below are from those tests.

An extreme example of cross-sensitivity or interference can be seen in the series of tests in which the large sensor array was exposed to flavor extracts (Figure 3.1). Initially, most of the flavor extracts tested contained a relatively large amount of alcohol (36 to 88% in the samples examined). The alcohol matrix dominates the signatures; consequently, the signatures of three commercially prepared flavor extracts (almond, peppermint, and lemon) are very similar to that of ethanol, although there are some secondary differences that may indicate a degree of uniqueness in the signature of each mixture. Given the fact that the TGS 812 sensor was intentionally designed to be highly sensitive to alcohol, it might be possible to magnify the differences with a different combination of sensors or operating conditions. This situation serves to illustrate that no one sensor array is sufficiently general that it can resolve all identification problems.

At the other end of the cross-sensitivity spectrum is the relatively benign effect of water when a much more reactive substance such as alcohol is present (Figure 3.2). In this test the discrete sensor array was exposed to lemon extract (88% alcohol) carried in dry air and the same extract carried in moist air (created by passing the air stream through a flask containing water and then through a flask containing lemon extract). For all practical purposes the resulting signatures are the same, indicating that the presence of water vapor has little effect on the




Figure 3.1. Signatures of flavor extracts are masked by that of the ethanol present as a diluent.




Figure 3.2. The presence of water vapor has little effect on the signature of lemon extract.

signature if one or more of the constituents of the mixture is substantially more reactive than water. This result is encouraging, given the ubiquity and variability of water vapor in most environments.

Between the two extremes of cross-sensitivity presented above are hexane/ethanol and gasoline/ethanol mixtures. In our previous work we discovered, surprisingly, that relatively small amounts of ethanol or methanol significantly affected the signature of gasoline, (Figure 3.3). The effect was relatively insensitive to the amount of alcohol present but was different for each of the two alcohols.

In Figure 3.4, a mixture of 2% ethanol in hexane was tested to compare with pure ethanol and pure hexane. The resulting signature shows a resemblance to hexane at the higher temperature sensors at which hexane gives the greater response and a resemblance to ethanol at the lower temperature sensors at which ethanol gives the greater response. This response is not unexpected; the response of any one sensor is the sum of its responses to the individual components, and, because these responses are exponential, they are dominated by the most reactive component at that temperature. Notice that the response is nonlinear; at the lower temperatures (sensors 4 to 9) the response to hexane/2% ethanol is very close to that of pure ethanol.

This interesting behavior suggests the possibility that enough information is present in the combined signature to enable a pattern recognition scheme to be developed that could identify both the hexane matrix and the small amount of alcohol by their individual dominance of different parts of the signature.




Figure 3.3. Analysis of mixtures is inherently difficult, but can be tractable if the possibilities are limited.




Figure 3.4. Hexane and ethanol are more reactive at different temperatures. The signature of hexane/2% ethanol shows some characteristics of each.

Tests were performed on a set of hexane, ethanol, and hexane/2% ethanol signatures using cluster analysis and neural network algorithms. Both cluster analysis programs were able to sort ethanol and hexane signatures, but they were not able to form separate clusters for the group of hexane/2% ethanol signatures (Tables 3.1 and 3.2). (Note: These data were taken at different times, and control of the actual mixture was not performed, so it is possible that at longer times the ethanol had selectively evaporated, leaving a residual mixture that would behave more like pure hexane.) In the main, hexane/2% ethanol signatures taken after a longer time tended to be classed with the hexanes, and signatures taken after only 1 min were classed with the ethanols. The most successful neural algorithm for this problem was the Hamming network using derivative data sets (Table 3.3). A network trained with representative signatures from three alcohols (ethanol, methanol, and isopropanol) and two alkanes (hexane and heptane) consistently recognized all ethanol signatures and all hexane signatures. Of five hexane/2% ethanol signatures presented to this network, two were classed with ethanol, two with hexane, and one was classed as both hexane and ethanol. The Hamming network, trained on magnitude data, also easily classed pure ethanols and hexanes but was confused by the hexane/2% ethanol signatures. The Boltzmann network, trained with derivative data, performed reasonably well with classification of pure hexane and ethanol signatures, but also confused the hexane/2% ethanol signatures with other compounds.

It must be kept in mind that in a real sensing problem, if the possibilities are quite limited and known at the outset, the problem might

Table 3.1. Classification of Hexane and Ethanol Signatures^a by the K-means Algorithm

Composition of Clusters			
	Ethanol	Hexane	Hexane/ 2% Ethanol
No. of Clusters = 2			
Cluster 1		11	1
2	15	2	4
No. of Clusters = 3			
Cluster 1		4	
2		9	1
3	15		4
No. of Clusters = 4			
Cluster 1		4	
2		7	1
3		2	1
4	15		3
No. of Clusters = 5			
Cluster 1		7	1
2	4		1
3		2	1
4	11		2
5		4	

^aTotal number of ethanol signatures was 15; total number of hexane signatures was 13; total number of hexane/2% ethanol signatures was 5.

Table 3.2. Classification of Hexane and Ethanol Signatures^a by the Maximin-Distance Algorithm

Composition of Clusters			
		Ethanol	Hexane/ 2% Ethanol
Threshold = 0.50			
Cluster	1	15	9
	2		4
Threshold = 0.33			
Cluster	1	15	4
	2		1
	3		3
	4		8
	5		1
Threshold = 0.25			
Cluster	1	15	3
	2		1
	3		3
	4		7
	5		1
	6		1

^aTotal number of ethanol signatures was 15; total number of hexane signatures was 13; total number of hexane/2% ethanol signatures was 5.

Table 3.3. Classification of Hexane and Ethanol Signatures^a by the Hamming Network

Network Assignment			
	Ethanol	Hexane	Hexane/ 2% Ethanol
Input Pattern:			
Ethanol	15		
Hexane		12	
Hexane/ 2% Ethanol	2	2	1


^aThe network was trained with derivative data sets of three different alcohols, and two different normal hydrocarbon signatures.

not require pattern recognition at all. Consider again the signatures of gasoline and alcohol mixtures (Figure 3.3). If a sample is taken at the gas pump, it must be one of these three substances (it cannot possibly be water, ammonia, etc.). For this problem, a nine-element array is not necessary; if the curves are reproducible, any two sensors ought to be enough. Sorting the possible responses into at least two or three categories can be done with a simple analog or digital circuit (Figure 3.5). (Sorting into more than two categories will depend on the customer's goal: One might want to avoid any fuel containing alcohol, whereas another might tolerate ethanol but not methanol in the gasoline.)

