

Visual-Empirical Region-of-Influence Pattern Recognition Applied to Chemical Microsensor Array Selection and Chemical Analysis

G. C. OSBOURN,* J. W. BARTHOLOMEW,
A. J. RICCO, AND G. C. FRYE

Sandia National Laboratories,
Albuquerque, New Mexico 87185-1423

Received June 3, 1997

I. Introduction

There is much current interest in the use of pattern recognition (PR) techniques to identify and quantify chemical analytes on the basis of the multivariate responses of chemical microsensor arrays.^{1,2} PR techniques use previously measured array responses of the class categories of interest, called training set data, to infer the class category associated with new array measurements (test data). The class categories can be the chemical identities of the analytes and can also represent the analyte concentrations. Each sensor in the array provides the response along one dimension of a multidimensional response vector, and each array response is a well-defined geometric location in a multidimensional response space (usually called the feature space).

The application of PR techniques to microsensor-based chemical sensing is motivated by two advantages: (i) chemical identification is possible using only partially chemically selective microsensors rather than highly selective (and more difficult to develop) microsensors, and

(ii) PR can, in principle, identify a large number of chemical species using a small, fixed set of sensors. Many research groups have been exploring the combination of PR analysis and sensor array hardware to carry out portable, low-power chemical analyses. Many computational techniques are available for such analyses, including feature selection techniques (e.g., principal components analysis (PCA)), statistical-based methods for PR (e.g., SIMCA), and nonstatistical methods for PR (e.g., neural nets and k nearest neighbors (k -NN)). The literature on the use of these techniques for chemical sensing is reviewed in refs 1 and 2.

These techniques are widely used and are certainly useful in certain problem domains. However, the reader may be unaware that these popular techniques have limitations that can be undesirable for certain chemical sensing applications. We briefly review these limitations in the next section. Our intent is not to suggest that the popular techniques are useless, but rather to make clear the motivations for the development of a new PR approach that overcomes some of these limitations. Our PR method uses a new principle (visual-empirical region-of-influence (VERI)^{3,4}), derived from human vision research results, that is indeed designed to address these limitations. We present VERI as an effective alternative to PCA (for feature search), statistical PR, and neural net PR approaches for chemical sensor array applications. In the rest of this paper we outline the VERI PR method, emphasize the situations where its properties are particularly advantageous, and demonstrate its use in several phases of the development of a chemical sensor array system. Space limitations prevent our discussion of all of the sensor systems we have analyzed using VERI. Here we present results from arrays of coated surface acoustic wave (SAW) devices applied to the detection of volatile organic compounds (VOCs).

II. Motivation for an Alternative PR Technique

II.a. Selecting Sensors for the Array.

A key step in the successful development of a PR approach for any specific

Dr. G. C. Osbourn received the B.S. degree in physics and mathematics in 1974 and the M.S. degree in physics in 1975 from the University of Missouri at Kansas City and the Ph.D. degree in physics in 1979 from the California Institute of Technology. He joined the technical staff of Sandia National Laboratories in 1979 and is currently Manager of the Vision Science, Pattern Recognition and Multisensor Algorithm Department. Osbourn was selected as 1 of 100 outstanding young U.S. scientists by *Science Digest* in 1984. He received the 1985 Department of Energy E. O. Lawrence award and the 1993 American Physical Society International Prize for New Materials for initiating the study of a new class of semiconductor quantum wells and superlattices. He has coauthored 75 publications and has been awarded 6 patents on semiconductor devices and 1 patent on image processing. He is a fellow of the American Physical Society and a member of the Optical Society of America. His current research interests include the study of human vision, the development of robust pattern recognition and clustering algorithms, multivariate image classification applications, and the application of those techniques to chemical sensing problems.

J. W. Bartholomew received the B.S. degree in electrical engineering in 1982 and the M.S. degree in electrical engineering in 1983 from Kansas State University. He joined the technical staff of Sandia National Laboratories in 1984 and is currently a senior member of the technical staff of the Vision Science, Pattern Recognition and Multisensor Algorithm Department. He is a member of the Institute of Electrical and Electronics Engineers.

A. J. Ricco was born in Oakland, CA, in 1958. He received a B.S. in chemistry (magna cum laude) from the University of California, Berkeley, in 1980 and a Ph.D. in inorganic chemistry from the Massachusetts Institute of Technology in 1984. He joined Sandia National Laboratories' Microsensor Research Division in 1984, where his research currently involves chemical sensor arrays based on acoustic wave, optical, electrochemical, chemiresistor, and various microelectronic platforms. He focuses on new ways to utilize chemical and physical effects together with new materials in combination with microelectronic technology and pattern recognition to develop novel chemical sensors with enhanced capabilities.

G. C. Frye received his B.S. degree from the University of Colorado in 1982 and his Ph.D. from the University of Washington in 1986, both in chemical engineering. Since joining Sandia National Laboratories in 1987, his research and development interests have been focused on thin film materials, acoustic wave sensors, and environmental monitoring systems. He currently is using this background to develop miniaturized chemical analysis systems that provide sensitive and selective detection of target analytes such as chemical warfare agents. This development effort is utilizing microfabrication techniques to integrate sample collection, chromatographic separation, and chemically selective detection functions on chip.

