

A Portable Electronic Nose For Toxic Vapor Detection, Identification, and Quantification

B.R. Linnell, R.C. Young, T.P. Griffin, B.J. Meneghelli, B.V. Peterson, K.B. Brooks
Applied Chemistry Laboratory, NASA Kennedy Space Center, Florida 32899

Abstract

A new prototype instrument based on electronic nose (e-nose) technology has demonstrated the ability to identify and quantify many vapors of interest to the Space Program at their minimum required concentrations for both single vapors and two-component vapor mixtures, and may easily be adapted to detect many other toxic vapors. To do this, it was necessary to develop algorithms to classify unknown vapors, recognize when a vapor is not any of the vapors of interest, and estimate the concentrations of the contaminants. This paper describes the design of the portable e-nose instrument, test equipment setup, test protocols, pattern recognition algorithms, concentration estimation methods, and laboratory test results.

Keywords: Electronic nose, space program, hypergolic fuel, pattern recognition, classification, quantification.

Introduction

The National Aeronautics and Space Administration (NASA) at Kennedy Space Center (KSC) has developed a portable instrument capable of eight hours of continuous operation, using a commercial off-the-shelf (COTS) electronic nose (e-nose) with a Palm Pilot for the user interface. This prototype uses algorithms developed at KSC which allow it to classify unknown vapors with a high success rate, recognize when a vapor is not one of the vapors of interest, and accurately estimate the concentration of single vapors or mixtures of two vapors. The prototype unit has completed lab testing and is being field tested as a personnel safety monitor for hypergolic fuels.

An e-nose consists of an array of non-specific vapor sensors^[1]. In general, the sensor array is designed such that each individual sensor responds to a broad range of chemicals, but with a unique sensitivity relative to the other sensors. Chemical identification is achieved by comparing the sensor response pattern of an unknown vapor to previously established patterns of known vapors. Different sensor types (metal oxide semiconductor, polymer composite, quartz microbalance, etc.) have different advantages and disadvantages – for example, some sensors are more sensitive to specific vapors, while others are less prone to drift due to changes in ambient conditions (e.g., temperature, relative humidity (RH), pressure). Many commercial e-noses have been trained to assign a quality value to flavors and food products, diagnosis certain diseases, and detect chemical spills, among other applications; however, few have been used to identify specific compounds or to quantify the concentrations of the compounds. NASA at KSC has assessed the sensitivity of several commercially available and pre-production e-noses^[2]. One very sensitive e-nose was found which is capable of identifying all of the hypergolic fuels and VOCs described below at the minimum required concentrations.

The ability to monitor air contaminants in a closed environment, such as the Shuttle, the International Space Station (ISS), and future human missions to Mars or the moon is important to ensure mission success and the safety of astronauts. Continuous air monitoring could provide notification of adverse events such as chemical spills or leaks. This has been demonstrated as important because post-mission analyses of grab air samples from the Shuttle have confirmed the presence of on-board volatile organic contaminants (VOCs)^[3]. In addition, the Space Program and military use large quantities of hypergolic fuels such as hydrazine (Hz) and monomethyl hydrazine (MMH) as rocket propellant and fuels for

auxiliary power units (APU). These substances are very toxic and are suspected human carcinogens. Hypergolic fuels can contaminate cabin air when astronauts return from extravehicular activities.

Threshold values for Hz and MMH, as well as the VOCs used in this research that are typical of the more than 60 compounds identified in Shuttle air samples so far^[3], are shown in Table 1. Current off-the-shelf portable hypergolic fuel detectors not only require 10 to 20 minutes of exposure to detect 10 ppb concentrations, but they often react to other vapors (interferants) to give false positives, making these units unacceptable for many operations.

Vapor	Abrv.	Threshold (ppm)
Hydrazine	Hz	0.01 ^A
Monomethyl hydrazine	MMH	0.01 ^A
Acetone	Ace	22 ^B
Isopropyl alcohol	IPA	60 ^B
Methylethyl Ketone	MEK	10 ^B
Toluene	Tol	16 ^B
Xylene	Xyl	50 ^B

Table I – Vapors used in this research

^AAmerican Conference of Governmental Industrial Hygienists^[4]
^B7-day Spacecraft Maximum Allowable Concentration (SMAC)^[5,6]

Methods

Generation of Calibrated Standard Vapors

Test vapors were generated using commercial vapor generators with permeation devices (PD) (Kintek Model 360, Austin TX). Vapors from the Kinteks were blended with clean air from a temperature, relative humidity, and flow rate controller (Miller-Nelson Model HCS-40, Monterey CA), which provided dilution up to a factor of 50. The resulting airstream was then drawn into the e-nose at 2 L/min. All vapor concentrations were verified on a regular basis using standard methods^[7,8,9]. Detailed information about the test setup and calibration procedures have been documented elsewhere^[1].

E-nose Evaluation

A literature and market search of available e-noses was performed to identify instruments suitable for the applications described above^[1]. Several of the e-noses tested showed reasonable sensitivity to ppm levels of MMH, Hz, and the VOCs. However, of the instruments tested, only the i-Pen (Airsense Analytics, Germany) was able to respond to 10 ppb levels of Hz and MMH with a signal to noise ratio greater than 3, and whose vendor was willing to provide the communications protocols for the development of an e-nose interface with a Palm Pilot.

Prototype Design and Fabrication

The basic Original Equipment Manufacturer (OEM) unit was fitted with an air pump for drawing samples into the sensors, an 8-hour lithium-ion battery for portability, an inlet filter for establishing a baseline reading, and a Palm Pilot to minimize size, weight, and startup time. The lithium-ion battery was selected to provide at least 8 hours of continuous operation for the lowest weight. Figure 1 shows a schematic diagram of the prototype unit.

For hypergolic fuel detection, the filter consists of glass wool soaked in a solution of 50 % sulphuric acid and 50 % water (by volume), suctioned to dampness, and dried at 60 degrees C.

