6-1-2005

Gas Plume Species Identification in Airborne LWIR Imagery Using Constrained Stepwise Regression Analyses

David Pogorzala
Rochester Institute of Technology

David W. Messinger
Rochester Institute of Technology

Carl Salvaggio
Rochester Institute of Technology

John Schott
Rochester Institute of Technology

Recommended Citation

David Pogorzala, David Messinger, Carl Salvaggio, John Schott, "Gas plume species identification in airborne LWIR imagery using constrained stepwise regression analyses", Proc. SPIE 5806, Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XI, (1 June 2005); doi: 10.1117/12.603661; https://doi.org/10.1117/12.603661

Gas plume species identification in airborne LWIR imagery using constrained stepwise regression analyses

David Pogorzala, David Messinger, Carl Salvaggio, John Schott

Digital Imaging and Remote Sensing Laboratory
Chester F. Carlson Center for Imaging Science
Rochester Institute of Technology

ABSTRACT

Identification of constituent gases in effluent plumes is performed using linear least-squares regression techniques. Airborne thermal hyperspectral imagery is used for this study. Synthetic imagery is employed as the test-case for algorithm development. Synthetic images are generated by the Digital Imaging and Remote Sensing Image Generation (DIRSIG) Model. The use of synthetic data provides a direct measure of the success of the algorithm through the comparison with truth map outputs. In image test-cases, plumes emanating from factory stacks will have been identified using a separate detection algorithm. The gas identification algorithm being developed in this work is performed only on pixels having been determined to contain the plume. Constrained stepwise linear regression is used in this study. Results indicate that the ability of the algorithm to correctly identify plume gases is directly related to the concentration of the gas. Previous concerns that the algorithm is hindered by spectral overlap were eliminated through the use of constraints on the regression.

Keywords: gaseous plumes, hyperspectral, stepwise regression, target identification

1. INTRODUCTION

In recent years, thermal hyperspectral imagery has been employed in the detection, identification and quantification of gasesous plumes.[1–4] This trend has been driven by growing concerns in environmental monitoring and new advances in sensing technology. Many of the investigations into this problem use image processing techniques commonly applied in the visible (VIS) and near infrared (NIR) portion of the electromagnetic spectrum on longwave infrared (LWIR) imagery. Many of these studies were successful in their ability to detect plume pixels in a scene. A drawback of these studies are that all of the techniques employed required prior knowledge of which gases were present in the scene.

Marinelli, *et al.* (2000) identified two separate gases being released from a small stack on the ground in front of a warm building. They showed that a Spectral Matched Filter will detect the plume in the scene. The process requires the user to define the wavelengths of interest based on the location of features of the gas in question. Lisowski and Cook (1996) use normalized band differencing to detect SO_2 in aerial imagery. A user is required since the background is normalized to the absorption peak of SO_2 before subtraction. Harig, *et al.* (2002) compute a correlation coefficient between an atmosphere-subtracted pixel spectrum and the library absorption spectrum of a gas in specified spectral windows. This notion of using a subset of spectral windows for each material (*i.e.,* gas species) was established by Clark, *et al.* (1990). This is an attractive possibility for the gas problem since absorption features of gases are pronounced and often distinct from one gas to another. Young (2002) performed gas detection using a spectral matched filter and gas identification using both constrained and unconstrained linear regression. The basis vector set in his regression model is restricted to five gases, all at a single temperature. These gases were known to be the effluents from a refining plant in the scene. Young (2002) is successful at differentiating between plumes containing different gases using this method.