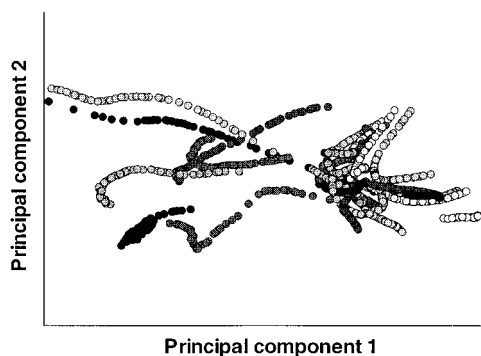


FIGURE 1. First 2 principal components of a training set of 28 individual volatile organic compounds (VOCs) and binary VOC mixtures obtained with a six-sensor array. The different chemicals are given different shades of gray. This set is referred to as chemical set B in section IV.b. Both our VERI PR method and a neural net analysis indicate that these 28 chemical classes are fully separated by the full six-sensor array responses, yet these first 2 principal components fail to separate (distinguish) the classes. This example illustrates that the first few principal components do not necessarily indicate the effectiveness of a set of sensors for pattern recognition.

chemical sensing problem is the selection of sensors to include in the array. Ideally, these sensors provide maximal separation of the clusters of multidimensional array responses for the distinct target chemicals. It is a common practice for statisticians to use the PCA method to guide the choice of features from a larger candidate set.² PCA computations provide useful qualitative insights and guidance for removing redundancies from among very large numbers of potential sensor combinations. Each principal component direction is a linear combination of the original sensor responses, with the first few components indicating directions for which the variations in the training data are largest. PCA can provide a reasonable and simple model of the spread of responses, i.e., Gaussian variation along a few principal directions. However, the first few principal components may contain large contributions from all of the original sensors. Thus, the PCA method does not directly reduce the number of sensors³ required for the array system, and it is common practice to retain all sensor contributions to the small set of principal components selected. PCA does not provide quantitative comparisons that address the many issues (e.g., what are the best three- and four-sensor arrays from a set of N sensors, and how different are their PR performances?) in making sensor downselections.

Furthermore, the PCA method is not guaranteed to provide the best (or even a good) choice of features for *distinguishing classes*.⁵ The multivariate directions which contain most of the variations in the data distributions need not be the directions along which complicated class shapes actually separate. Further, the visual separation of classes provided by the first few principal components may be quite misleading in some cases, as we illustrate in Figure 1. These normalized data, which we describe more fully in section IV.b, show the first two principal components of the six-sensor array responses to 28 distinct volatile organic compound mixtures. As can be seen, these two PCA components indicate that the chemi-

cal classes are poorly separated. However, both our VERI analysis and neural network results for these data agree that this chemical set is completely separated in the full six-sensor response space. This example illustrates that the responses of modern chemical microsensors do not all transform readily into simple Gaussian distributions along principal axes. PCA-selected features will not produce optimal, or perhaps even acceptable, PR performance for such data.

A more computationally intensive approach to sensor selection is to use a PR method, combined with some optimization search procedure, to hunt for good choices of the sensors. The PR method directly evaluates each potential array choice that is selected by the optimization procedure, so that array selection is appropriately based on actual PR performance. Difficulties arise if the PR method requires the user to provide design decisions or design calculations (to make the method operational) for each sensor set considered. This is the case for neural networks, for which training calculations are needed for each new set of sensors. This additional computational overhead restricts the number of sensor combinations that can be examined. Also, the user may inconsistently train the network for different sensor sets, or inconsistently choose reject thresholds, thereby biasing the sensor selection. These difficulties would be eliminated if the PR method required no user inputs and no optimization computations for each sensor combination considered.

II.b. Statistical vs Nonstatistical PR Approaches. Statistical PR methods have a sound mathematical foundation and are powerful approaches to apply to PR problems *if the statistical assumptions underlying these methods are valid for the application*. Statistical approaches typically are based on PCA models of the data distributions, and thus require that sensor responses can be approximately transformed into Gaussian distributions along principal axes. As indicated above, this is not necessarily true for modern microsensors. Such sensors present a choice in system design—either eliminate these types of sensor responses from consideration or consider a different PR methodology. Statistical approaches determine the PR result by selecting the class with minimum probability of error. This requires that the probabilities/frequencies of future occurrence are known for all target chemicals in the application environment. This can be true for certain manufacturing and quality control applications. However, there are many important chemical sensing applications for which frequencies of occurrence of target chemicals in an uncontrolled environment are completely unknown. These and other issues have led to a continuing interest in developing nonstatistical PR methods, e.g., neural nets. We limit our discussion to nonstatistical methods for the remainder of this Account.


II.c. Uncontrolled Chemical Environments. Exposure to Unknown Chemicals. Neural networks have become a popular choice for carrying out the PR task in chemical sensing applications. The back-propagation method is often used to “train” the networks, so that they will produce network output responses that indicate the

selection of one of the target chemicals as the best match to any input. Such networks can be quite effective at detecting subsequent exposures to one of the target chemicals, and can work well for applications in which tight control exists over which chemicals are presented to the system. However, the performance of these networks on data from uncontrolled environments may be quite inadequate. Arrays in uncontrolled environments typically encounter chemicals other than those in the target set, and the PR methodology must be capable of recognizing unexpected chemicals correctly (i.e., as unknowns). We illustrate the performance of a neural net from the OLPARS software package version 7.3 (available from PARS Government Systems Corp., La Jolla, CA) for such a case in Figure 2. This is a net with a single hidden layer. Figure 2a shows hypothetical training data obtained for three classes using two feature measurements (the x and y axes). As is evident by direct visual inspection, this training data set consists of classes that are partially overlapped, classes that are "touching", and a surrounding volume that belongs to none of the classes. Figure 2b shows the class volumes computed for these data using the neural net approach. The class volumes were determined by examining the class identity of a grid of vectors, and shading the entire pixel surrounding each grid with the gray shade of the corresponding data points of the same class in Figure 2a. Note that the class volumes are *unbounded*. This simply indicates that the neural net always returns the best match to one of the target chemicals in the training set, even when the actual match may be quite poor. In other words, *any* arbitrary array response will be recognized as one of the training set chemicals by the neural net.