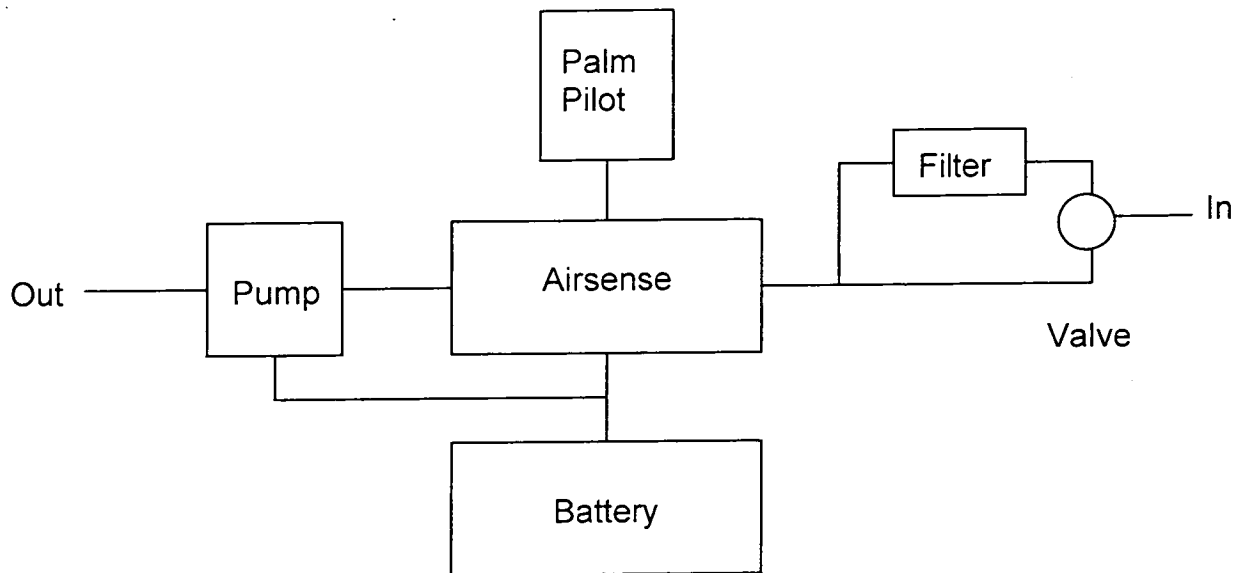


Figure 1. Prototype schematic

The user interface consists of three buttons, “Sniff”, “Classify”, and “Exit”. Upon pressing the Sniff button, a graph is displayed showing the three largest sensor responses as the data is gathered for 90 seconds. Once the sniff is completed, the Classify button runs the identification and quantification algorithms developed at KSC, and tells the user which class the vapor belongs to and its estimated concentration, or that the vapor is “unknown” (i.e., not one of the known vapors in the model). The Palm Pilot then monitors the sensors until it has determined that they have returned to their baseline values, and are ready for another sniff. This stabilization time can take from one to ten minutes, depending on the concentration of the previous exposure.

Single-Vapor Identification Algorithm

Code was written for the Palm Pilot to implement a standard quadratic classifier^[10]. Given class i with mean μ_i and covariance matrix Σ_i , an unknown sample x is assigned to that class with the smallest value of

$$(x-\mu_i)^T \Sigma_i^{-1} (x-\mu_i) + \ln(|\Sigma_i|)$$

where x^T indicates a vector transpose and $|\Sigma_i|$ indicates the determinant of Σ_i . For all of this research, the data consisted of the e-nose sensor responses 90 seconds after the start-of-sniff^[11].

After a class has been selected, the square of the Mahalanobis distance^[10] $(x-\mu_i)^T \Sigma_i^{-1} (x-\mu_i)$ from the example to the estimated class is calculated, and compared to a statistically determined threshold. If the example is too far from the assumed class, the example is rejected as an unknown vapor.

Single-Vapor Quantification Algorithm

To estimate the concentration of a single vapor, the model data was plotted to show sensor response as a function of concentration, shown in Figure 2. This data was then fitted (using the Levenberg-Marquardt nonlinear least-squares algorithm, hereafter "LM") to the formula $y=A(1 - e^{-Bx})$, to find the values for the parameters A and B which best relate the sensor response y to the vapor concentration x . This formula was determined to be appropriate because the sensor response should go to zero as the concentration goes to zero, and the sensor response should saturate at high enough concentrations. This formula was found to be the best of those tested.

When presented with an unknown sample, the inverse of this formula was then used to convert each of the ten sensor responses into ten concentration estimations. Many different techniques were explored to convert the ten estimates into a single value, including taking the mean, the mean weighted by the quality of the curve fit, and multiple linear regression, but it turned out that using the estimate of one particular sensor was usually the best. Which sensor to use depended on the vapor, and was determined by analyzing the model data.

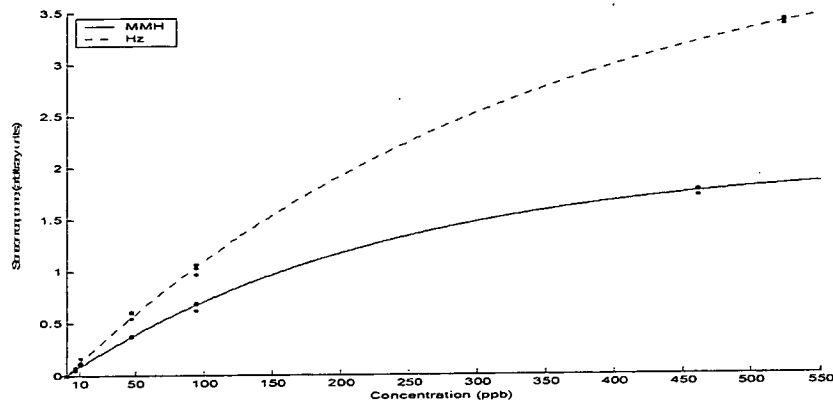


Figure 2. Concentration curve fitting (one sensor)

Note that with this approach, if the unknown vapor is misidentified, the concentration estimate will be invalid.

Two-Vapor Identification/Quantification Algorithm

In order to build a model of how the sensors respond to a mixture of two vapors, sensor responses were gathered for each vapor of interest, and for all pairs of vapors (in a 50/50 mix), across a range of known concentrations (see Figure 3 and Table VII). For each pair of vapors, the response of each sensor is plotted vs. the concentration (creating a surface, also shown in Figure 3), and an equation of the form $z=A + (B*x^C + D*y^E)^F$ is fitted to the data (again using the LM algorithm) to find the least-squares values of the parameters A,B,C,D,E,F which relate the concentrations x and y to the sensor response z . This formula has been found by other researchers^[11-12] to be the best model of sensor response to mixtures of vapors, and was the best of those tested for this research.