Uniqueness

Implicit in any pattern recognition scheme is the assumption that signatures are unique. But because this is a catalytic oxidation process,¹⁵ it is reasonable to expect that families of related organic compounds might have similar signatures. This proposition was tested with the IGAS chip in both the thermal gradient (Figure 3.6) and the distributed catalyst (Figure 3.7) designs. Not surprisingly, similar compounds do indeed have similar signatures. This phenomenon can also be observed in the large sensor array with hexane and heptane (Figure 3.8).

The neural network algorithms often confused signatures of chemically-similar substances, such as heptane and heptane or ethanol, methanol, and isopropanol, but almost always at least classified all



R_c : Resistor calibrated to balance the bridge when the detector "sniffs" pure gasoline.

R_{Sensor} : Resistance of the sensor.

R_b : Bridge resistors.

Red LED lights when alcohol is present in the gasoline mixture.

Green LED lights during idle condition before the gasoline is present.

Figure 3.5. Gasohol detector circuit. This circuit is in its most elemental form. It might be necessary to amplify the bridge output to light the LEDs; likewise, it might be necessary to add one or more sensor circuits for redundancy and to compensate for concentration.




Figure 3.6. Responses of sensor TOS-2 exposed to seven different gases.² Heater power was kept constant at about 2.5 W. Similar compounds have similar signatures, suggesting that the behavior is dominated by the most active functional group on the molecule and is relatively insensitive to chain length.




Figure 3.7. Responses of sensor TOS-7 to eight different gases.² The active area was uniformly heated in this case, so the different signatures arise from differences in the way each gas interacts with the various catalysts. Notice that the effect of ammonia on the sensor is strongest in the ruthenium area of the chip, instead of the platinum area (as with the other gases). This is the sort of information to look for when designing a sensor array for a particular application. Heater power was 3.9 W.




Figure 3.8. Comparison of heptane and hexane signatures in the large sensor array. Readings for each were taken five minutes after the start of the gas stream.

members of one group together (refer to Appendix B). The Hamming network always correctly identified the normal hydrocarbons as normal hydrocarbons, and in most cases successfully distinguished hexane from heptane. It was slightly more successful identifying individual alcohols. The Boltzmann network usually grouped the alcohols successfully as alcohols, but it had difficulty with the normal hydrocarbons, although with water as part of a derivative training set, normal hydrocarbon group classification was reasonably successful. The Hopfield network performed poorly with one particular data set, but could recognize the alcohol group with the magnitude data sets.

When training vectors included water signatures, some unhappy consequences were observed in the Boltzmann and Hopfield networks trained with magnitude data. Water signatures, though dramatically different from those of the other substances, were not identified as successfully. In the same tests, the Boltzmann network failed miserably with the normal hydrocarbons and could not even identify all of its own training vectors. One explanation for the poor performance is that the addition of the water vector, with its much greater magnitude, looked greatly different from the collective alcohol and normal hydrocarbon signatures, which share many of the same bits; consequently, differences between their signatures were reduced in significance to the network. This factor suggests the importance of training the network with equally "different" patterns as much as possible.

Although uniqueness of signatures is considered desirable, there might be situations in which the signature of an unknown substance has some similarities to a known substance, which could provide some

chemical information. The signatures of both oil of clove and oil of wintergreen show an interesting increase in sensor array response at the middle temperature zones (Figure 3.9). Perhaps this response is due to the benzene ring present in both compounds. In fact, the signature of naphthalene, another aromatic compound, has interesting similarities in the intermediate temperature range (sensors 5 to 7) (Figure 3.10). This suggests the possibility that part of the signature is more important or contains more chemical information in some cases. One might, therefore, consider analyzing only certain parts of the signatures or possibly weighting some points more heavily. Although these kinds of sensors would not have the distinguishing power of a mass spectrometer or a gas chromatograph, there are applications in which the ability to sense a trend, such as that described above with the phenolic compounds, would be very useful.

Effects of Concentration

In many chemical engineering situations, it is important to know quantitatively the concentration of one or more components when the identities of all components are known. It has already been shown that the signatures of hexane/ethanol mixtures do not follow a rule-of-mixing behavior. To explore further the behavior of two component mixtures over a wider concentration range, water/alcohol solutions (0, 10, 20, 50, and 100% ethanol) were tested (Figure 3.11), in the large sensor array.




Figure 3.9. Two aromatic compounds show interesting similarities. An indication of the reproducibility of the signatures is shown by the two solid lines for oil of clove representing two measurements separated by a week during which signatures of many other substances were collected.




Figure 3.10. Signatures of naphthalene vapor from whole and crushed moth balls. The effect of concentration is to raise or lower the signature without affecting its shape. Note the similarity of the middle of the signature (a steep, monotonic decrease in resistance of sensors 5 through 7) to the signatures of the other aromatics shown in Figure 3.9.




Figure 3.11. Signatures of several water/ethanol mixtures. The signature of 50% ethanol (not shown) is virtually identical to that of pure ethanol. Note how the ethanol signature dominates over that of water even at low concentrations.

It was observed that the signature of each mixture had virtually the same shape, but the absolute value of the responses was nonlinear with composition. The 50% and the 100% ethanol gave virtually identical responses; concentrations of ethanol between 50% and 100% would therefore be difficult to distinguish. The response to the 10% ethanol mixture is about one-fourth of the response to pure ethanol (relative to the response to water) and the response to 20% ethanol is about one-half of the response to pure ethanol.

Generally, the cluster analysis algorithms gave expected results for the data; that is, the more variable and distant water signatures formed groups opposed to one or two groups of signatures containing any amount of ethanol (Tables 3.4 and 3.5). Without selectively tuning the threshold value of a particular grouping by the maximin-distance algorithm, the nonlinear response of the array poses a problem for distinguishing smaller concentration differences. The K-means algorithm with five clusters performs better, as shown by the separation of pure alcohols (and the high-concentration ethanol/water signature) from the lower-alcohol signatures. One of the 10% ethanol/water patterns is grouped with pure ethanol, but the reading for this particular signature was taken at an earlier time than the other 10% readings and is considered a transient response. Better results might be achieved by converting the data to logarithms before processing, although this type of sensor will always exhibit a more variable and less sensitive response to high concentrations of water and special techniques might have to be employed to deal accurately with mixtures at this end of the spectrum. This situation also leads to the consideration of replacing some of the

Table 3.4. Classification of Water and Ethanol Signatures^a by the K-Means Algorithm

Composition of Clusters						
		100% Ethanol	50% Ethanol	20% Ethanol	10% Ethanol	100% Water
No. of Clusters = 2						
Cluster	1	2	1	1	4	
	2					5
No. of Clusters = 3						
Cluster	1	2	1	1	4	
	2					3
	3					2
No. of Clusters = 4						
Cluster	1	2	1		1	
	2			1	3	
	3					2
	4					3
No. of Clusters = 5						
Cluster	1	1	1			
	2	1			1	
	3					2
	4					3
	5			1	4	

^aInput patterns consisted of 2 each 100% ethanol, 1 each 50% ethanol, 1 each 20% ethanol, 5 each 10% ethanol, 5 each 100% water.