Detection of outliers is often provided through the supplemental use of a threshold value on the magnitude of the neural net output. Figure 2c shows the class volumes that result when a large reject threshold is applied to the output node with the strongest response. The black regions are those where input vectors are rejected. Note that these new reject regions are unbounded also, yet they do not correspond to the outlier volumes. As seen in Figure 2b, smaller thresholds (or no threshold) produce smaller or nonexistent reject regions. No choice of threshold value produces the correct two-dimensional volumes associated with the outliers. Unfortunately, the degraded performance that may result from the use of simple reject thresholds may go undetected by the user for higher-dimensional array data which cannot be plotted for inspection.

We note that another popular nonstatistical method, i.e., k -NN, produces similar results on these data. The k -NN method produces class volumes (not shown) that are unbounded without thresholding, and adding thresholding to the results produces reject regions which do not match the true class volumes associated with the data.


II.d. Applications with Partially Indistinguishable Target Chemicals. Warning of Unreliable PR Results. Figure 2 also illustrates another potential problem with the use of neural networks for applications involving the



(a)



(b)



(c)

FIGURE 2. (a) Hypothetical training set from two sensors containing three classes (indicated by shades of gray). (b) Neural net class volumes computed for the training set in (a). The gray values of the training points are used to indicate the class identification computed at each location, i.e., at each two-sensor vector response. A reject threshold of 0.5 is used, but this and smaller thresholds have little effect on the class volumes. The class volumes are unbounded, so that any responses of the two sensors will be identified as one of these three known classes. (c) Same as (b), but with the value of the reject threshold increased to 0.9. Regions of reject response, labeled black, are now present between the classes. No value of the reject threshold can produce neural net outlier detection that matches the class shapes of this simple example.

discrimination of closely related chemicals. It may occur, as in Figure 2, that the best sensors available to the user are incapable of distinguishing all of the target chemicals over all concentration ranges, in which case it becomes necessary to deal with the possibility of confusions between similar chemicals. All PR techniques will necessarily produce errors in such cases if they simply select the best class match. For certain applications, it may be unacceptable to make such errors. For example, frequently confusing a rarely encountered, but high-conse-

quence chemical (say, a chemical warfare agent) with a harmless, commonly encountered chemical will produce frequent alarms that are unacceptable. In such cases, simply reporting the best class match is insufficient. It would be advantageous to obtain a warning from the PR method when the detection of the high-consequence chemical is uncertain, and to know which other chemical(s) are also possible matches to the input. One could attempt to obtain this sort of information through the use of a threshold on the neural net output, but Figure 2 illustrates that no choice of threshold value produces the correct two-dimensional volumes associated with the unreliable (multiple-class match) output detection. Further, there is the added uncertainty of distinguishing an outlier rejection from a multiple-class-match detection.

The examples in sections II.b, II.c, and II.d illustrate that the reliable detection of unknown chemicals and the appropriate warning of multiple class matches can require multivariate class volume shapes that are not properly produced through the choice of a single threshold value on the neural net output. In general, unnecessary class identification errors may result from any PR methodology that relies on a single, user-selected threshold to determine the multivariate class volume shapes and sizes.

III. VERI PR

The situations described above motivated the development of the VERI approach to PR. VERI is a nonparametric/nonstatistical PR technique. As such, it is applicable to any kind of sensor responses (e.g., nonlinear, including nonmonotonic, as a function of a single chemical concentration and nonadditive as a function of multi-component mixture concentrations). VERI provides information beyond the traditional PR task of finding the best class match for any multivariate input vector of measured features. The automated VERI method, rather than the user, also implicitly determines multivariate threshold functions for the two key issues discussed in sections II.c and II.d: whether the best class match is a sufficiently good match or an outlier; whether other class matches are as acceptable as the best class match, so that the best match is only an unreliable guess. The VERI method automatically produces bounded class volumes with sizes, shapes, and interclass overlaps (if any) that are consistent with the sampling of sensor array measurements used to represent each chemical class. As with any PR method, the user is responsible for acquiring a good set of training data that contain the range of sensor responses to be encountered in the application.

The completely automated VERI approach is a useful alternative for PR analyses in several ways. The method can handle any type of sensor response. The method provides true multivariate determination of the detection of outliers and unreliable (multiple) class matches for each application, regardless of the complexity of the sensor response distributions. VERI PR is particularly effective for comparisons of performance of alternative sensor selections for designing a new array. VERI automated

sensor array optimization allows the PR performance achievable from different sensor sets to be compared consistently without user input. The automated VERI approach allows an inexperienced user to correctly apply PR to problems in which unknown chemicals or incompletely distinguishable chemicals may be encountered, but without the need to supply any algorithm design inputs, algorithm optimization calculations to make the method functional (e.g., neural net back-propagation training), or multivariate reject thresholds into the algorithm.