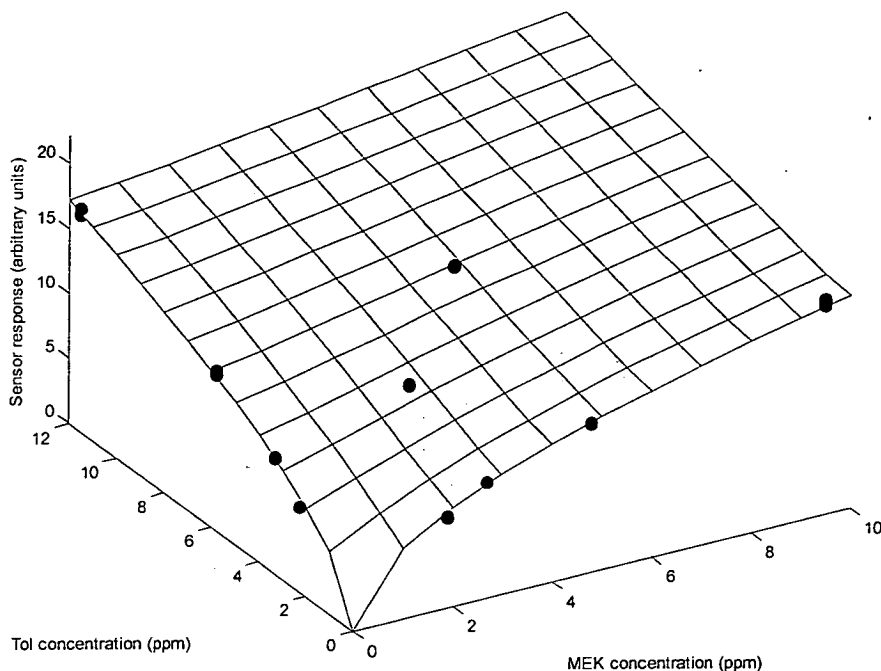


Figure 3 – Two-vapor concentration model surface for MEK+Tol, sensor #9

Given an unknown example, the first step is to calculate the difference between each sensor value of the unknown and the modeled surface for that sensor, for all possible vapor pairs (Ace+IPA, MEK+Tol, etc.). The lowest point on this error surface represents the most likely concentrations which produced the given sensor value, for that sensor and vapor pair. Next, for each vapor pair, the multiple error surfaces (one per sensor) are combined to create a cumulative error surface. The best method found so far is to simply add the error surfaces, although performance was enhanced by leaving out sensors which either showed very small responses or were saturated.

For each vapor pair's cumulative error surface, the minimum point on that surface must be found. To facilitate rapid algorithm development, finding the optimum point has been done so far by sampling values on a grid, then using the minimum value.

The "quality" of each minimum for all vapor pairs was then determined. The best metric turned out to be the value of the surface at the minimum, although various normalizations of this value were tried as well. The vapor pair whose minimum has the best quality is then selected as the identified classes. The location of the minimum within that pair's error surface determined the estimated concentrations.

For example, given an unknown which is a mixture of 5.8 ppm Tol and 4.0 ppm Xyl, the cumulative error surfaces are shown in Figure 4 for Ace+IPA and Tol+Xyl. As can be seen, the minimum point on the Tol+Xyl surface is much lower than the minimum point on the Ace+IPA surface, so Tol+Xyl would be chosen as the most likely vapor pair. Within the resolution of the grid, the minimum point on the Tol+Xyl error surface occurs at (5.81, 3.99), which would be taken to be the most likely concentrations of the vapors.

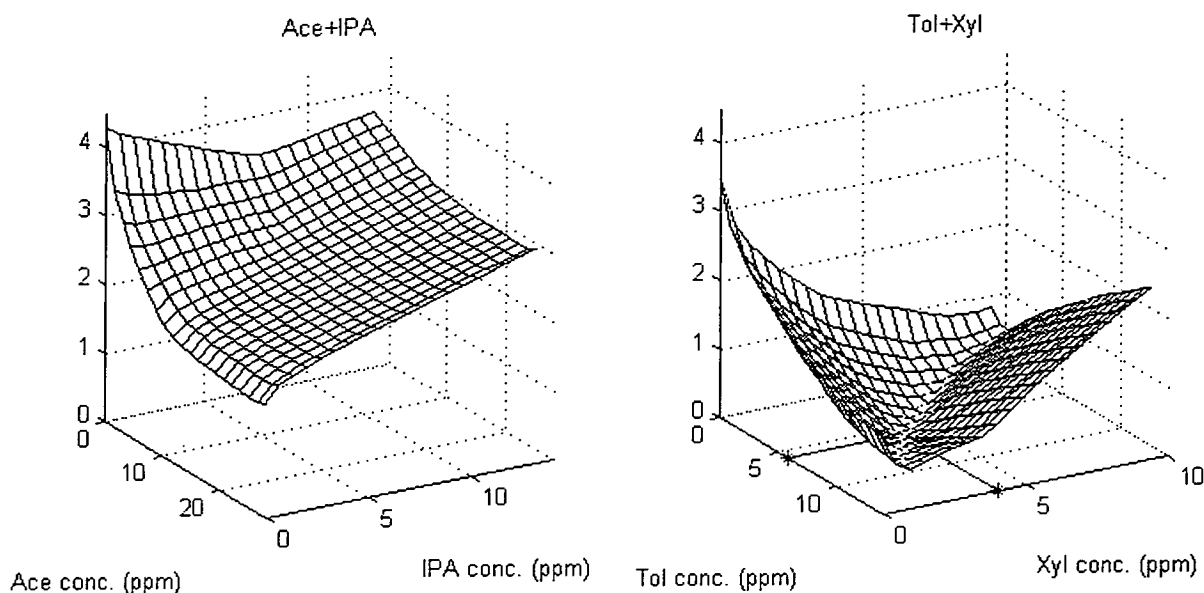


Figure 4 – Picking vapor pairs and concentration estimates from cumulative error surfaces

Training and Validation Data

The i-Pen was trained using four concentrations of each vapor at three RHs, shown in Table II. In all cases, a lab computer was used to acquire the training data for the models. Due to time constraints, mixture vapor models were only gathered for Ace+IPA, MEK+Tol, MEK+Xyl, and Tol+Xyl.

Vapor	Conc. Used (ppm)	RHs (%)
MMH	0.0072, 0.047, 0.095, 0.52	50, 70, 85
Hz	0.011, 0.048, 0.095, 0.46	50, 70, 85
Ace	4.7, 6.7, 12, 23	20, 50, 85
IPA	2.2, 3.2, 5.6, 11	20, 50, 85
MEK	1.9, 2.7, 4.8, 9.5	20, 50, 85
Tol	2.3, 3.3, 5.8, 11	20, 50, 85
Xyl	1.6, 2.3, 4.0, 7.9	20, 50, 85

Table II – Model Vapor Concentrations and Humidities

For the hypergolic fuels, validation data using Hz and MMH vapors at various concentrations and RHs (to be described in the Results section) were then gathered using the Palm Pilot, as it would be used in the field.

A lab computer was used to acquire the VOC validation data, which also spanned a range of concentrations and RHs (also described in Results). However, the VOC validation data was gathered some time after the original model data, and as will be seen in the Results section, may have been significantly different than the model data due to sensor “drift”, a common problem with e-noses.