Table 3.5. Classification of Water and Ethanol Signatures^a by the Maximin-Distance Algorithm

Composition of Clusters						
		100% Ethanol	50% Ethanol	20% Ethanol	10% Ethanol	100% Water
Threshold = 0.50						
Cluster	1	2	1	1	5	2
	2					3
Threshold = 0.25						
Cluster	1	2	1	1	5	
	2					1
	3					3
	4					1
Threshold = 0.10						
Cluster	1	2	1	1	5	
	2					1
	3					1
	4					1
	5					1
	6					1

^aInput patterns consisted of 2 each 100% ethanol, 1 each 50% ethanol, 1 each 20% ethanol, 5 each 10% ethanol, 5 each 100% water.

sensors in the array for others, that handle water vapor measurements with more consistency.

Little attention was given to applying the neural network algorithms to this kind of data because it was felt that the curve shapes of mixtures between the pure water and pure ethanol curves were too similar and could not provide useful information. Success was declared if the signature was classed as water when the input pattern was more water than alcohol and classed as alcohol when the input pattern was 50% alcohol or greater. In this approach, the networks trained with derivative data sets were more successful than the ones trained with the magnitude data. This fact might indicate that the neural networks can successfully distinguish subtle changes in slope caused by changes in mixture concentration.

A suggested approach for increasing the success of using neural networks for concentration analysis is to train the networks with patterns for specific concentrations and to use a training vector, which is as large as practical, to increase resolution of individual patterns. A characteristic of the Boltzmann and Hopfield networks as implemented on the ANSim™ program, namely, the intermediate "gray" scale candidates that appear as next likely choices, might be exploited to interpolate the concentration value.

Reproducibility

Another implicit assumption is that the sensor array is stable enough over the long term so that once a known signature is learned, the array can produce that signature repeatedly at any time in the future when it is exposed to the same substance. Signatures of ethanol collected sporadically over 9 months' time (Figure 3.12), seem at first glance to have significant differences. However, examination of this series of patterns with some of the derivative data networks yielded surprising results. When the signatures were converted to a more useful form (Appendix A), the network consistently recognized them as ethanol after training with randomly selected signatures.

Reproducibility can also be looked at with regard to the pattern recognition behavior. In some of the Boltzmann and the Hopfield networks, the final response does not converge to the same answer every time. This happened most often within groups of like compounds; for example, sometimes when confronted with a hexane signature, the Hopfield network would converge to hexane, other times to heptane, and convergence to both at once also occurred. This is a known feature of neural network behavior,²¹ especially when the network becomes saturated or the training vectors have too many similarities. This might be an indication that the training vectors need to be redesigned or the sensor array or that the front-end signal processor should be altered to give more distinct signatures for the different substances.




Figure 3.12. Signatures of ethanol over a 9-month period. Surprisingly, most of these signatures were identified correctly as ethanol by one of the Boltzmann networks and one of the Hopfield networks trained with derivative data, and one of the Hamming networks trained with magnitude data.

Dynamic Behavior

All signatures presented so far represent steady-state conditions in which the measurement is made after the sensor array has been exposed to the test atmosphere long enough for the signature to become reasonably stable. The dynamic response to a changing atmosphere is generally not well understood but is nonetheless, quite important for real applications. First, one would like to know how long it takes the array to form a recognizable signature upon exposure to a sample. Second, Siegel³ suggested that dynamic measurements might supplement the steady-state signatures of compounds that are similar structurally and allow them to be distinguished. A limited number of dynamic measurements were done primarily to explore the first of these two issues. Training vectors consisting of patterns assumed to be steady-state responses were tested against series of readings taken at prescribed times. The Hamming network consistently performed well for both derivative data and magnitude data when only alcohols and normal hydrocarbons formed the training vectors. The Boltzmann network, trained with the five hydrocarbon training vectors, did fairly well using the derivative data, and it also performed well with magnitude data when trained with the additional water signature. Usually, the true identity of the signature was determined by the fourth or fifth minute that the air stream carrying the gas had been flowing through the sensor array.

With the Hamming network, there was very little confusion during the entire time period.

The potential usefulness of Siegel's proposition is shown by the dynamic responses of the large array to methanol and isopropanol (Figures 3.13 and 3.14). Although the signatures of these two compounds have many similarities and are misidentified by some of the neural networks, their dynamic responses are much different. Note that for isopropanol the resistance of sensor 4 decreases with time during the first 4 min, whereas for methanol, the resistance of sensor 7 increases over the entire 20 minutes of exposure. Obviously, the dynamic behavior is a potential source of additional information to enhance the uniqueness of the signatures.




Figure 3.13. Time response of the methanol signature in the large sensor array.




Figure 3.14. Time response of the isopropanol signature in the large sensor array.

CHAPTER 4

CONCLUSIONS

It was found that methodologies from several disciplines can be combined to form a new field of development devoted to solving some important and potentially wide-ranging gas sensing problems. The determination of concentrations and, more importantly, the analysis of mixtures by using metal-oxide sensor arrays coupled with pattern recognition schemes were shown to be feasible. The conclusions in this study can serve as a basis for the development of a selective chemical detection instrument that does not pretend to replace sensitive and general-purpose instruments such as gas chromatographs and mass-spectrometers, but discovers its niche in specific problem areas where size, portability, field operation, low power, and/or rapid response time are desirable characteristics.

Specific conclusions can be summarized as follows:

1. An array of different chemical sensors can provide significantly more detailed information about its environment than can a single sensor.
2. A simple experimental sensor array can be constructed with discrete conventional sensors operated at different temperatures. Because different compounds react on catalytic surfaces with different

characteristic activation energies, the collection of responses for a given substance can form a recognizable signature.

3. A miniature sensor array can be constructed by hybrid circuit or other techniques to minimize the overall size and power consumption. This concept was demonstrated in a small device with a continuous gas sensing film with varying surface catalytic properties. By measuring the response at different points along the film with an array of electrodes, one can obtain electrical signatures that are characteristic of particular gases. Variations in catalytic properties from point to point on the miniature sensor can be achieved by establishing a thermal gradient, by distributing several different catalysts, or both.

4. The gas sensitive film can be deposited by several methods, including screen printing of fritted oxide mixtures or the decomposition of tin chloride. The relative effectiveness of the catalysts with respect to oxidation of organic compounds studied here, in decreasing order of effectiveness are: platinum, rhodium > ruthenium > palladium > iridium, osmium.