This work attempts to use stepwise linear regression techniques to identify the gases in *pre-detected* plumes in a synthetic hyperspectral image. The resulting algorithm is meant to run without user interaction and with

Further information: Chester F. Carlson Center for Imaging Science, Rochester Institute of Technology, 54 Lomb Memorial Dr., Rochester, NY 14623, pogo@cis.rit.edu

Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XI, edited by Sylvia S. Shen, Paul E. Lewis, Proceedings of SPIE Vol. 5806 (SPIE, Bellingham, WA, 2005) · 0277-786X/05/$15 · doi: 10.1117/12.603661

no prior knowledge of which gases are likely to be present in the scene. Standard matrix regression is not implemented since it forces a solution for each basis vector in the model regardless of whether or not that gas is present in the plume. Rather, stepwise regression, an iterative approach that eliminates needless basis vectors from the model, is used. Stepwise regression has been shown to be successful at selecting the appropriate basis vectors (endmembers) in unmixing hyperspectral VIS/NIR pixels.[6] Constraints are placed on the regression in this study and will be discussed in detail below.

2. METHODOLOGY

2.1. Data

Synthetically generated imagery is the source of data for this study. Data generated by the Digital Imaging and Remote Sensing Image Generation (DIRSIG) model is attractive since the user has control over the types and quantities of gases being released by plumes in the scene. This allows for images to be generated that are inherently "ground-truthed". The user populates the scene with three-dimensional objects generated in CAD design software, and attributes these objects with field-measured properties such as spectral reflectance, thermal conductivity and spectral emissivity. Atmospheric effects are modeled with MODTRAN. For more information on DIRSIG see Schott, *et al.* (1999).[7]

The plume model within DIRSIG is the Jet Propulsion Laboratory (JPL) Plume Model.[8] Spatially, the plume is constructed by modifying the gas concentration. In the downwind direction the concentration follows an exponential decay. In the across-track direction, as well as across the z-axis of the plume, the concentration follows a Gaussian distribution. The width of the distribution increases as the plume travels further downwind. The result represents a time average of a plume rather than an instantaneous snapshot of a dynamic plume of gas. The JPL model gives depth to the plume by defining it as a cloud of points, each carrying with it the temperature and concentration of the gas at that location. When DIRSIG traces a ray from the sensor down through the plume, it encounters these points and integrates the radiative transfer along the path of the ray. This approach can reproduce phenomenology that a single slab model can not.

The DIRSIG model was used to create a test image. The scene is of a region in northern Rochester, NY consisting of a high school, athletic fields, and several surrounding blocks of a residential neighborhood. The high school has been modified by the addition of two factory-type stacks releasing two separate plumes. One plume contains NH_3 gas and the other freon-114. The stacks were placed such that the two plumes do not overlap at any point downwind. The presence of hundreds of homes, trees, and numerous roads provide a large amount of background clutter in the scene. The spectral response of the simulation was chosen to model the SEBASS (Spectrally Enhanced Broadband Array Spectrometer System) sensor.[9] A single band at $10.73\mu m$ from the plume image is shown in Figure 1. The two plumes are clearly seen moving from the middle of the scene to the lower-right corner.

Three copies of this scene were produced. All three images are constructed with the same objects, atmosphere and other properties. The only difference is the release rate of the plumes; 50 g/s for Case 1, 2.5 g/s for Case 2, and 0.25 g/s for Case 3. The images measure X lines by Y samples by $J = 128$ bands, from $7.5\mu m$ to $13.6\mu m$. The values for X and Y vary across the three scenes. Case 1, measures 256×256, Case 2, is 128×128, and Case 3, is 200×200. The change in dimensions is a result of a desire to reduce algorithm run-times.

The gas absorption spectra used to generate the scene were taken from a standard library. Thirty-one gases were used, most of which were measured at three temperatures, 5, 25 and 50° C. This resulted in 88 total library spectra. Both plumes in the scene use the 25° spectra. Thirty-one gases from this library are also used as the basis vector set in the stepwise regression discussed in Section 2.3.