Results

Hypergolic Fuels, Single-Vapor

In order to determine how well the classes can be differentiated, some estimate of the classifier's future performance is required^[13].

If all the data is used to build the model, and is then also used to estimate the success rate (known as "Resubstitution", abbreviated "Resub"), the estimate will generally be too optimistic. This problem is usually solved by using techniques such as "N-Fold Cross Validation", which sets aside part of the data, builds a model with the remaining data, and uses the first part to estimate the performance. N different portions are set aside, and the N estimates are then averaged (N=3-10 is typical). If N equals the number of examples, the result is the "Leave-One-Out" estimator (LOO), which is often called simply "cross validation". However, when there is a relatively large amount of data, all estimation methods will return approximately the same value^[14].

The same methods were used to calculate the average absolute percent relative error in quantification ($100 * |estimate - true| / true$, hereafter "%error"); however when estimating the concentration it is assumed that the vapor is correctly identified. To avoid artificially inflated error rates, two values are shown in Table III. The first is the mean %error for all concentrations greater than 10 ppb, and the second is the mean absolute error (in ppb) for the 10 ppb examples. This is because errors of just a few ppb have very large %errors at 10 ppb, which is at the limit of detection for the e-nose. These results are across all vapors, concentrations, and humidities. Figure 5 shows the %error for each concentration estimate.

Result	Resub	LOO
Classification success	99 %	98 %
Quantification error	3.5 %, 2.5 ppb	4.1 %, 2.5 ppb

Table III – Hypergolic single vapor estimated performance

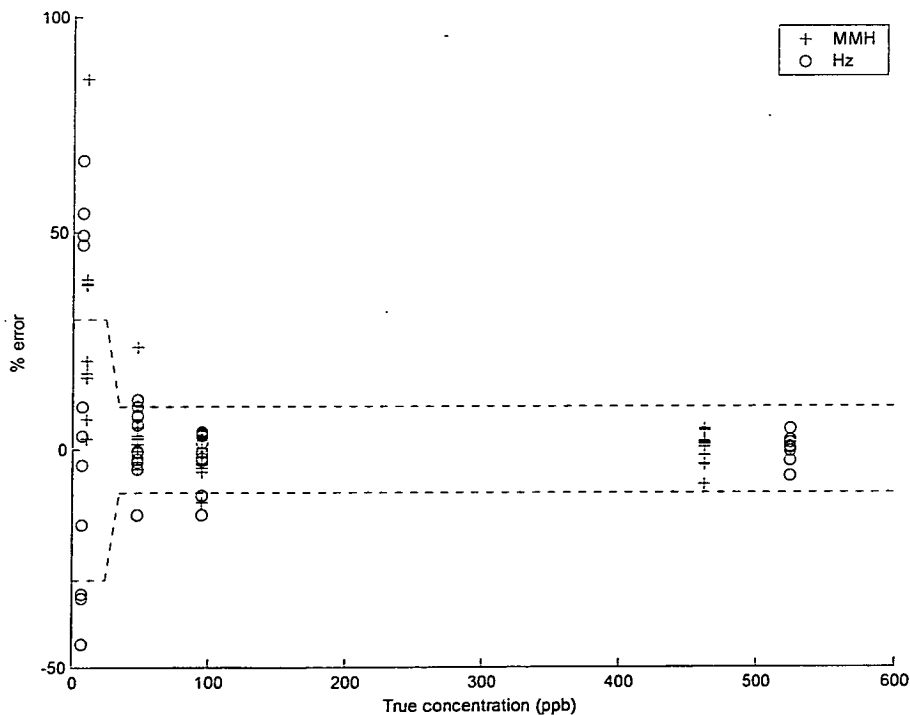


Figure 5 – Hypergolic single vapor concentration estimation results (LOO)

All validation tests were performed at 70 % RH. Table IV shows the summary statistics for each vapor (MMH, Hz) and concentration, and Figure 6 shows the concentration estimation errors. The mean concentration errors shown are only for those examples that were correctly classified, and are again separated into 10 ppb and greater-than-10-ppb values. Note that the 250 ppb vapors were not part of the training data, but are accurately identified and quantified.

Vapor	Std. Conc. (ppb)	% correctly identified	Mean %error or error
Hz	7.2	50	0.4 ppb
Hz	95	100	5.5 %
Hz	250	100	1.3 %
Hz	524	100	2.8 %
MMH	11	83	2.6 ppb
MMH	95	100	10 %
MMH	250	100	9 %
MMH	461	100	1.3 %
Overall	-----	92	5.0 %, 1.8 ppb

Table IV – Hypergolic single vapor validation performance

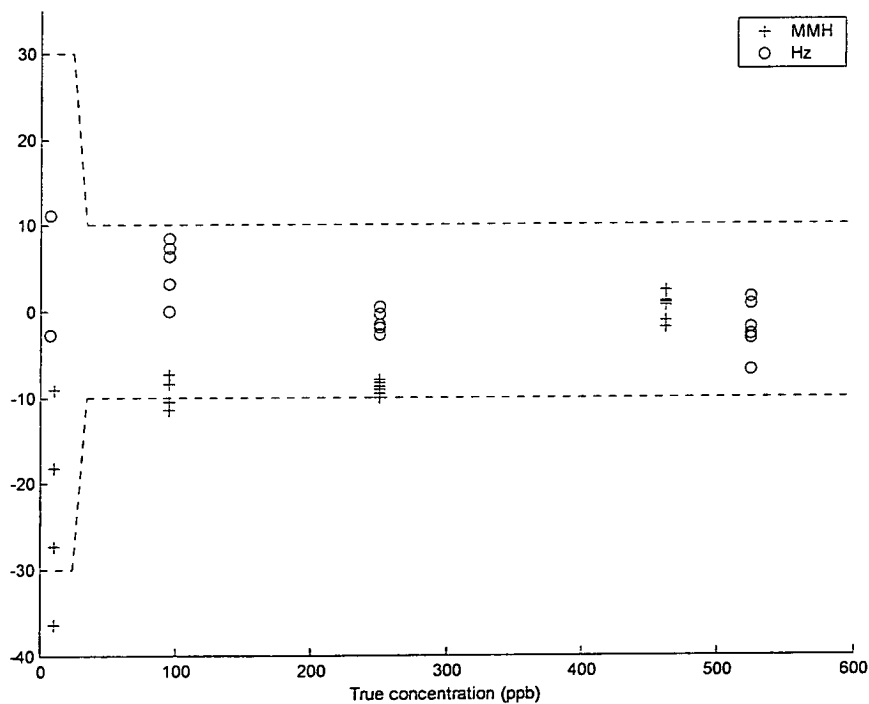


Figure 6 – Hypergolic single vapor concentration validation results

Organic, Single-Vapor

The results are shown in Table V, across all vapors, concentrations, and humidities. When estimating the concentration, it is assumed that the vapor has been correctly identified. For this and all subsequent VOC results, only the mean %error is given since none of the lowest concentrations were near the limit of detection. Figure 7 shows the concentration estimation errors.