The VERI PR method is based on the VERI clustering technique.³ The reader is invited to visit our web site⁴ to obtain a detailed description of the method and a discussion of the vision research that forms the foundation of the approach. The VERI method is a multivariate clustering method with two essential properties. First, it produces clustering results without any user inputs. This is the crucial property that allows us to carry out fully automated PR. Second, the technique provides clustering results for two-dimensional data that agree with typical (measured) human judgments. This provides assurance that the VERI class boundaries are computed in a way that is consistent with typical human clustering judgments, where such judgments can be made. We have recently posted a "user's guide" at our website, which also describes the use of VERI for general PR applications,⁴ and the reader can consult this reference for all other details of the use of this technique.

The key idea underlying VERI PR is to regard the assignment of class membership of a new data vector as a VERI clustering problem. We assume that the user has acquired a set of training data, obtained using some collection of sensors, which contains multiple concentrations of the chemical classes that are to be detected. This sensor set may include many more candidates than the final array subset ultimately selected for the application. We also assume that the user has attempted to include training data representative of the entire range of variability that is expected to be encountered subsequently in the test data. VERI clusterings are computed between any new data vector and the training data vectors. The class of the new data vector is inferred from the class identities of the training vectors that cluster with it.

VERI requires a preprocessing step—a leave-one-out analysis of the training data. This step "holds out" each array response in the training data and examines it as if it were a new measurement. This step evaluates whether class overlaps exist and whether the classes are adequately sampled.⁴ This step is also the one used to search for an optimal selection of sensors for a new array application. The presence of class overlaps is indicated if any of the held-out training points cluster with a remaining training point of a different chemical class. If the final selected sensor set does not completely separate the chemical classes, then the training vectors that are in the class-overlap regions are flagged and labeled with the additional chemical class(es) that overlap with them. Flagging training vectors in a class overlap region ensures that the




FIGURE 3. Class volumes, computed using VERI, for the training data in Figure 2a. The gray values of the training points are used to indicate the class identification computed at each location, i.e., at each two-sensor vector response. In addition, VERI finds two other types of volumes: outlier (black) and multiple class matches where classes touch or overlap (lightest gray). The class volumes are bounded and follow the shape of the set of training vectors, the outlier volume properly surrounds the classes, and the multiple-class-match volumes correctly indicate the presence and locations of the touching and overlapping classes. These class volumes are obtained by VERI directly from the training data; i.e., no user inputs or parameter adjustments are supplied to the algorithm.

overlap-class volume computed on subsequent PR data is complete, as described in ref 4.

After the leave-one-out calculations are performed and the sensor set is selected, VERI PR on new test data can be carried out. VERI clusterings of each new array response vector with the training data are computed, and the class membership(s) of the training data vectors that group with the test vector are assigned to the test vector. Three possible cases occur^{3,4} for the class assignment of each new array response vector.

(i) Unique class identified. The new vector is clustered only with training data of a single chemical class and is assigned that same chemical identity.

(ii) Unknown class. The new vector is clustered with none of the training data and is designated an outlier (unknown: not in the target set of chemicals).

(iii) Multiple matching classes. The new vector is clustered with training data that contain more than one class identity. This can occur either by clustering with multiple training vectors that have differing classes or by clustering with a single training vector that was flagged as being in a class-overlap volume in the leave-one-out preprocessing computation. These vectors occur where the training-set class distributions overlap or touch, and cannot be assigned a single chemical identity without potential error. The potential class matches consist of the complete list of classes encountered through the clusterings.

The VERI method does not explicitly compute class boundaries in feature space, but the training data set implicitly defines such boundaries through the VERI clusterings. Vectors which do (do not) cluster with training data vectors in a given class are (are not) in the associated class volume. The performance of VERI can be better visualized by examining the class volumes that result from example data, and in Figure 3 we present the class volumes for the example data of Figure 2a. As before, the class boundaries were determined by examin-

ing the class identity of a grid of vectors, and shading the entire pixel surrounding each grid vector according to the computed class membership. Note that the VERI method produces distinguishable class volumes, outlier (black) volumes for which no target chemical is an acceptable match, and class-overlap volumes (labeled as the lightest gray shade in the figure). These class volumes are obviously quite different from those produced by our example neural net for any choice of net threshold. We also note that the VERI results are quite different from k -NN results (not shown) on these same data.

Several general properties of VERI class volumes can be seen from this example. First, the class volumes always surround each of the training vectors, and the class volumes are always bounded (for training sets with two or more vectors). This property guarantees that outlier test vectors can be detected in all applications of the VERI method. The VERI algorithm *automatically* defines the class volumes in a manner that is locally consistent with the sampling density of the class distribution.⁴ This is a sensible result, as knowledge of the class boundary position cannot be known to better spatial accuracy than the local sampling separations of the training data. Note that the class overlap regions are also consistent with typical human judgment, and provide direct warning of potential class confusion while also providing all acceptable class matches. We again emphasize that the bounded class volumes are implicitly achieved by VERI with *no* user-adjusted parameters or threshold values. All that is required is the training set.