2.2. Radiance Model

The standard model for a "background" pixel in the LWIR spectral region, where the pixel radiance, $L(\lambda)$, is a linear combination of C background materials, is expressed as,

$$L_s(\lambda) = \left(\sum_{j=1}^{C} \beta_j \, \epsilon_j(\lambda) \, B(\lambda, T_s) \right) \tau_{atm}(\lambda) \; + \; L_u(\lambda), \tag{1}$$

Figure 1. Single-band image of synthetic DIRSIG data at $10.73\mu m$. Both plumes can be seen moving from the middle of the scene to the upper-right corner. The NH_3 plume is the left-most plume while the freon-114 plume is on the right.

where β_j is the mixing coefficient for the j^{th} background material, $\epsilon_j(\lambda)$ is the spectral emissivity of the j^{th} background material, $B(\lambda, T_s)$ is the Planck radiance function evaluated at the surface temperature, T_s, τ_{atm} is the atmospheric transmission and L_u is the upwelled radiance. Reflected downwelled radiance has been excluded from Equation (1) due to the fact that many materials have a high, uniform emissivity in the LWIR.[10] This translates into a low, uniform reflectivity that does not have a significant impact on the sensor-reaching radiance.

When writing an expression for the radiance of a pixel containing a plume (*i.e.,* a "target" pixel) we can make several approximations. We assume that the plume self-radiation does not significantly reflect off of the ground and back up to the sensor. We neglect any scattering by the plume and any entrainment of ambient air within the plume. We also assume that all of the individual gases within the plume are at the same temperature. These assumptions allow us to write the sensor-reaching radiance for a plume pixel as

$$L_p(\lambda) = \left(\left[\sum_{j=1}^{C} \beta_j \, \epsilon_j(\lambda) \, B(\lambda, T_s) \right] \, \tau_p(\lambda) \, + \, \epsilon_p(\lambda) \, B(\lambda, T_p) \right) \, \tau_{atm}(\lambda) \, + \, L_u(\lambda), \tag{2}$$

where the subscript p denotes quantities related to the plume. Equation (2) shows that the self-emitted radiance from the background is attenuated by the plume. Both this radiance and the self-emitted radiance from the plume are then attenuated by the atmosphere.

Recalling Kirchoff's Law, we can equate the absorption of a gas to its emission as

$$\epsilon_p(\lambda) = \sum_{i=1}^{D} c_i \, k_i(\lambda) \tag{3}$$

where c_i is the column density and k_i is the absorption spectrum of gas i, and the sum is over a total of D possible gases in the plume. In addition, conservation of energy means that if the gas is optically thin, we can

set $\tau = 1 - \epsilon$. Substituting these changes into Equation (2) results in the complete model for a pixel containing multiple background materials and multiple gases as

$$L(\lambda) = \left(\left[\sum_{j=1}^{C} \beta_j \, \epsilon_j(\lambda) \, B(\lambda, T_s) \right] \left[1 - \sum_{i=1}^{D} c_i \, k_i(\lambda) \right] \; + \; \left[\sum_{i=1}^{D} c_i \, k_i(\lambda) \right] B(\lambda, T_p) \right) \tau_{atm}(\lambda) \; + \; L_u(\lambda). \quad (4)$$

From this point on, we will drop the atmospheric terms from the radiance model for simplicity. This does not indicate that atmospheric compensation has been done to the data. For these studies, the images have not been compensated for atmospheric effects.

The background term in Equation (4) contains $2C + 2D + 2$ unknown terms; C background ratios, β_j, C background emissivities, $\epsilon_j(\lambda)$, D column densities, c_i, D gas absorption spectra, $k_i(\lambda)$, the background temperature T_s, and the plume temperature T_p. In order to solve this equation for c_i by regression with a basis vector set populated with $k_i(\lambda)$, we seek to eliminate the other unknown terms. The need to solve for the $2C + 1$ background terms is eliminated by clustering the background terms into an *effective* self-emitted, surface-leaving radiance, $\bar{B}(\lambda, T_s)$. This was performed by modeling the background on a per-pixel basis. The Maximum Distance method (Max-D), an endmember extraction algorithm, was run on all non-plume pixels.[11] Fifteen endmembers were extracted using this algorithm. Since the endmembers were extracted on non-plume pixels, any spectral features due to the effluent gases are not present. After these endmembers are extracted, the algorithm addresses each plume pixel individually. These endmembers are then fit to the plume pixel using the partially-constrained stepwise regression. The goal here is not for the endmembers to model the pixel exactly. Significant residuals are not only expected but also desired in this case due to the presence of the plume. Allowing the coeffecients to take on any value eliminated the residuals too effectively, thus prompting the use of constraints on the regression. These coeffecients were then multiplied by the endmembers to arrive at the background estimate, $\bar{B}_{(x,y)}(\lambda, T_s)$, where (x, y) indicates the image coordinates of the pixel in question.