5. Much work remains to be done in the development and optimization of the sensor arrays themselves, particularly with regard to power consumption, thermal management, and materials stability.

6. No matter what technique is used to physically manufacture the sensor array, to be useful the output of the array must be deconvoluted by matching the pattern of responses to those produced by known substances.

7. In sensors examined thus far, the resulting signatures were similar within the functional group; that is, the ketones exhibited similar

responses that were distinct from alcohols and normal hydrocarbons, and the hydrocarbons and alcohols had similar responses within their groups. The behavior of these compounds is more strongly influenced by the presence or absence of reactive groups such as CO or OH than by the hydrocarbon chain length. Signatures of other substances, such as phenolic compounds, have similar features that appear to be related to the chemical structures they have in common.

8. The signature of a mixture appears to be the superposition of the signatures of its constituents, but because the effect is exponential, the signature of the more reactive component will dominate. Therefore, if the substance that we seek to detect is the more reactive component, we will enjoy enhanced sensitivity (e.g., alcohol in gasoline). On the other hand, diluents such as alcohol can completely mask more subtle components such as essential oils.

9. Water has a relatively weak effect on the sensor array. This characteristic has two implications: first, humidity has little effect if we are looking at substances that are much more reactive than water; and second, if the array is properly configured, it might be possible to study flavor notes in coffee or tea.

10. Several pattern recognition approaches have been examined to determine their applicability to gas sensor arrays. Each technique was found to have its inherent strengths and weaknesses. No one technique was the best for all cases studied. The difficulty of the analysis is directly related to how much is known at the outset (number of possible species present, expected concentrations, ranges of responses, etc.).

11. Neural networks, as used in this study, have been successfully applied to the identification of gas signatures. The poorer performance of the Hopfield network was probably due to the small size of the pattern vectors, which shared many bits in common. This problem might be overcome by increasing the number of bits in the patterns. The Boltzmann network, a better performer, not only returns a "best" choice, but also indicates which patterns were "next best" choices by gray-scale shading. These intermediate values might be useful in interpolating gas concentrations. The Hamming network was the best performer for signatures that were input more or less directly (magnitude data). This performance is encouraging because it means a much simpler front-end signal processor for an actual device. This network also had some success with the correct identification of more than one compound in a mixture (using derivative data), and this fact indicates some exciting prospects for the future analysis of mixtures. Much more work is required in data conversion approaches and network tuning to realize the full potential of neural networks for olfactory analysis.

12. The clustering techniques had success in separating hexane signatures from ethanol and high concentrations from low concentrations in ethanol/water mixtures. These tests, too, were of a preliminary nature, and much remains to be done to fully appreciate their power in gas signature analysis. Either algorithm could be successfully applied to the analysis of concentration.

13. Preliminary tests on alcohol/gasoline mixtures hold the promise that this technology can indeed be applied to important industrial problems. In fact, the problem of detecting alcohol in gasoline can be solved by a two-sensor array, in which each sensor has a different set of responses to the possible gasoline/alcohol combinations.

14. Both the sensors themselves and the pattern recognition engines must ultimately be application-specific. No one system will be general enough to classify completely unknown mixtures from an infinite universe of possibilities; conversely, if relatively simple problems can be identified, the classification procedure can be more easily developed. Each of the algorithms studied can conceivably be implemented as an application-specific integrated circuit.

LIST OF REFERENCES

LIST OF REFERENCES

1. P. K. Clifford, *Selective Gas Detection and Measurement System*, U. S. patent, 4,542,640, 1985.
2. R. J. Lauf, B. S. Hoffheins, and C. A. Walls, *An Intelligent Thick-Film Gas Sensor: Development and Preliminary Tests*, ORNL/TM-10402, Oak Ridge National Laboratory (1987).
3. M. W. Siegel, Carnegie-Mellon University, personal communication (1985).
4. M. Kaneyasu, A. Ikegami, H. Arima, and S. Iwanage, "Smell Identification Using a Thick-Film Hybrid Gas Sensor," *IEEE Trans. Components, Hybrids, and Manufacturing Technology* **CHMT-10** (2) 267-73 (1987).
5. S. B. Crary, *Recent Advances and Future Prospects for Tin-Oxide Gas Sensors*, General Motors Report GMR-5556, Warren, Mich. (1986).
6. J. R. Stetter, S. Zaromb, W. R. Penrose, T. Otagawa, J. Sinclair, and J. Stull, "Portable Instrument for the Detection and Identification of Air Pollutants," *Proc. National Symposium on Recent Advances in Pollutant Monitoring of Ambient Air and Stationary Sources*, EPA-600/9-84-019, 1984.
7. H. Wolhtjen, A. Snow, and D. Ballantine, "The Selective Detection of Vapors Using Surface Acoustic Wave Devices," *Transducers '85, Proc. Third International Conference on Solid-State Sensors and Actuators*, Philadelphia, 1985, pp. 66-70.
8. R. Mueller, E. Lange, and A. Hinterstocker, "Multidimensional Sensors for Gas Analysis," *Ibid*, pp. 81-4.
9. R. C. Hughes, W. K. Schubert, T. E. Zipperian, J. L. Rodriguez, and T. A. Plut, "Thin-Film Palladium and Silver Alloys and Layers for Metal-Insulator-Semiconductor Sensors," *J. Appl. Phys.* **62** (3) 1074-83 (1987).

10. W. P. Carey, K. R. Beebe, E. Sanchez, P. Geladi and B. R. Kowalski, "Chemometric Analysis of Multisensor Arrays," *Sensors and Actuators* **9** 223-34 (1986).
11. W. P. Carey, K. R. Beebe, and B. R. Kowalski, "Selection of Adsorbates from Chemical Sensor Arrays by Pattern Recognition," *Anal. Chem.* **58** (1) 149-53 (1986).
12. S. Karpel, "Gas Sensors Based on Tin Oxide," *Tin and its Uses* (149) International Tin Research Institute, Uxbridge, UK (1986).
13. B. S. Hoffheins, R. J. Lauf, and M. W. Siegel, "Intelligent Thick-Film Gas Sensor," *Proc. 1986 Symposium on Microelectronics*, Atlanta, 1986, pp. 154-60.
14. M. W. Siegel, L. Wong, R. J. Lauf, and B. S. Hoffheins, "Dual Gradient Thick-Film Metal Oxide Gas Sensor," *Transducers '87, Proc. Fourth International Conference on Solid-State Sensors and Actuators*, Tokyo, 1987.
15. N. Taguchi, *Method for Making a Gas-Sensing Element*, U. S. patent, 3,625,756, 1971.
16. C. A. Harper, *Handbook of Thick Film Hybrid Microelectronics*, McGraw-Hill, New York, 1974, p. 118.
17. B. S. Hoffheins and R. J. Lauf, "Use of Chemical Sensor Arrays for Food and Fragrance Analysis," AICHE Annual Meeting, Washington, DC, 1988 (to be published).
18. R. W. Moncrieff, *The Chemical Senses*, CRC Press, Cleveland, Ohio, 1967.
19. Figaro Engineering, Inc. "Figaro Gas Sensor TGS" (product literature) Tokyo, 1976.
20. B. S. Hoffheins and R. J. Lauf, "Gas Sensor Arrays for Olfactory Analysis: Issues and Opportunities," *Sensors Expo Proceedings*, Chicago, 1988, pp. 205.1 - 205.8.