The heavy reliance of all PR techniques on the training data necessitates careful selection of the training data set used for each chemical. VERI uses the training vector distribution alone to estimate the class volume of each chemical class. For example, assume that a set of just three responses to a given chemical, at low, moderate, and high concentrations, is found to be highly collinear. If the researcher understands the chemistry of the problem (hypothetically, relatively weak physical adsorption being the dominant factor), she/he may reasonably conclude that all intermediate concentrations fall on the same line. The VERI algorithm does not make any such assumption of "wide-range linearity" (an assumption that could be detrimental as well as beneficial). Therefore, such a training data set must be augmented by a series of interpolated points lying between the experimentally determined values. In the event that the data are not found (or expected) to be collinear, or if the operator is not expert in the likely chemical interactions, training data must be obtained with adequate concentration resolution to define the desired operating range without significant gaps.

IV. VERI Applied to Chemical Sensing

In this section we report on a series of studies that combine VERI PR with coated surface acoustic wave (SAW) device technology to achieve robust identification of sets of VOCs and water over wide concentration ranges.

Robust monitoring of these chemicals is of interest for a number of industrial and DOE applications.⁶ Two distinct sets of SAW coatings and three sets of VOCs were used to generate array responses. The first film set, developed by Ricco, Crooks, Allred, and co-workers, are self-assembled monolayer (SAM) films and plasma-grafted polymer films (PGFs). These films and the data acquisition method are described in a companion paper.⁷ The two sets of chemicals examined using these devices were 12 individual VOCs (DIMP (diisopropylmethylphosphonate), DMMP (dimethylmethylphosphonate), acetone, benzene, toluene, carbon tetrachloride, trichloroethylene (TCE), cyclohexane, isooctane, methanol, pinacolyl alcohol, and 1-propanol) (section IV.a below) and water, and 6 VOCs (acetone, DIMP, isooctane, 1-propanol, TCE, and toluene) and their 21 binary combinations (section IV.b below). We will refer to these VOC chemical sets as set A and set B, respectively. The second SAW film set, developed by Frye, Grate, and co-workers, are the polymer films: poly(diphenoxyphosphazene) (PDPP), poly(epichlorohydrine) (PECH), a fluorinated polysiloxane (OV-215), ethyl cellulose, poly(isobutylene) (PIB), a cyanide-modified polysiloxane (OV-275), poly(ethylene glycol adipate) (PEGA), poly(ethylenimine) (PEI), poly(vinyl propionate) (PVPR), a hydrophobic fluoropolyol (FOX), poly(vinyl acetate) (PVA), polyvinyltetradecanal, and hydrogen bond acid modified polysiloxane (HBA). Polymer deposition techniques and the data acquisition method are also described elsewhere.⁸ The chemicals examined by these devices were water and nine individual VOCs: acetone, chloroform, hexane, 2-propanol, methyl ethyl ketone (MEK), perchloroethylene (PCE), toluene, carbon tetrachloride, and TCE (section IV.a below). We will refer to these as chemical set C.

It is common for differently coated sensors to yield ranges of responses that are quite different in magnitude. We equalize the responses of each sensor by a multiplicative scale factor, so that the largest absolute training data response of each sensor is unity. Subsequent input responses must be scaled by these same factors for PR analysis.

An essential ingredient for the PR development process is the examination of an adequate set of target chemical training data that represents the expected variability of the test measurements over the lifetime of the sensors. One source of variability, which is not convenient to obtain experimentally (with newly fabricated sensors), is the change in SAW response sensitivity that can occur as the SAW films age during sensor use. The physics of such effects is not understood in sufficient detail to provide a quantitative model of this variability. Empirical data on a subset of aged sensors indicate that response aging can be approximated by the original data reduced by a multiplicative factor. We produced augmented training and test sets of data, in which the original sensor responses are simply reduced by a multiplicative factor for a given amount of aging. We do not have available the relative aging rates for all the different SAW films. To avoid selecting arbitrary (e.g., equal) aging rates for each

film which might give uncharacteristic results, we examine an ensemble of sensitivity loss combinations, where each array response in the original training data set is assigned a randomly chosen sensitivity loss for each sensor. The randomly selected sensitivity loss for each sensor is constrained to be less than some ceiling value (e.g., 10% loss), and this ceiling can be varied to study the overall effects of increasing aging. We note that this model is a worst-case situation, and that it also effectively represents worst-case noise/variability from other sources, e.g., uncontrolled SAW array temperature in field use. Thus, our aging model can be considered as a general numerical representation of any source of sensor response variability.

We achieve and demonstrate robust chemical recognition using the two types of SAW arrays in several steps: (i) We first optimize the choice of SAW coatings, as a function of array size, for identifying the individual VOCs. This step finds those combinations of films that produce multidimensional data vectors for the different chemicals that are most distinct. (ii) For cases where we wish also to detect simple VOC mixtures, we then evaluate the performance of one of the best SAW arrays on a new and larger set of individual VOCs and binary VOC mixtures, and we also evaluate the (modeled) effects of SAW film aging (or system noise) on the PR performance. (iii) Although not a required part of the PR process, we also estimate the response of the arrays to unknown chemicals by a "leave-one-CLASS-out" procedure, where each VOC chemical set is held out of the training set and tested as an unknown chemical class. (iv) Finally, we discuss two of the many alternative approaches for estimating VOC concentrations.