Substituting this for the background term, and distributing through the plume transmission term, we arrive at,

$$L(\lambda) = \bar{B}_{(x,y)}(\lambda, T_s) - \bar{B}_{(x,y)}(\lambda, T_s) \sum_{i=1}^{D} c_i \, k_i(\lambda) \; + \; \left[\sum_{i=1}^{D} c_i \, k_i(\lambda) \right] B(\lambda, T_p). \quad (5)$$

We can now subtract the first background term from both sides and factor out the plume emissivity term to arrive at,

$$L(\lambda) - \bar{B}_{(x,y)}(\lambda, T_s) = \sum_{i=1}^{D} c_i k_i(\lambda) \left[B(\lambda, T_p) - \bar{B}_{(x,y)}(\lambda, T_s) \right]. \quad (6)$$

The only parameters in Equation (6) that need to be addressed are the plume temperature, T_p, and the surface temperature, T_s. Specifically, we must take into account the difference between these two values. Assuming a relatively flat surface emissivity spectrum, the plume is seen in emission if this *temperature contrast* is positive. If the surface is warmer than the plume and the contrast is negative, then the plume is in absorption. The difference between emission and absorption is critical as it dictates the "direction" of the gas features (peaks or troughs) in the pixel. Since we are attempting to identify the gases by fitting a set of basis vectors to the pixel spectrum, we need to ensure that the basis vector set accounts for both emission and absorption. We do so by defining a temperature contrast term, ΔT, and replacing T_p by $(T_s \pm \Delta T)$. We estimate the surface temperature by inverting $\bar{B}_{(x,y)}(\lambda, T_s)$ to solve for the spectral brightness temperature, $T_{s,(x,y)}(\lambda)$. The maximum of $T_{s,(x,y)}(\lambda)$ is selected as the surface temperature estimate, $\widehat{T}_{s,(x,y)}$. Replacing this for T_p in Equation (6) brings us to the final equation for a plume pixel,

$$L(\lambda) - \bar{B}_{(x,y)}(\lambda, T_s) = \sum_{i=1}^{D} c_i k_i(\lambda) \left[B(\lambda, \widehat{T}_{s,(x,y)} \pm \Delta T) - \bar{B}_{(x,y)}(\lambda, T_s) \right]. \quad (7)$$

Source	dof	SS	MS
Total SS	J	$\mathbf{x}'\mathbf{x}$	
SS Due to Regression (SSR)	N	$\mathbf{f}'\mathbf{A}'\mathbf{x}$	
SS About Regression (SSE) (residuals)	$J-N$	$\mathbf{x}'\mathbf{x}-\mathbf{f}'\mathbf{A}'\mathbf{x}$	$MSE=s^2$
SSR Reduced Model	$N-1$	$\mathbf{f}'_{N-1}\mathbf{A}'_{N-1}\mathbf{x}$	
SSR Extra Term	1	$\mathbf{f}'\mathbf{A}'\mathbf{x}-\mathbf{f}'_{N-1}\mathbf{A}'_{N-1}\mathbf{x}$	MS_{extra}