Result	Resub	LOO
Classification success	100 %	100 %
Quantification error	5.3 %	5.9 %

Table V – Organic single vapor estimated performance

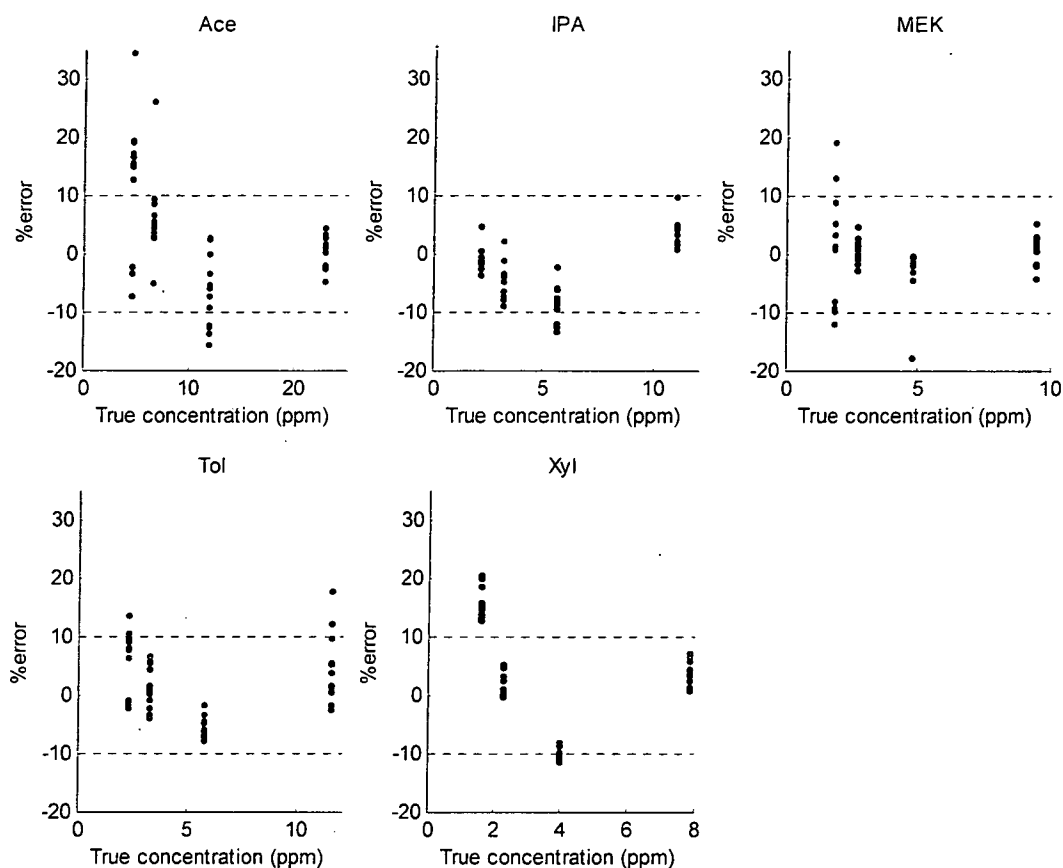


Figure 7 – Organic single vapor concentration estimation results (LOO)

A wide but not exhaustive combination of concentrations and humidities were gathered for validation. Table VI shows the summary statistics for each vapor and concentration, and Figure 8 shows the %errors in concentration estimation. When estimating the concentration, it is assumed that the vapor has been correctly identified.

Vapor	Std. Conc. (ppm)	% correctly identified	Mean %error
Ace	4.7	100	12.7
Ace	6.7	100	7.2
Ace	12	100	4.0
Ace	23	100	1.6
IPA	2.2	100	36.3
IPA	3.2	83	27.1
IPA	5.6	100	25.4

IPA	11	100	15.5
MEK	1.9	100	36.8
MEK	2.7	100	25.8
MEK	4.8	100	32.6
MEK	9.5	100	27.9
Tol	2.3	100	4.3
Tol	3.3	100	1.8
Tol	5.8	100	4.6
Tol	11.5	100	9.5
Xyl	1.6	0	6.3
Xyl	2.3	50	7.5
Xyl	4.0	100	4.1
Xyl	7.9	50	15.1
Overall	-----	91	14.8