IV.a. Optimizing Selection of Array SAW Films. We describe the optimization results used to select a "good set" of SAW films for VOC discrimination. We seek to identify the smallest subsets of sensors with the best PR performance. The smallest arrays are desirable from a practical standpoint in that they reduce hardware complexity and system cost. Our approach is to carry out a VERI leave-one-out computation for each array combination and directly compare the PR performance. This is only possible for a limited number N of sensor alternatives, as the number of arrays of m sensors that can be formed from N available sensors is $N!/(m!(N-m)!)$. However, for $N < 20$, this optimization approach is feasible using 1997 workstation computing power. To date, the sensor array optimization tasks we have undertaken have all involved fewer than 20 SAW coating options at a time. If much larger sets of sensors are available to select from, we recommend using the VERI leave-one-out computation on array selections picked by an optimization search algorithm (e.g., genetic algorithms). We note that, after a single array is selected, subsequent VERI PR on the individual array can be carried out in real time (typically less than a few seconds per input vector) on PC hardware.

Figure 4 shows a histogram of the VERI optimization results for small arrays that can be selected from nine SAM and PGF films on the individual VOCs of chemical set A.

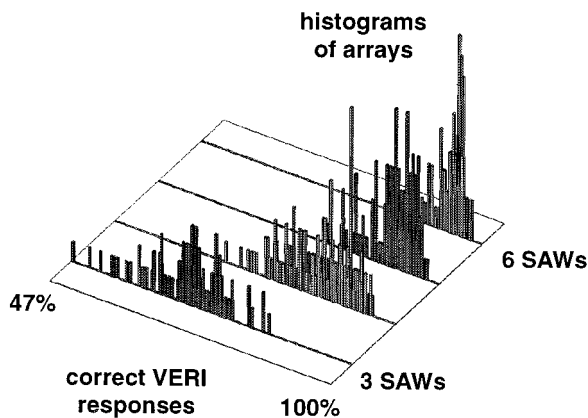


FIGURE 4. Results of the VERI array optimization study for chemical set A. A histogram of arrays is given as a function of the percentage of correctly identified training vectors in leave-one-out computations. The arrays are separated out by size, with results for arrays with three to six sensors shown here.

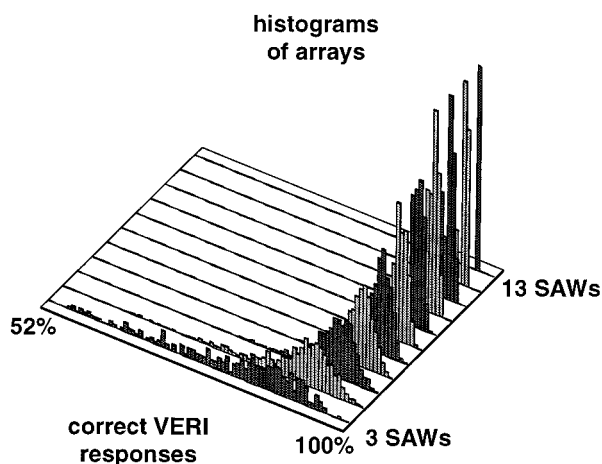


FIGURE 5. Results of the VERI array optimization study for chemical set C. A histogram of arrays is given as a function of the percentage of correctly identified training vectors in leave-one-out computations. The arrays are separated out by size, with results for arrays with three to thirteen sensors shown here.

The VERI percent correct classifications on the training set A for each of the potential arrays is obtained, and the number of array combinations that yield a particular percentage of incorrect chemical identifications is plotted for arrays with 3–6 films. Over 200 array combinations were examined. The overall PR performance improves as films are added, and there are array combinations exhibiting nearly 100% correct identifications for four or more films in the arrays. The best arrays produce no incorrect identifications, but do not reach exactly 100% correct because of a small number of outliers which we did not remove before these studies. Figure 5 shows a similar histogram for the arrays that can be selected from 13 polymer films on the individual VOCs of chemical set C. Over 8000 array combinations were examined. Again, the overall performance improves as films are added. These results demonstrate that there are more array selections that provide good chemical discrimination performance as the size of the array increases. However, there are also many poor array selections to avoid, particularly when the




FIGURE 6. Same as Figure 5, but using training data with up to 50% sensitivity loss (or noise) added to the sensor responses.

array is small, as these arrays provide dramatically degraded chemical discrimination performance.

It is also of interest to estimate what effects SAW film sensitivity loss (or noise) has on the overall performance of all potential arrays. While we cannot afford to compute the results for numerous aging amounts (this is done for a single array in the next section), here we choose a large aging ceiling of 50% to intentionally create significant degradations in all of the array results. This allows an observation of the trends to be expected. We carry out this additional computation on the polymer films of Figure 5, and show the resulting histograms for a ceiling of 50% sensitivity loss in Figure 6. The key result that is apparent from this histogram is that chemical discrimination performances of the best large arrays have degraded much more than those of the best small (five- to six-sensor) arrays. Figures 5 and 6, taken together, indicate that there is an optimum polymer array size of five to six polymer-coated SAWs for achieving *and maintaining* the best long-term VOC discrimination. These results demonstrate the importance of carefully selecting the number and identities of partially selective microsensors for a chemical sensing application.