21. *ANSim User's Manual*, Science Applications International Corp., San Diego, 1987.
22. R. O. Duda and Peter E. Hart, *Pattern Classification and Scene Analysis*, Wiley, New York, 1973, p. vii.
23. M. A. Sharaf, D. L. Illman, and B. R. Kowalski, *Chemometrics*, Wiley, New York, 1986, p. v.
24. D. Coomans and I. Broeckaert, *Potential Pattern Recognition in Chemical and Medical Decision Making*, Research Studies Press, Letchworth, UK, 1986, p. vii.
25. J. J. Hopfield and D. W. Tank, "Neural' Computation of Decisions in Optimization Problems," *Biol. Cybern.* **52** 141-52 (1985).
26. D. L. Massart and L. Kaufman, *The Interpretation of Analytical Chemical Data by the Use of Cluster Analysis*, Wiley, New York, 1983.
27. G. O. Allgood, personal communication (1989).
28. J. T. Tou and R. C. Gonzalez, *Pattern Recognition Principles*, Addison-Wesley, Reading, Mass., 1974, pp. 94-7.
29. Ibid., pp. 92-4.
30. C. Jorgensen and C. Matheus, "Catching Knowledge in Neural Nets," *AI Expert* 30-41 (December 1986).
31. M. Verleysen, B. Sirletti, A. Vandemeulebroecke, and P. G. A. Jespers, "Neural Networks for High-Storage Content-Addressable Memory: VLSI Circuit and Learning Algorithm," *IEEE Solid-State Circuits* **24** (3) 562-9 (1989).

Supplementary References

S. Amari, "Field Theory of Self-Organizing Neural Nets," *IEEE Systems, Man, and Cybernetics SMC-13* (5) 741-48 (1983).

A. G. Barto, R. S. Sutton, and C. W. Anderson. "Neuronlike Adaptive Elements That Can Solve Difficult Learning Control Problems," *IEEE Systems, Man, and Cybernetics SMC-13* (5) 834-46 (1983).

M. A. Cohen and S. Grossberg. "Absolute Stability of Global Pattern Formation and Parallel Memory Storage by Competitive Neural Networks," *IEEE Trans. Systems, Man, and Cybernetics SMC-13* (5) 815-26 (1983).

D. Desieno, *Bidirectional Associative Memory Demonstration Program*, Logical Designs Consulting Inc., San Diego, 1987.

K. Fukushima, S. Miyake, and T. Ito. "Neocognitron: A Neural Network Model for a Mechanism of Visual Pattern Recognition," *IEEE Trans. Systems, Man, and Cybernetics SMC-13* (5) 826-34 (1983).

T. Hirschfeld, J. B. Callis, and B. R. Kowalski, "Chemical Sensing in Process Analysis," *Science* **226** 312-18 (1984).

J. J. Hopfield and D. W. Tank, "Computing with Neural Circuits: A Model," *Science* **233** 625-33 (1986).

R. A. Jacobs, "Increased Rates of Convergence Through Learning Rate Adaptation," *Neural Networks* **1** 295-307 (1988).

W. P. Jones and J. Hoskins. "Back-Propagation: A Generalized Delta Learning Rule," *Byte* 155-62 (October 1987).

G. Josin, "Neural-Network Heuristics: Three Heuristic Algorithms that Learn from Experience," *Byte* 183-92 (October 1987).

T. Kohonen, *Self-Organization and Associative Memory*, Springer-Verlag, Berlin, 1984.

B. Kosko, "Constructing an Associative Memory," *Byte* 137-44 (October 1987).

B. Kosko, "Bidirectional Associative Memories," *IEEE Trans. Systems, Man, and Cybernetics* **18** (1) 49-60 (1988).

R. P. Lippmann, "An Introduction to Computing with Neural Nets," *IEEE ASSP Magazine* 4-22 (April 1987).

E. Rich, *Artificial Intelligence*, McGraw-Hill, New York, 1983.

D. L. Rumelhart, J. L. McClelland, and the PDP Research Group, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, vol. 1., MIT Press, Cambridge, Mass., 1986.

G. Schöner and J. A. S. Kelso. "Dynamic Pattern Generation in Behavioral and Neural Systems," *Science* **239** 1513-20 (1988).

P. Süskind, *Perfume: The Story of a Murderer*, Penguin, New York, 1987.

T. Torioka and N. Ikeda. "Pattern Separating Functioning of Two-Layered Random Nerve Nets with Feedforward Inhibitory Connections," *IEEE Trans. Systems, Man, and Cybernetics* **18** (3) 358-66 (1988).

APPENDIXES

APPENDIX A

PATTERN RECOGNITION METHODOLOGY

PATTERN RECOGNITION METHODOLOGY

Much of the historical and current effort in pattern recognition design²² has focused on image processing and speech and language analysis, not on olfactory problems. However, the observer must make sense of gas sensor data through measurement comparisons and graphical means, and these do, in fact, constitute "scenes" of information. In light of this fact, many pattern recognition schemes become likely candidates for automatically processing these data, and the underlying chemical or physical principles of the gas sensor response are of no consequence to their design or operation. Collectively, the mathematical and technical approaches to chemical data analysis are called "chemometrics."²³ Many of these are intensely statistical and probabilities and can involve lengthy computation time,²⁴ so the use of these for this study was rejected on the grounds that they would be impractical for eventual use in application-specific hardware.

The approach in this study was to initiate the development of a general methodology based on the theory that selected pattern recognition techniques can be applied to any array that gives signatures. It is assumed that the sensor input to the pattern recognition system would be standardized in some way so that a general pattern recognition test bed could be used to define the parameters of an application-specific identification system. A coordinating function of the test bed would be to help eliminate or add sensors to get a "good" solution for a selected classification problem. (Hopfield²⁵ says that from an engineering or

economic standpoint what we really desire is not the "best" solution, but a "good" one that works.) The defined parameters could then be used in real-time software-controlled systems, embedded microprocessors, or in an application specific integrated circuit.