Table VI – Organic single vapor validation performance

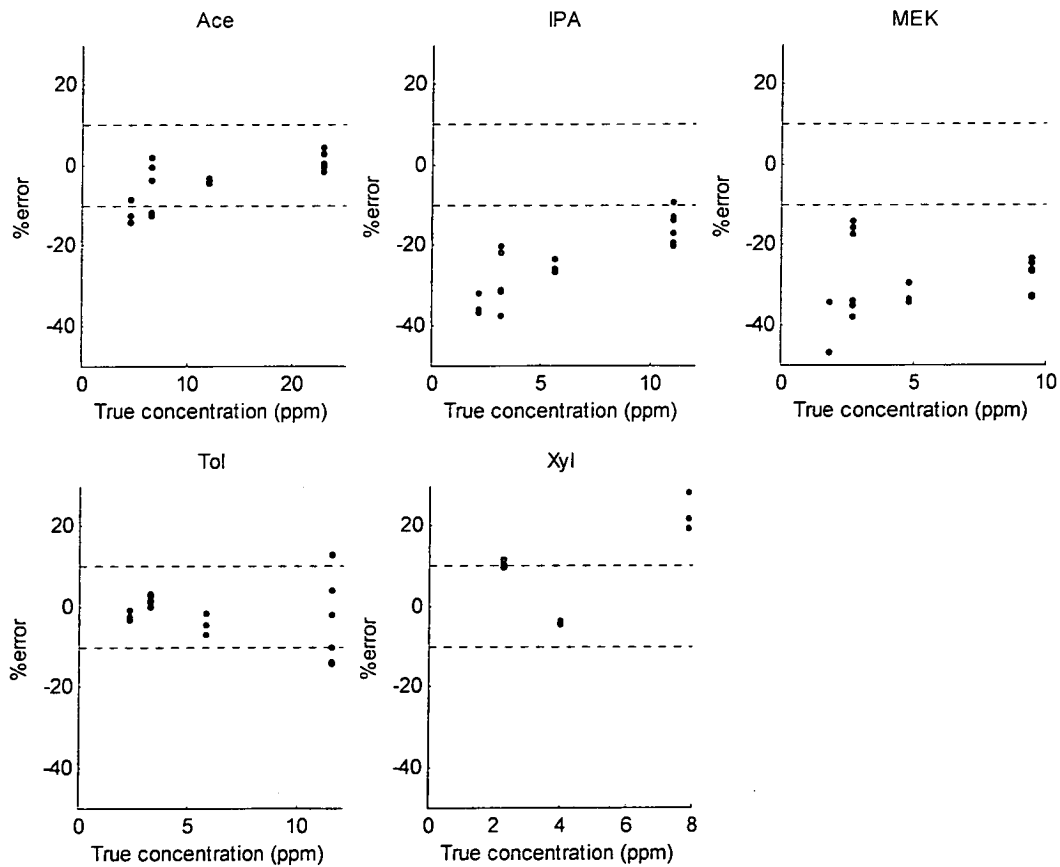


Figure 8 – Organic single vapor concentration validation results

A question arose as to whether the poorer results of the validation data could be due to sensor drift. While the validation data was started only two weeks after the model data was finished, the success rate in classification shown in Table VI is deceiving, because the Mahalanobis distance is not taken into account. When that is done, it turns out that every example in the validation set is considered different

enough from the model to be classified as an “unknown”. Therefore the validation set might not be similar enough to the model for the results in Table VI to be considered relevant.

Organic, Two-Vapor

All examples were taken at 50 % RH. Table VII shows the summary results, where “Both Classified” indicates the percentage of the unknowns for which both vapors were correctly identified. For single-vapor unknowns, “Both Classified” indicates that the second vapor’s concentration was determined to be zero. In all cases, “Both Classified” means the unknown was identified completely correctly. Unlike the previous estimation results, these results are given by vapor and concentration to show trends better. The concentration errors are shown in Figure 9.

Because the LM algorithm took significantly longer to process the two-variable surface equation, Leave-One-Out testing was not performed. Resubstitution was determined to be an acceptable estimate here, because there were 52 examples used to create each surface, and the removal of any one of them would not significantly alter the resulting model parameters of the surface. As can be seen in the other estimated performance results above (Tables III and V), Resubstitution and Leave-One-Out give very similar results when there are a large number of examples.

Vapor #1	Conc. (ppm)	Vapor #2	Conc. (ppm)	Both Classified (%)	Mean %error
Ace	4.7			100	23.6
Ace	6.7			100	8.4
Ace	12.0			100	7.5
Ace	23.0			100	5.8
IPA	2.2			100	21.1
IPA	3.2			100	9.7
IPA	5.6			100	3.9
IPA	11.0			100	2.1
MEK	1.9			100	5.3
MEK	2.7			100	11.1
MEK	4.8			100	9.4
MEK	9.6			100	4.2
Tol	2.3			100	13.2
Tol	3.3			100	10.0
Tol	5.8			100	12.7
Tol	11.5			100	5.8
Xyl	1.6			100	6.4
Xyl	2.3			100	8.5
Xyl	4.0			100	12.4
Xyl	7.9			100	1.1
Ace	4.7	IPA	2.2	100	6.3
Ace	6.7	IPA	3.2	100	10.9
Ace	12.0	IPA	5.6	100	18.1
Ace	23.0	IPA	11.0	100	8.7
MEK	1.9	Tol	2.3	no data	no data
MEK	2.7	Tol	3.3	100	13.3
MEK	4.8	Tol	5.8	100	12.6

MEK	9.6	Tol	11.5	100	14.2
MEK	1.9	Xyl	1.6	100	15.0
MEK	2.7	Xyl	2.3	100	26.3
MEK	4.8	Xyl	4.0	100	12.4
MEK	9.6	Xyl	7.9	100	14.2
Tol	2.3	Xyl	1.6	100	9.8
Tol	3.3	Xyl	2.3	100	11.6
Tol	5.8	Xyl	4.0	100	4.9
Tol	11.5	Xyl	7.9	100	21.3
Overall	----	----	----	100	10.2

Table VII – Organic Two Vapor Estimated Performance (Resubstitution)

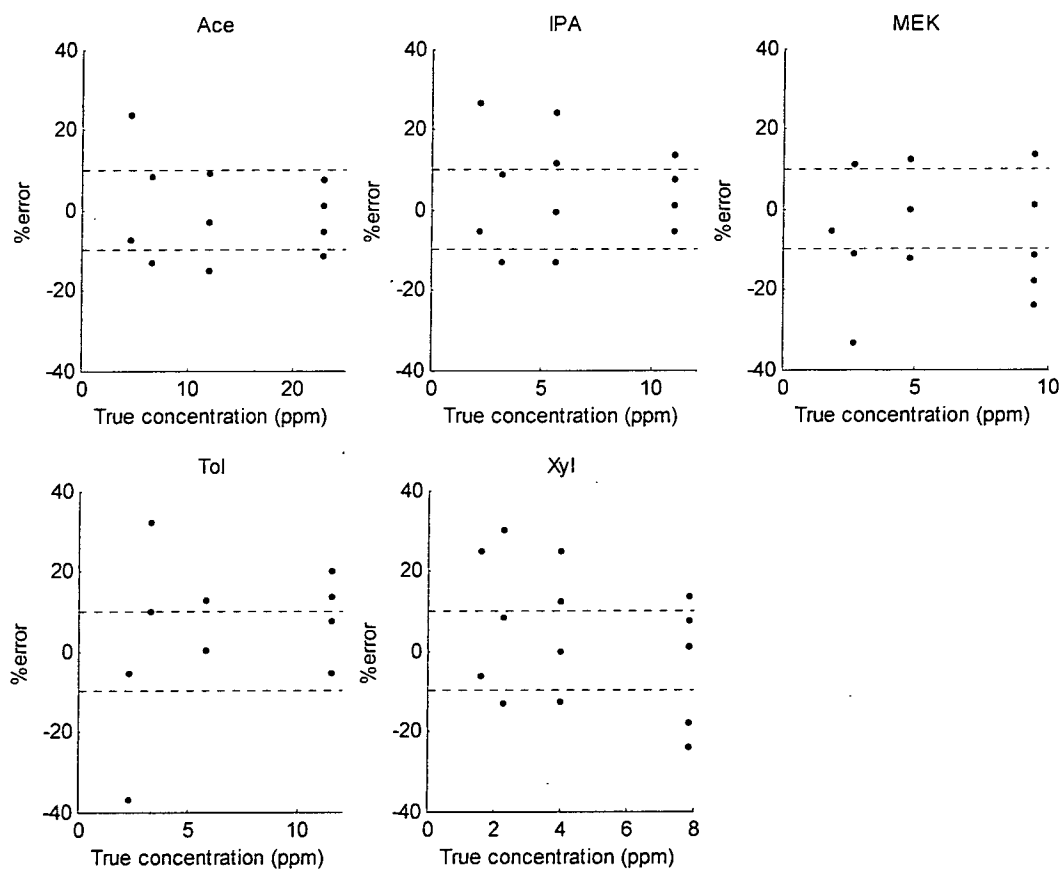


Figure 9 – Organic Two Vapor Concentration Estimated Performance (Resubstitution)