IV.b. Effects of Device Aging on Distinguishing Individual VOCs and VOC Binary Mixtures. We selected one of the best six-film arrays from the optimization of set A to investigate the ability of this approach to treat both individual VOCs and VOC binary mixtures, and to further quantify the effects of SAW sensitivity aging. The six films are monolayers self-assembled on Au films from $\text{SH}(\text{CH}_2)_{15}\text{CH}_3$, $\text{SH}(\text{CH}_2)_{10}\text{COOH}$, and $\text{SH}(\text{CH}_2)_{10}\text{COO}^-/\text{Cu}^{2+}$ and plasma-grafted films fabricated from a 5-min base layer of acrylic acid plus a 30-min graft of acrylic acid, a 5-min base layer of polyisobutylene plus a 15-min graft of acrylic acid, and a 30-min plasma-polymerized layer made from eugenol (see ref 7 for an explanation of the film chemistry and synthetic procedures). Chemical set B was used for this study. Set B contains fewer individual VOCs than set A, but set B also contains all binary combinations of the individual VOCs. Binary VOC mixture data were acquired, starting from a high concentration/low concentration mixture. By increasing the concentra-




FIGURE 7. VERI leave-one-out results for one of the best six-sensor arrays, identified from the results of Figure 4, on chemical set B. Each of the chemical classes is identified by number here: (1) acetone, (2) acetone/DIMP, (3) acetone/water, (4) acetone/isooctane, (5) DIMP, (6) DIMP/water, (7) DIMP/isooctane, (8) DIMP/1-propanol, (9) water, (10) water/isooctane, (11) water/1-propanol, (12) water/TCE, (13) isooctane, (14) isooctane/1-propanol, (15) isooctane/TCE, (16) isooctane/toluene, (17) 1-propanol, (18) 1-propanol/acetone, (19) 1-propanol/TCE, (20) 1-propanol/toluene, (21) TCE, (22) TCE/acetone, (23) TCE/DIMP, (24) TCE/toluene, (25) toluene, (26) toluene/acetone, (27) toluene/DIMP, (28) toluene/water. The results include up to 0.1% (i.e., essentially new sensors), 8%, and 12% sensitivity losses. The errors become significant by 12% sensitivity loss for this chemical set.

tion of the low component while decreasing the concentration of the high component, the opposite low/high mixture limit is reached.⁷ Aging (noise) effects were added as described at the beginning of section IV, and three ceiling values of 0.1% (essentially new sensors), 8%, and 12% were used.

Figure 7 shows a histogram of incorrect chemical identification of water, individual VOCs, and binary mixtures (numbered along the *x* axis, with the chemicals given in the caption), with the ceiling values of the sensitivity shifts given along the *z* axis. The average incorrect identifications over the entire chemical set are given by the results labeled "avg". The results on new sensors (0.1% aging/noise ceiling) show that good chemical differentiation on both individual and binary VOC mixtures is achieved using the six-sensor array optimized for performance on a set of individual VOCs. These results are on the same data used for the PCA results shown in Figure 1. The results show that the performance is maintained for modest drops in film sensitivities that might occur due to SAW film aging, but that ceiling values of greater than 10% lead to large chemical identification error rates.

We show the aging effects on the individual VOCs from set A for comparison in Figure 8. The chemicals are numbered along the *x* axis, with the chemicals given in the caption. These had aging ceilings of 0.1%, 12%, and 28%. As expected, the individual VOC discriminations are far more robust than those with binary mixtures. These results directly indicate the distinct stability requirements that must be met by the sensors over time to maintain the chemical recognition ability of individual VOCs and simple VOC mixtures.

IV.c. Response to Unknown VOCs. We also examined the ability of the VERI PR to recognize array responses

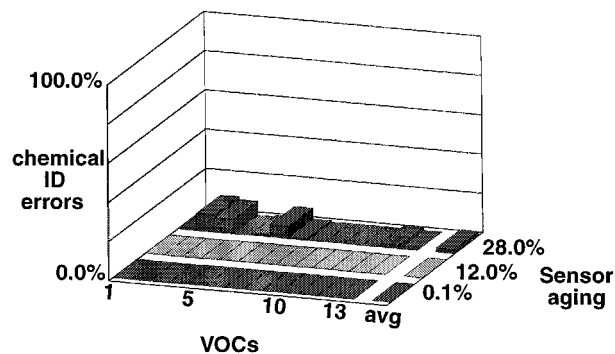


FIGURE 8. VERI leave-one-out results for one of the best six-sensor arrays, identified from the results of Figure 4, on chemical set A. Each of the chemical classes is identified by number here: (1) acetone, (2) benzene, (3) carbon tetrachloride, (4) cyclohexane, (5) DIMP, (6) DMMP, (7) water, (8) isooctane, (9) methanol, (10) pinacolyl alcohol, (11) TCE, (12) toluene, (13) 1-propanol. The results include up to 0.1% (i.e., essentially new sensors), 12%, and 28% sensitivity losses. The errors remain insignificant even up to 28% sensitivity loss for this chemical set. Detection of individual VOCs is much more robust against sensor aging or noise compared to combined detection of individual VOCs and their binary mixtures.

from chemicals outside the training set ("unknowns"). The ideal approach would be to do this with newly measured chemicals, but such data were not available at the time of this work. Our alternative was to perform what we call a leave-one-CLASS-out analysis. Each of the 28 classes of set B was held out from the training set, 1 at a time, and presented to the VERI algorithm as an unknown type of chemical. The errors are the percentage of the unknown (held out) chemical responses that were incorrectly identified as 1 of the remaining 27 target chemicals. We find that the average error across the entire set is 4%, with 21 of the held-out classes giving no errors. The good recognition performance on the remaining "known" chemical classes was also maintained. The VERI method provides effective and automatic recognition of nearly all of the unknown outliers without user threshold adjustments. Failure to detect outliers by the VERI method is not due to inappropriate class volumes, but instead is due to unknown chemicals which give responses that are too close to one of the target chemical class volumes.