As with any chemical analysis system that relies on pattern recognition, the approach requires that the sensor array be "calibrated" by exposing it to known compounds, thereby developing a library of signatures. Unknown samples would then be identified by the signature recognition capabilities of the network, which would be chosen for its ability to analyze such signatures in real time. The application-defined system or network of the necessary size could be implemented as a single very large-scale integration (VLSI) chip or programmed in a microprocessor.

Just as one sensor array will not satisfy all applications, it is likely that one pattern recognition method will not work for every sensing problem or every sensor array; the recognition system will have to be customized, too. However, it will be easier to make successive models once the first one has been made.

Some existing pattern recognition tools were examined for their usefulness in resolving signatures collected from the gas sensor arrays. Two cluster analysis techniques, the K-means and maximum-distance algorithms were chosen to compare with results given by Hamming, Hopfield, and Boltzmann machine neural network simulation algorithms. (No claim is made that these techniques are the only ones that are applicable or that they are necessarily the best solutions for these data.) Resistance readings from the sensor arrays could be input directly in the

cluster analysis techniques, but the neural networks required a scene, or binary picture, as input, and so a preliminary transformation of the data had to be devised. Descriptions of the pattern recognition algorithms and data conversion methods follow.

Classical Methods: Cluster Analysis

Once plotted, the signatures of gas sensor data appear to be separable by clustering techniques (refer to Tables 3.1, 3.2, 3.4, and 3.5 on pages 55, 56, 69, and 70 respectively). Cluster analysis techniques are not unknown to the analysis of chemical data.²⁶ Also, Allgood²⁷ suggested that the gas sensor array data would be easily processed by either the K-means or the maximum-distance algorithms.

For this study, several data sets were processed by K-means and the maximum-distance algorithms as implemented by Allgood.²⁷ These tests were of a preliminary nature intended to demonstrate the ability of the algorithms to separate the data into logical groups.

The K-means algorithm²⁸ forces the data sets into the number of groups specified by the designer. That is, if the designer specifies four groups, the algorithm iteratively processes the data sets so that the members of each group are closer to their calculated cluster center than to the cluster center of any other group. Results of the K-means processing for the hexane/ethanol data and the water/ethanol data are given in Tables 3.1 and 3.4. One situation that occurs as a result of the

relatively large variation of the water responses is that some of the water signatures will be grouped by themselves. The designer of a system might wish to scale the data logarithmically to minimize this situation, or to recognize that certain groups will have to be lumped together as one. In an actual application, a test bed could define the cluster centers for the possible outcomes, and these could be programmed into the application-specific hardware; alternatively, the K-means algorithm could be implemented as part of a calibration procedure in the instrument, so that any effects of sensor drift could be minimized.

The maximum-distance algorithm²⁹ as implemented by Allgood²⁷ determines the number of clusters formed by use of a threshold value, which gives a maximum distance that an individual member of the group can be from the calculated cluster center. The designer specifies the threshold value to be tolerated, and the algorithm iteratively processes the data sets until all groups are defined. Results of using this algorithm with the hexane/ethanol and the water/ethanol data are given in Tables 3.2 and 3.5. Effects of the larger variation of water signatures and to a lesser extent the hexane and alcohol signatures are noticed in the formation of more than one group containing different signatures of the same compound (the same thing is noticed with the K-means algorithm). This technique could be incorporated into an instrument design in the same way as the K-means algorithm. Depending upon the chemical substances involved, clusters might be assigned different threshold values to accommodate variable shifts in sensor responses to a particular compound.

Neural Network Designs

It seems reasonable to apply the concept of neural networks as a means of pattern classification for gas sensor arrays. Several features of neural networks, summarized recently by Jorgensen and Matheus,³⁰ promise advantages over the chemometric approach. (1) Recall of information learned by the system is highly resistant to hardware failure (loss of a memory bit or two in traditional computers can be catastrophic). (2) The abstraction of data occurs automatically as a byproduct of learning. (3) Pattern recognition occurs in parallel and reconstructively. (4) The neural networks can exhibit adaptive features, which select and generate their own pattern features from exposure to the inputs. (5) Neural networks can capture patterns occurring both in time and space and operate in discrete or continuous representation modes. The attributes mentioned above could lead to features in a pattern recognition system that are not possible now. The reconstructive capability of a neural network should be much more forgiving of sensor drift and missing information in a gas signature; alternatively, or in addition, the neural network, programmed to behave dynamically, could overcome the drift problem and information loss. Also, an analog neural scheme, because of its ability to handle parallel computations rapidly,²⁵ might successfully address the identification problem in the overlapping patterns of gas mixtures.

Another reason for choosing neural networks as the basis of a pattern recognition system is because many workers, such as Verleysen

et al.,³¹ are designing and building neural networks in VLSI chips. This type of signal processing device is precisely the sort of component needed in a low-power and portable gas sensing system. It is a small, fast, and low-power device. (A Hopfield network, implemented in a VLSI chip, was proposed by Verleysen and coworkers. The chip is 64 mm², consists of 128 neurons that converge to a stable state in 150 ns, and dissipates about 100 mW.)

Several paradigms are available for testing and evaluation in the SAIC ANSim[™] Artificial Neural Systems (ANS) Simulation program, version 2.3. The program consists of two parts, the ANSpec^{™*} compiler for developing application code of the supplied ANS paradigm, and the ANSim[™] graphics-oriented, menu-based program for the user to develop ANS models based upon ANS paradigms. The ANSim[™] program allows for network training, monitoring, and storing for later use. Available paradigms include Adaptive Resonance Theory-Gray Scale, Back Propagation, BP Shared Weights, Hopfield, Boltzmann, and Hamming networks. The ANSim[™] system hardware used in the present work consists of a IBM AT[™] clone with an SAIC Delta Floating Point Processor^{™†} (22 MFlop AT bus compatible processors). Microsoft Windows^{™‡} is required for the operation of the user interface.

*ANSpec[™] is a trademark of Science Applications International Corporation San Diego, California.

†Delta Floating Point Processor[™] is a trademark of Science Applications International Corporation.

‡Microsoft Windows[™] is a trademark of Microsoft Corporation Redmond, Washington.

The Hopfield, Boltzmann, and Hamming networks were chosen because they work as autocorrelators that retrieve the patterns most like the input pattern presented to them. The operation of these networks also most closely fits the conception of how a real-time gas sensor array detector would work; that is, the stored patterns would be the possible choices with which the input pattern would be compared. All of these networks are trained with a set of training vectors that consist of bit patterns. Individual bits are valued either -0.5 or 0.5. In this study, the three networks were trained with the same sets of training vectors.