Table 1. Sum of Squares ANOVA Table

2.3. Stepwise Linear Regression

Unmixing a hyperspectral pixel into its component endmembers (basis vectors) can be treated as a linear regression. The standard model for regression in matrix form is[12]

$$\mathbf{x}=\mathbf{A}\mathbf{f}+\boldsymbol{\epsilon}, \tag{8}$$

where $\mathbf{x}$ is a $(J\times 1)$ pixel vector, $\mathbf{A}$ is a $(J\times N)$ matrix of N basis vectors, $\mathbf{f}$ is a $(N\times 1)$ vector of basis vector abundances, and $\boldsymbol{\epsilon}$ is a $(J\times 1)$ error term. The abundance coefficients, $\mathbf{f}$, are solved for as

$$\mathbf{f}=\mathbf{A}^{\dagger}\,\mathbf{x}, \tag{9}$$

where $\mathbf{A}^{\dagger}$ denotes the pseudo-inverse of $\mathbf{A}$ and is defined as $\mathbf{A}^{\dagger}=(\mathbf{A}'\mathbf{A})^{-1}\,\mathbf{A}'$.

Stepwise regression begins by populating the matrix $\mathbf{A}$ with each of the M candidate basis vectors in turn. An F-test is used to determine which candidate vector best fits the data. The $M-1$ remaining vectors are then added to $\mathbf{A}$ in a similar fashion. The significance of adding a second vector to the regression model is determined using the Analysis of Variance (ANOVA) technique described below.[6,12] Once the model contains $N=2$ vectors, the $M-2$ remaining vectors are tested and, if necessary, one is added. After a third vector is added, each of the $N=3$ vectors in the model are removed in turn to ensure that each are contributing to the fit of the model. Candidate vectors are added and subtracted in this fashion until the fit of the model to the pixel spectrum can not be improved.

The ANOVA technique mentioned above calculates how much of the total variation about the mean (Sum of Squares (SS)), $\mathbf{x}'\mathbf{x}$, is due to the regression model (SSR), $\mathbf{f}'\mathbf{A}'\mathbf{x}$, and how much is from random errors (SSE), $\mathbf{x}'\mathbf{x}-\mathbf{f}'\mathbf{A}'\mathbf{x}$.[12] The ANOVA structure is depicted in Table 1. The Mean Squared (MS) column is defined by dividing the SS by the corresponding degrees of freedom. In stepwise regression, ANOVA is useful in breaking the SSR of the N-element model down into the SSR due to the previous $(N-1)$-element model and the SSR due to the addition of the new term. If we define the $(N-1)$-element model as

$$\mathbf{x}=\mathbf{A}_{N-1}\mathbf{f}_{N-1}+\boldsymbol{\epsilon}, \tag{10}$$

then we can see how this fits into the ANOVA analysis in Table 1.

The MS terms from Table 1 can be used to create a statistic to test the quality of the regression model. This statistic is based upon the distribution of the error term, $\boldsymbol{\epsilon}$. It has been shown that if the model errors are independent and distributed as $\epsilon\sim N(0,\sigma^2)$, then the sum of the squares of these errors will be Chi-squared distributed with J degrees of freedom abbreviated as χ^2_J.[12] Thus, we conclude that if the model is a good fit, then the sum-squared errors are Chi-squared distributed. If the model is a poor fit, then the errors will not follow this distribution. We formulate this as a hypothesis test, defining a null hypothesis that the abundance

coefficient for the i^{th} basis vector, $\mathbf{f}_i$, is zero. The corresponding alternate hypothesis is that the coefficient is non-zero,

$$H_0: \quad \mathbf{f}_i = 0$$
$$H_1: \quad \mathbf{f}_i \neq 0.$$

As long as the null hypothesis is true, then the ratio of the mean-squared errors,

$$\frac{MS_{extra}}{MSE} \tag{11}$$

follows an $F_{1,J-N}$ distribution.[6] For this work the F-statistic was calculated at a probability of 0.99. When attempting to add a basis vector to the model, the ratio in Equation (11) is computed. If the ratio is smaller than the F-statistic then the null hypothesis is true and the vector is rejected. If the ratio is larger than the threshold, then it is added to the model. The algorithm terminates when no vectors are either added to or subtracted from the model.