We also examined the performance of the OLPARS neural net on the same six-sensor chemical training data. The neural net had 6 inputs, a maximum available number of nodes of 50 per layer for the next 2 layers, and 28 outputs. During network training by back-propagation, the error rate was observed to rapidly drop initially as a function of the number of training iterations, and then essentially saturated. Training was terminated at the beginning of the saturation portion in an attempt to avoid "overtraining". When the net was trained on the full training data, the percent correct without a reject threshold was similar to the VERI result. This confirms the general effectiveness of the VERI sensor optimization selection in achieving good separation of the chemical classes. The net was retrained for each training set with one class held out in the leave-one-CLASS-out analysis. As discussed above, we use a nonzero reject threshold on

the largest net output to achieve outlier detection. Unfortunately, the neural net recognition performance on the *known* classes simultaneously falls as the reject threshold is increased. At the highest threshold values the left-out class errors decrease to 18%, but the errors on the remaining known classes rise to 28%. At the low threshold limit, essentially all left-out class data are erroneously identified as known classes. The neural net outlier detection results are quite poor, as expected. There is no choice of a reject threshold which compensates for the unbounded class volumes of the single-hidden-layer net; hence, true multivariate detection of unknown chemicals as outliers is not available with this method via thresholding. This is one of the key performance distinctions between VERI and three-layer neural net methods.

IV.d. Concentration Estimates. Once the chemical identity of the individual VOC or simple VOC mixture has been determined by VERI, there are many alternatives for estimating the chemical concentration(s). We describe two alternatives here. A simple method is to select (in the preprocessing stage) the single, most sensitive sensor for each target chemical, and use an analytic fit to that single-sensor response to calibrate future measurements of the identified target chemical response. This simple approach is preferred when it is effective. For more difficult problems, say ternary VOC mixtures at high concentrations for which the sensor responses may be *nonadditive* functions of the three mixture component concentrations, we have used VERI directly to estimate the component concentrations. We use the list of all of the training vectors that cluster with a new input to compute averages of the mixture component concentrations of these training vectors. The result is a set of component concentration estimates, along with standard deviation estimates of the uncertainties of these estimates. We find that this approach works well on model ternary VOC mixture data that we have constructed by combining measured individual VOC responses additively and non-additively. However, the accuracy of this approach decreases rapidly with decreasing size of the training set. Much larger ($\times 10$) training sets are required to achieve 1% accuracy in the concentration estimates than are needed to identify the chemicals in our model ternary VOC study. The choice of method for estimating the chemical concentrations will necessarily be influenced by the details of the user's actual application.

V. Conclusions

We have presented the VERI method as an alternative to PCA (for sensor selection), statistical PR, and neural net

PR analyses of chemical sensor array responses for certain classes of chemical sensing applications, and we have illustrated the use of VERI and coated SAW devices to recognize a variety of individual VOCs and simple VOC mixtures while automatically rejecting unknown chemicals. Our future work will emphasize the analysis of more complex mixtures and the optimized selection of arrays from much larger candidate sets of sensors.

The HBA SAW film material was kindly provided by Jay Grate of Pacific Northwest National Laboratory. This work was sponsored by the U.S. Department of Energy under Contract DE-AC04-94AL85000. Sandia is a multiprogram laboratory operated by Sandia Corp., a Lockheed Martin Co., for the U.S. Department of Energy.

References

- (1) Gardner, J. W.; Bartlett, P. N. Pattern Recognition in Gas Sensing. In *Techniques and Mechanisms in Gas Sensing*; Moseley, P. T., Norris, J., Williams, D. E., Eds.; IOP Publishing: Bristol, England, 1991; see also references therein.
- (2) Hierlemann, A.; Schweizer-Berberich, M.; Weimar, U.; Kraus, G.; Pfau, A.; Göpel, W. Pattern Recognition and multicomponent analysis. In *Sensors Update*; Bates, H., Göpel, W., Hesse, J., Eds.; VCH: Weinheim, Germany, 1996.
- (3) Osbourn, G. C.; Martinez, R. F. Empirically defined regions of influence for clustering analyses. *Pattern Recognit.* **1995**, *28* (11), 1793.
- (4) A detailed description of VERI and a discussion of the human vision research that underlies the method can be found at our website address <http://www.sandia.gov/1100/1155Web/users.htm>.
- (5) Hand, D. J. *Discrimination and Classification*, John Wiley and Sons: New York, 1981. Also, see: Nadler, M.; Smith, E. P. *Pattern Recognition Engineering*; John Wiley and Sons: New York, 1993.
- (6) Frye, G. C.; Martin, S. J.; Cernosek, R. W.; Pfeifer, K. B. Portable acoustic wave sensor systems for on-line monitoring of volatile organics. *Int. J. Environ. Conscious Manuf.* **1992**, *1*, 37.
- (7) Ricco, A. J.; Crooks, R. M.; Osbourn, G. C. *Acc. Chem. Res.* **1998**, *31*, 289.
- (8) Frye, G. C.; Blair, D. S.; Schneider, T. W.; Mowry, C. D.; Colburn, C. W.; Donovan, R. P. Development and evaluation of on-line detection techniques for polar organics in ultrapure water. *J. Inst. Environ. Sci.* **1996**, *39*, 30.

AR970070J