Each of the networks has slightly different characteristics that might render one or another more appropriate for a specific kind of gas sensor data, but this difference had to be determined experimentally. The basic network is the Hopfield network,²¹ which consists of units (neurons) that are connected to other units by weights determined during the training session by autocorrelating the input patterns. The state of the neuron is affected by its hard threshold nonlinearity interacting with its weighted connections to the other neurons. When trained, the network responds to an (unknown) input pattern by synchronously updating the connections made by the input in the weighted network until the network converges to a minimum energy state. The final state may not depict one of the original training patterns if the training patterns share too many bits in common. The Hopfield network requires the most care in training pattern configuration to ensure a stable and accurate network.

The Boltzmann machine network²¹ represents an improvement to the Hopfield network by an annealing process analogous to the method of

strengthening metals. A temperature is identified during each cycle through the network. A unit has an energy value equal to the weighted sum of all its connections to other neurons and will remain stable according to the Boltzmann probability if its energy is very high or very low at high temperature cycles of the system. It switches states if the cycle temperature is high and it has a midrange energy state. As the system goes toward lower temperatures, the fluctuations become infrequent and a minimum energy state is achieved. An annealing schedule consists of several cycles run at different temperatures to drive the network into a deep local minima.

The Hamming network²¹ is also improvement to the Hopfield net. The Hamming distance is used to compute a score between the input to the network and each of the training vector sets in the network. The network is then cycled until it converges on one of the training set patterns. This network always returns one of the patterns of the training set, unlike the Boltzmann and Hopfield networks.

Data Conversions for the Neural Network Programs

The main strategies of data conversion were to (1) make the patterns different enough to be distinguished by the network and (2) keep the pattern size fairly small (to minimize processing time and system size). Upon examination of the raw data, differences are fairly obvious to the observer, but to the neural networks, the general trends

between one signature and another might not appear all that different. The neural networks used require the input in the form of a bit pattern of -0.5 and 0.5 values. Size of the bit pattern matrix is defined by the user.

Two types of data conversions were performed. In the first, called magnitude data, the base-ten logarithms of the resistance values were plotted in a 9x10 or a 9x16 grid (Figure A.1). Columns were ordered by sensor number, left to right; rows corresponded to sensor response, as indicated. In general, only one block per column was filled in, although in training vectors several patterns were added together in hopes that variation and drift in signatures could be better accommodated. In the second type of data conversion, called derivative data, the bit pattern was determined by calculating the slopes between adjacent sensor responses and plotting the base-ten logarithm of these in 8x10 or 8x14 grids (Figure A.2). If the slope is negative, the signature is plotted as a histogram beginning at the middle of the grid and extending down a length equal to its magnitude. If positive, the histogram will start in the middle and extend upward.

Four sets of training vectors were constructed. The first two are of magnitude data, shown in Figures A.3 and A.4. The first of these has isopropanol, methanol, ethanol, heptane, hexane, and water training vectors; the second is a scaled-down version of the first that does not include the water vector. There are two sets of training vectors containing derivative data, shown in Figures A.5 and A.6. As with the magnitude training sets, one includes the water profile and the other consists of only the three alcohols and the two normal hydrocarbons.




Figure A.1. Magnitude data plots for four ethanol signatures (left) and one ethanol signature (right).






Figure A.2. Derivative data plots for an ethanol signature.


*Negative slopes are plotted as negative histograms, positive slopes are plotted as positive histograms.




Isopropanol




Methanol




Ethanol



Heptane




Hexane




Water


Figure A.3. Training set for magnitude data (9 x 16 grid).




Isopropanol




Methanol



Ethanol




Heptane




Hexane


Figure A.4. Training set for magnitude data (9 x 10 grid).




Isopropanol




Methanol



Ethanol




Heptane




Hexane


Figure A.5. Training set for derivative data (8 x 10 grid).




Isopropanol




Methanol




Ethanol



Heptane



Hexane



Water

Figure A.6. Training set for derivative data (8 x 14 grid).

(Note that water signatures, as seen in Chapter 3, can be 2 to 3 decades higher than those of alcohols and normal hydrocarbons, and the resulting bit patterns are dramatically different from those of the other substances.) As much as possible, the training vectors were selected from signatures that appeared to be representative of the particular compound.

APPENDIX B

NEURAL NETWORK RESPONSE DATA

NEURAL NETWORK RESPONSE DATA

The following figures give results of each neural network for each of the four training sets. One can consider two levels of success: Complete success would be indicated by all correct identifications of the actual compound (all points would lie along the diagonal), whereas partial success would be indicated by correct placement of an unknown into at least its correct chemical class (all points are within the large bold boxes on the diagonal). Correct identification of the network for its own training vectors is considered a minimal expectation of performance. For this reason, the box corresponding to the training vectors is also highlighted.

The solid-black circles in the boxes represent the network's classification of an input pattern every time that pattern was presented to the network. Open circles indicate that the classification was sometimes chosen, but not always. (The network converged to different answers when presented with the same signature more than once.) The dashes represent a less-than-positive classification, and this was indicated in the network simulation program by an intermediate gray-scale value. (The gray-scale value had to lie closer to a "yes" response for the identification in question than a "no" response.) Some network classifications simultaneously returned more than one positive response to one input pattern. (For this reason, the number of solid circles in some horizontal rows is greater than the number of different signatures for those compounds.)

The number of input patterns for a particular compound is shown in parentheses to the right of the compound name.

The first three figures (Figures B.1, B.2, and B.3) show the responses of each network trained with the five-element magnitude data set. The second three figures (Figures B.4, B.5, and B.6) give the responses of each network to the six-element magnitude data set. The third three figures, (Figures B.7, B.8, and B.9) are the responses of each network to the five-element derivative data set, and the last three figures (Figures B.10, B.11, and B.12) indicate the network's response to the six-element derivative data set.

Assignment by the Neural Network







		Isopropanol	Methanol	Ethanol	Heptane	Hexane	Hex/2% EtOH	Training Set
Actual Identity of Input Pattern	Isopropanol (5)							
	Methanol (5)							
	Ethanol (9)							
	Heptane (4)							
	Hexane (9)							
	Hex/2% EtOH (3)							
	Training Set (5)							

Figure B.1. Classification by Hopfield network trained with five magnitude data patterns.

Assignment by the Neural Network

		Isopropanol	Methanol	Ethanol	Heptane	Hexane	Hex/2% EtOH	Training Set
Actual Identity of Input Pattern	Isopropanol (5)							
	Methanol (5)							
	Ethanol (9)							
	Heptane (4)							
	Hexane (9)							
	Hex/2% EtOH (3)							
	Training Set (5)							

Figure B.2. Classification by Boltzmann network trained with five magnitude data patterns.