At this point we can draw similarities between the final plume-pixel radiance equation of Equation 7 and the the standard equation for matrix regression of Equation (8). The abundance fractions to be solved for, $\mathbf{f}$, are the relative gas concentration path-lengths, c_i, and the matrix of basis vectors, $\mathbf{A}$, is the absorption spectrum multiplied by the plume-background radiance contrast, $k_i(\lambda)\left[B(\lambda, \widehat{T}_{s,(x,y)} \pm \Delta T) - L_{(x,y)}(\lambda)\right]$. In addition, by defining various positive and negative values for ΔT, we can easily account for both the emission and absorption of the plume gases in our matrix of basis vectors.

By implementing the stepwise routine as described above, we are performing an unconstrained regression. While this does result in an algorithm with a very fast runtime, the downside of this is that we need a way to ensure that the coeffecients do not allow the regression to "overfit" the pixel's spectrum. Some degree of error is expected when performing such a regression, but allowing the model to take on any coefficient value necessary to minimize this error may lead to unrealistic results. In order to achieve this safeguard, rather than using the unconstrained pseudo-inverse solution of (9), we instead impose a partial constraint on the coefficients with each calculation of $\mathbf{f}$. By partially constraining the regression, we force the solution of all coefficients in $\mathbf{f}$ to be positive. Once the stepwise routine has arrived at a final set of basis vectors, we solve for the final coefficients using fully-constrained regression, in which all coefficients are both positive and sum to 1.0.

3. RESULTS

3.1. Algorithm Output

The algorithm was applied to only those pixels known to contain plume gases. A by-product of a DIRSIG image is a "truth map" that indicates the spatial location of the plume. This was used to mask out non-plume pixels. In applying the algorithm to real imagery, it is assumed that the plumes have been detected by a separate process. The raw output from the algorithm was an image cube of size $X \times Y \times M$, where X and Y are dependant upon the size of the input image. Each of the $M = 440$ bands in the output detection maps correspond to one of the candidate basis vectors fed into the regression. This value for M is the result of multiplying the 88 gas spectra from the library by five values of ΔT; -10, -5, 0, +5 and +10°C. For pixel (x, y), the value for each band was set equal to the abundance coefficient solved for at the end of the stepwise regression routine. For all bands not included in the final regression model, their value was left at 0.0.

This image cube was then refined in a secondary routine that "collapses" the output down to 31 bands, one band per gas species available in the library. The value assigned to pixel (x, y) in, for example, the freon-114 band, is the average of the values of all 15 basis vector bands in the raw output image that correspond to freon-114. These averages were then normalized by the sum of every value in all 440 bands for that particular pixel. This was done to scale all results between 0.0 and 1.0 for ease of analysis. Once the output has been reduced to 31 bands, the "spectral" dimension was changed to a "Gas Index". Each "band" now represents the returns from one particular gas species. Table 2 clarifies this by explaining which gas corresponds to which gas index value.

Gas Index	Symbol	Gas Name	Gas Index	Symbol	Gas Name
1	acrol	Acrolein	17	f125	Freon-125
2	bamine	Butylamine	18	f12	Freon-12
3	C_6H_5F	Fluorobenzene	19	f134a	Freon-134a
4	C_6H_6	Benzene	20	f218	Freon-218
5	CCl_4	Carbon Tetrachloride	21	H_2O	Water
6	CH_3Cl	Methyl chloride	22	HCHO	Formaldehyde
7	CH_3SH	Methyl mercaptan	23	HCl	Hydrogen chloride
8	CH_4	Methane	24	hyd	Hydrazine
9	CO_2	Carbon dioxide	25	NH_3	Ammonia
10	CO	Carbon monoxide	26	phg	Phosgene
11	dclm	Dichloromethane	27	SF_6	Sulfur hexafluoride
12	dclp12	1,2-Dichloropropane	28	SO_2	Sulfur dioxide
13	dclp13	1,3-Dichloropropane	29	tce1122	Tetrachloroethane
14	edb	1,–Dibromoethane	30	tce	Trichloroethylene
15	edc	1,2-Dichloroethane	31	vcl	Vinyl chloride
16	f114	Freon-114			

Table 2. Table of gas indices, the gases they represent and their abbreviations. The abbreviations used are drawn from the naming convention used by the gas absorption spectral library.