Validation data was gathered at concentrations that were not included in the model, for both single-vapor and two-vapor mixtures. All examples were taken at 50 % RH. Results for each example are summarized in Table VIII, where “One Classified” indicates that one vapor was correctly identified and one was misidentified, and “None Classified” means that both vapors were misidentified. For single-vapor unknowns, “One Classified” means either that the true vapor was misidentified, or the second vapor’s estimated concentration was non-zero. The %error was calculated only for those non-zero vapors which were correctly identified. Figure 10 shows the concentration estimation errors.

Vapor #1	Conc. (ppm)	Vapor #2	Conc. (ppm)	Both Classified (%)	One Classified (%)	None Classified (%)	%error
Ace	13.2			0	100	0	55.8
Ace	17.5			0	100	0	40.5
IPA	6.3			100	0	0	15.3
IPA	8.3			100	0	0	4.8
MEK	5.4			0	100	0	49.5
MEK	7.2			0	100	0	42.9
Tol	6.5			0	100	0	49.5
Tol	8.7			0	100	0	45.3
Xyl	4.5			0	100	0	46.3
Xyl	6			0	100	0	33.4
Ace	13.2	IPA	6.3	100	0	0	42.6
Ace	17.5	IPA	8.3	100	0	0	35.7
MEK	5.4	Tol	6.5	0	100	0	5.3
MEK	7.2	Tol	8.7	0	100	0	9.7
MEK	5.4	Xyl	4.5	100	0	0	17.9
MEK	7.2	Xyl	6	100	0	0	9.6
Tol	6.5	Xyl	4.5	100	0	0	26.8
Tol	8.7	Xyl	6	100	0	0	21.5
Overall	----	----	----	31	69	0	31.7

Table VIII – Organic Two Vapor Validation Performance

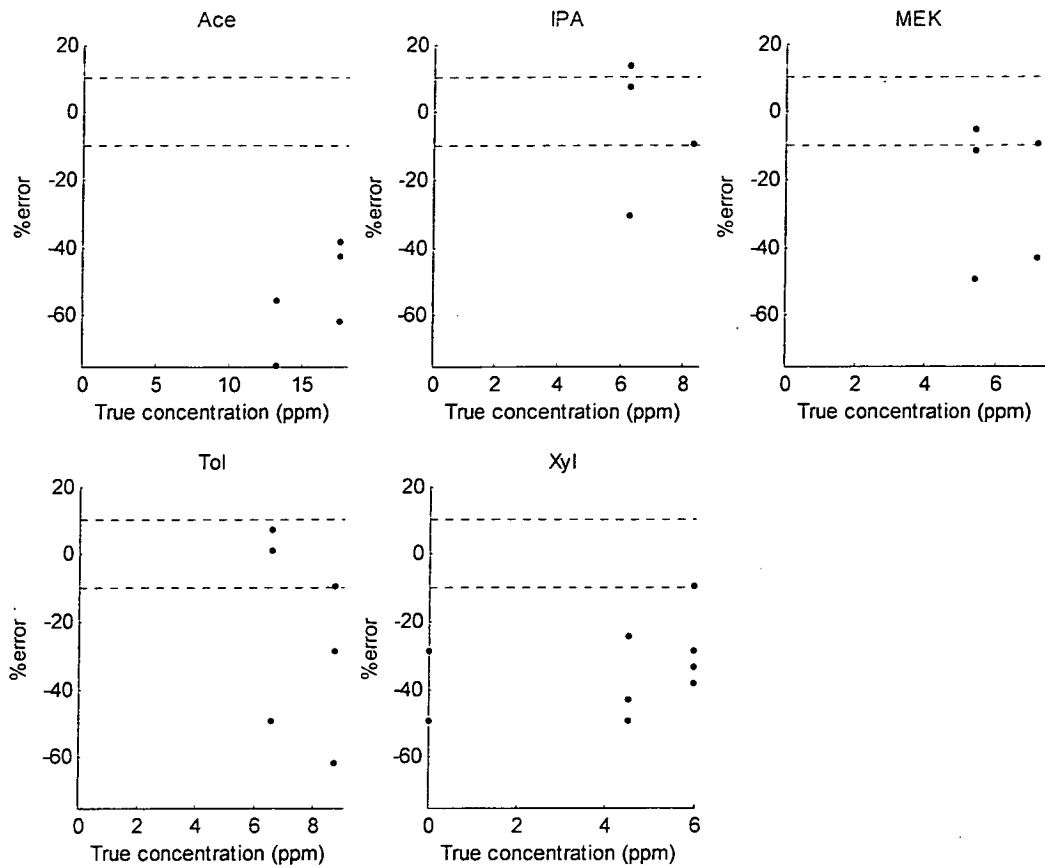


Figure 10 - Organic Two Vapor Concentration Validation Performance

This data was taken three to four months after the original model data. As with the single-vapor VOC validation data, all of the validation examples used here (single and mixed) were significantly different enough from the model data to be considered “unknowns” when the Mahalanobis distance is included, and so these results may not be relevant.

Resistance to False Positives

Four samples were taken of each of the organic vapors at the minimum concentrations listed in Table II, and classified using the hypergolic fuel training samples as the model. As can be seen in Table IX, the organic vapors are identified as not belonging to either MMH or Hz because the Mahalanobis distance is too large (vapors with distances greater than 5.1 are considered “unknown”).

Vapor	Mahalanobis distances
MMH	1.4 - 2.9
Hz	0.8 - 3.8
Ace	328.2 - 365.7
IPA	41.7 - 50.0
MEK	157.0 - 164.1
Tol	178.6 - 193.6

Xyl	572.6 - 633.4
-----	---------------

Table IX – False positive results

Discussion

Single Vapors

In all cases, the estimated performance was slightly better than the validation performance. The estimated classification success rate was 98-100 %, and the validation success rate was 91-92 %. Excluding the 10 ppb hypergols, the %error in concentration estimation was about 3-6 % except for the validated organics, where it was 15 %. However, as discussed above, the relevance of the VOC validation data is very questionable.