Assignment by the Neural Network

		Isopropanol	Methanol	Ethanol	Heptane	Hexane	Hex/2% EtOH	Training Set
Actual Identity of Input Pattern	Isopropanol (5)							
	Methanol (5)							
	Ethanol (9)							
	Heptane (4)							
	Hexane (9)							
	Hex/2% EtOH (3)							
	Training Set (5)							

Figure B.3. Classification by Hamming network trained with five magnitude data patterns.

Assignment by the Neural Network

	Actual Identity of Input Pattern					
	Isopropanol	Methanol	Ethanol	Heptane	Hexane	Water
Isopropanol (5)						
Methanol (5)						
Ethanol (9)						
Heptane (4)						
Hexane (9)						
Water (3)						
Hex/2% EtOH (3)						
Ethanol/Water (5)						
Training Set (6)						

Figure B.4. Classification by Hopfield network trained with six magnitude data patterns.

Assignment by the Neural Network

Actual Identity of Input Pattern	Isopropanol	Methanol	Ethanol	Heptane	Hexane	Water	Training Set
	Isopropanol (5)						
	Methanol (5)						
	Ethanol (9)						
	Heptane (4)						
	Hexane (9)						
	Water (3)						
	Hex/2% EtOH (3)						
	Ethanol/Water (5)						
	Training Set (6)						

Figure B.5. Classification by Boltzmann network trained with six magnitude data patterns.

Assignment by the Neural Network

		Isopropanol	Methanol	Ethanol	Heptane	Hexane	Water	Training Set
Actual Identity of Input Pattern	Isopropanol (5)							
	Methanol (5)							
	Ethanol (9)							
	Heptane (4)							
	Hexane (9)							
	Water (3)							
	Hex/2% EtOH (3)							
	Ethanol/Water (5)*							
	Training Set (6)							

Figure B.6. Classification by Hamming network trained with six magnitude data patterns.

*For one pattern, the network returned no classification.

		Assignment by the Neural Network						
		Isopropanol	Methanol	Ethanol	Heptane	Hexane	Hex/2% EtOH	Training Set
Actual Identity of Input Pattern	Isopropanol (9)							
	Methanol (8)							
	Ethanol (8)							
	Heptane (8)							
	Hexane (8)							
	Hex/2% EtOH (5)							
	Training Set (5)							

Figure B.7. Classification by Hopfield network trained with five derivative data patterns.

		Assignment by the Neural Network						
		Isopropanol	Methanol	Ethanol	Heptane	Hexane	Hex/2% EtOH	Training Set
Actual Identity of Input Pattern	Isopropanol (9)							
	Methanol (8)							
	Ethanol (8)							
	Heptane (8)							
	Hexane (8)							
	Hex/2% EtOH (5)							
	Training Set (5)							

Figure B.8. Classification by Boltzmann network trained with five derivative data patterns.

Assignment by the Neural Network

		Isopropanol	Methanol	Ethanol	Heptane	Hexane	Hex/2% EtOH	Training Set
Actual Identity of Input Pattern	Isopropanol (9)							
	Methanol (8)							
	Ethanol (8)							
	Heptane (8)							
	Hexane (8)							
	Hex/2% EtOH (5)							
	Training Set (5)							

Figure B.9. Classification by Hamming network trained with five derivative data patterns.

Assignment by the Neural Network

		Isopropanol	Methanol	Ethanol	Heptane	Hexane	Water	Training Set
Actual Identity of Input Pattern	Isopropanol (9)							
	Methanol (9)							
	Ethanol (14)							
	Heptane (8)							
	Hexane (12)							
	Water (none)							
	Hex/2% EtOH (1)							
	Ethanol/Water (5)							
	Training Set (6)							

Figure B.10. Classification by Hopfield network trained with six derivative data patterns.




Figure B.11. Classification by Boltzmann network trained with six derivative data patterns.

Assignment by the Neural Network

Actual Identity of Input Pattern	Isopropanol	Methanol	Ethanol	Heptane	Hexane	Water	Training Set
	Isopropanol (9)						
	Methanol (9)						
	Ethanol (14)						
	Heptane (8)						
	Hexane (12)						
	Water (none)						
	Hex/2% EtOH (1)						
	Ethanol/Water (5)*						
	Training Set (6)						

Figure B.12. Classification by Hamming network trained with six derivative data patterns.

*For one pattern, the network returned no classification.

VITA

Barbara Hoffheins was born in Bellefonte, Pennsylvania, on March 29, 1955. She attended grade schools in Roanoke and Herndon, Virginia, and was graduated from Anthony Wayne High School in Whitehouse, Ohio, in 1972. She pursued a study of forestry at The University of Tennessee, Knoxville, and The University of Idaho, Moscow, from 1972 until 1975. In 1977, she attended and completed the course of study at Tennessee Institute of Electronics in Knoxville. In 1978, she began employment at the Y-12 Plant in Oak Ridge as an electrical apprentice and completed this Department of Labor training in 1980. Later, the same year, she entered the electrical engineering program at Tennessee Technological University in Cookeville and in 1983 received the Bachelor of Science degree in Electrical Engineering. Since 1983 she has been employed by the Oak Ridge National Laboratory as a development engineer in the Instrumentation and Controls Division. During this employment, she has collaborated on the following papers:

1. "Intelligent Thick-Film Gas Sensor," B. S. Hoffheins, R. J. Lauf, and M. W. Siegel, *Proc. 1986 Symposium on Microelectronics*, Atlanta 1986 pp. 154-60.

2. "Dual Gradient Thick-Film Metal Oxide Gas Sensor," M. W. Siegel, L. Wong, R. J. Lauf, and B. S. Hoffheins, *Transducers '87, Proc. Fourth International Conference on Solid-State Sensors and Actuators*, Tokyo, 1987.

3. *An Intelligent Thick-Film Gas Sensor: Development and Preliminary Tests*, R. J. Lauf, B. S. Hoffheins, and C. A. Walls, ORNL/TM-10402, Oak Ridge National Laboratory (1987).

4. "Use of Chemical Sensor Arrays for Food and Fragrance Analysis," B. S. Hoffheins and R. J. Lauf, AIChE Annual Meeting, Washington, DC, 1988 (to be published).

5. "Gas Sensor Arrays for Olfactory Analysis: Issues and Opportunities," B. S. Hoffheins and R. J. Lauf, *Sensors Expo Proceedings*, Chicago, 1988, pp. 205.1-205.8.

6. "An Omnidirectional Antenna for a Personal Rf Dosimeter," B. S. Hoffheins, M. R. Moore, M. L. Bauer, *Proc. 1989 Symposium on Microelectronics*, Baltimore, Maryland (to be published).