3.2. Analysis of Output

It is difficult to determine the overall effectiveness of a particular configuration of the algorithm by simple visual inspection of the consolidated, 31-band output described above. Looking at a single band of the output image gives no indication of how the other bands look as well. In addition, if one were to simultaneously display several, or even all 31 bands of the output image, it is difficult to come to a quantitative conclusion as to which gas species is most likely present in the plume by simply analyzing the output plume maps. Even the use of the "z-profile" tool in ENVI on a particular pixel is not sufficient because it can not take into account the performance of the algorithm on the entire set of plume pixels.

With this in mind, a system of analysis was developed in which regions of interest (ROI's) were defined in ENVI based on the DIRSIG truth maps. One of the truth images generated when DIRSIG simulates a scene is a map of the Plume Column Density for each gas species in the scene. This map is a greyscale image representing the column density of that particular gas in units of ppm-m. Regions of interest were made for ranges of these column densities. Values from 0.0–0.9999 constituted one region, 1.0–9.9999 a second, 10.0–99.9999 a third, 100.0–999.9999 a fourth, and 1000.0–9999.9999 a fifth. In referring to these ROI's , the values defining the ranges will be rounded for simplicity. Each column density range was not present in each test case. These ROI's were applied to the 31-band output images. Using the ROI tool in ENVI, the mean value of all pixels within each ROI was calculated on a per-band basis. This produced a series of scatterplots that depicted the average abundance coefficient as a function of gas index.

On the surface, it would seem that gas indicies with high returns in any ROI would be likely to be in the plume. This, however, is not always the case. At regions of low concentration, the background-subtracted spectrum has very weak absorption features. Any mistake in the modeling of the background can easily overshadow actual plume gas features. In addition, sensor noise becomes problematic as it too is on the order of gas absorption features. Because of this, the returns for low-concentration ROI's are more likely to identify incorrect gases and are therefore often discounted. If, for example, gas 16 has the largest return in only the highest-concentration ROI, while other gases are more prominent in other ROI's, we can still conclude that gas 16 is the correct gas.

In fact, we expect the strength of the returns to diminish along with the plume concentration. Gases whose returns increase as concentration decreases can be ruled out as a possibility.

3.3. Results for the Freon-114 Plume

Figures 2 through 4 show that the algorithm is able to successfully identify freon-114 in regions where the plume has a high concentration path-length. The Case 1 image, which has the highest release rate, has the best overall returns for gas index 16, freon-114. Only in the 1–10 ROI are there other gases that approach the return strength of freon-114. In the Case 2 results, freon-114 has the highest return in only the 100–1000 ROI. However, keeping in mind the above discussion of giving more weight to the highest-concentration ROI, and since there are sizeable returns for both the 10–100 ROI and the entire plume, we can conclude that freon-114 is the gas in the plume.

In the results for the Case 3 imagery, we can not be confident that freon-114 is present in the plume. Although it does have the highest returns for the 10–100 ROI, the returns are near zero for all other ROI's. Despite our assertion that high-concentration ROI's are more important, we can not conclude that it is freon-114, and not some other gas, that is in the plume. It is interesting to note however, that Figure 4 indicates that gas 18 is in the plume. According to Table 2, gas 18 is freon-12, a different type of freon. Since freon-12 is in the same family as freon-114, it has a similar molecular structure and therefore a similar absorption spectrum. The other gas with large returns across all ROI's in the Case 3 results is gas 28, SO_2. However we would not conclude that sulfur dioxide is likely to be in the plume because the largest return is in the 0–1 ROI. The strength of the returns are actually the opposite of what we would expect in a correct identification. In all, it is encouraging that at low gas concentrations, the algorithm was able to generically identify freon in the freon-114 plume.