Other studies^[15-16] which have attempted to quantify single-vapor e-nose data have shown poor performance in identification and/or quantification. One reason for our good results is due to our method of using a very accurate standard statistical pattern recognition technique to first identify the vapor, whereas other studies used non-standard methods to identify the vapor. Another improvement comes from the choice of formula the concentration data is fitted to. Other studies used cubic splines or polynomials, which do not model the actual underlying trend of the data very well.

Mixture Vapors

For the mixtures, the validation results differed significantly from the estimated results. The estimated results show a 100 % success in identifying both vapors, with an average %error of 10 %, while the validation results show only a 31 % chance of identifying both vapors, with an average %error of 32 %. However, as discussed above, the VOC mixture validation data is significantly different than the model data, and is probably not relevant.

Previous studies^[17-21] which have attempted to quantify e-nose data for vapor mixtures have had similar or better accuracy in quantification, but they all started with the assumption that the identity of the vapors in the mixture was already known. The algorithm presented here has the unique ability to identify the mixtures as well as quantify them.

Future Research

Future enhancements include using more sophisticated searches of the two-vapor cumulative error surface such as gradient descent, which has been determined to be optimal since the error surface has only one minimum. Studies should be made to see if the two-vapor algorithm can realistically be scaled to three or more vapors. Sensor drift needs to be studied, as well as methods for minimizing its effects.

Conclusions

A prototype portable e-nose capable of detecting 10 ppm MMH and Hz has been developed at KSC NASA. It is capable of detecting, identifying, and quantifying vapors in only 90 seconds, with a 1 to 10 minute recovery time, depending on the concentration of the exposure. This unit classifies single vapors with 90-100 % accuracy, and quantifies them with an average of about 5 % error, except at the limit of detection (10 ppb), where the error is less than 3 ppb. The prototype also shows excellent resistance to false positives, and may be trained to detect, identify, and quantify virtually any vapor of interest, within the detection limits of the sensors. It also shows great promise in being able to accurately identify and quantify mixtures of two vapors and possibly more.

References

- (1) H. Nagle, R. Gutierrez-Osuna, S. Schiffman, "The How and Why of Electronic Noses", IEEE Spectrum, September 1998, pp. 22-34.
- (2) R. Young, W. Buttner, B. Linnell, "Electronic Nose for Space Program Applications", Sensors and Actuators B, Proceedings of the Ninth International Meeting on Chemical Sensors, Part II, IMCS-9, Boston MA, p.7-16, July 2002.
- (3) J. James, T. Linero, H. Leano, J. Boyd, P. Covington, Volatile Organic Contaminants Found in the Habitable Environment of the Space Shuttle: STS-26 to STS-55, Aviation, Space, and Environmental Medicine, September 1994, pp. 851-857.
- (4) ACGIH, *Documentation of the Threshold Limit Values and Biological Exposure Indices*, 5th ed, American Conference of Governmental Industrial Hygienists, Cincinnati, Ohio, 2002.
- (5) Spacecraft Maximum Allowable Concentrations for Airborne Contaminants, JSC20584, NASA Johnson Space Center, May 1990.
- (6) Safety Requirements Document for International Space Station Program, SSP50021, NASA Johnson Space Center, December 12, 1995.
- (7) J. Wyatt, S. Rose-Pehrsson, T. Cecil, K. Crossman, N. Mehta, R. Young, "Coulometric Method for the Quantification of Low-Level Concentrations of Hydrazine and Monomethylhydrazine", American Industrial Hygiene Association Journal, June 1993, pp. 285-292
- (8) "Determination of Concentrations of N₂H₄ or MMH Vapor in Nitrogen or Air by the Coulometric Titration Method", Applied Chemistry Laboratory, Kennedy Space Center, internal document.
- (9) P. Taffe, S. Rose-Pehrsson, J. Wyatt, "Material Compatibility with Threshold Limit Value Levels of Monomethyl Hydrazine", NRL Memorandum Report 6291, October 1988.
- (10) R.O. Duda and P.E. Hart, "Pattern Classification and Scene Analysis", John Wiley & Sons, New York, 1973.
- (11) L. Lundstrom, "Hydrogen-sensitive MOS-structures, Part I : Principles and Applications", Sensors and Actuators, Vol. 1, p. 403-426, 1981.
- (12) P.K. Clifford and D.T. Tuma, "Characteristics of Semiconductor Gas Sensors. I : Steady-state Gas Response", Sensors and Actuators, Vol. 3, p. 233-254, 1982/1983.
- (13) G. Toussaint, "Bibliography on Estimation of Misclassification", IEEE Transactions on Information Theory, Vol. 20, pp. 472-479, 1974.
- (14) N. Glick, "Additive Estimators for Probabilities of Correct Classification", Pattern Recognition, Vol. 10, pp. 211-222, 1978.
- (15) M.A. Ryan, H. Zhou, M.G. Buehler, K.S. Manatt, V.S. Mowrey, S.P. Jackson, et al., "Monitoring Space Shuttle Air Quality Using the Jet Propulsion Laboratory Electronic Nose", IEEE Sensors Journal, Vol. 4, No. 3, p. 337-347, June 2004.

(16) L. Carmel, N. Sever, D. Lancet, D. Harel, "An eNose Algorithm for Identifying Chemicals and Determining Their Concentration", *Sensors and Actuators B*, 93, p. 77-83, 2003.

(17) H. Zhou, M.A. Ryan, M.L. Homer, "Nonlinear Least Squares Based Algorithm for Identifying and Quantifying Single and Mixed Air Contaminants with JPL's Electronic Nose", to appear in *IEEE Sensors*.

(18) W.P. Carey, S.S. Yee, "Calibration of Nonlinear Solid-State Sensor Arrays Using Multivariate Regression Techniques", *Sensors and Actuators B*, 9, p. 113-122, 1992.

(19) G. Huyberegts, P Szecowka, J Roggen, B.W. Licznerski, "Simultaneous Quantification of Carbon Monoxide and Methane in Humid Air Using a Sensor Array and an Artificial Neural Network", *Sensors and Actuators B*, 45, p. 123-130, 1997.

(20) G. Faglia, F. Bicelli, G Sberveglieri, P Maffezzoni, P. Gubian, "Identification and Quantification of Methane and Ethyl Alcohol in an Environment at Variable Humidity by an Hybrid Array", *Sensors and Actuators B*, 44, p. 517-520, 1997.

(21) B.W. Jervis, J. Desfieus, J. Jimenez, D. Martinez, "Quantification of Gas Concentrations in Mixtures of Known Gasses Using an Array of Different Tin-Oxide Sensors", *IEE Proceedings-Science, Measurement, and Technology*, Vol. 150, No. 3, p. 97-106, May 2003.

Acknowledgements

Funding for these projects was provided by NASA, under the University Affiliated Spaceport Technology Development Contract.