3.4. Results for the NH_3 Plume

The results for the ammonia plumes in all cases are very similar to the freon-114 plumes. Figure 5 shows that gas 25, NH_3, is clearly the only gas being found by the algorithm. As with the freon-114 plume in Case 1, a second gas has equal returns in the 1–10 ROI, but this is again negligible. The Case 2 results show a clear indication of NH_3 in the 100–1000 ROI, along with the expected fall-off in strength across subsequent ROI's. Gas 18, freon-12, has significant returns across most ROI's, but we can discount it because it has equal returns across all ROI's except the largest. As for the Case 3 results, we can not come to a conclusive answer as to which gas is in the plume based on the results in Figure 7. Unlike the Case 2 results, in which the returns for NH_3 were much larger than for freon-12 in the 100–1000 ROI, the two gases have near equal returns across all ROI's. Ammonia has the higher return in the 10–100 ROI, but freon-12 is slightly larger in all other regions. In addition, the strength of the returns fall off for freon-12 across the ROI's, indicating that it is likely to be present in the plume. Because of this, we can not come to a definitive answer for Case 3.

It is of interest to note that at small concentration path-lengths, both plumes began to be falsley identified as the same gas, freon-12. The reason for this is not quite known. The fact that it is being falsley associated with the freon-114 plume is not surprising since both gases are very similar chemically. However it is not known why it is so prominent in the NH_3 plume. It is understood why the algorithm begins to fail at low concentration path-lengths. In regions where the gas signature is weak, our characterization of the background becomes more important. Any mismatch in the background model is now of equal or larger magnitude than the actual gas features themselves. This has less of an impact where the gas is at a high concentration path-length and therefore has prominent spectral features. This reinforces the idea that the primary obstacle to gas identification is finding a way to know exactly what the background spectrum is behind the plume gas.

4. CONCLUSIONS

Stepwise linear regression was used to identify the constituent gases in gas plumes. The regression was partially constrained within each iteration of the stepwise routine, while the final coefficients were solved for using fully-constrained regression. Three cases of synthetic imagery were generated for the study. Each case had two plumes, one containing freon-114, the other NH_3. The release rates of both plumes were the same in each case, and that rate decreased across the three cases. Absorption spectra of 31 gases were multiplied by five temperature contrast values to generate the library of potential basis vectors. A series of output detection maps were generated for

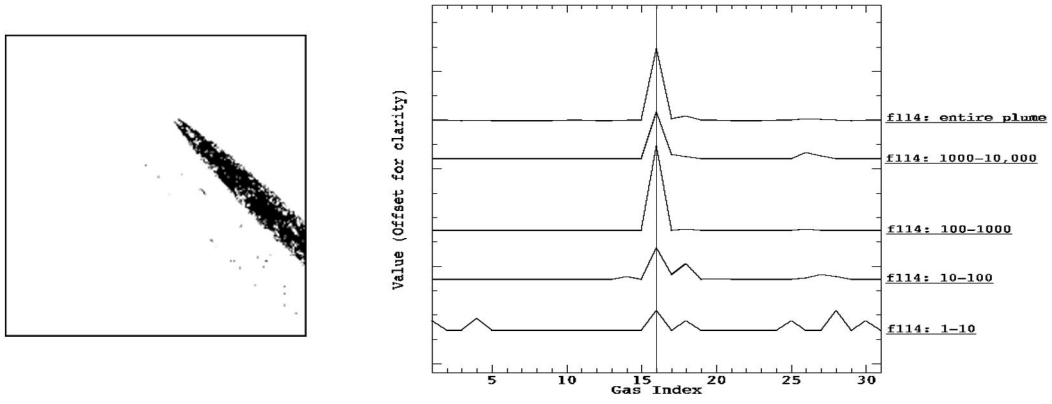

Figure 2. Results for the freon-114 plume of Case 1. The detection map for Gas Index 16 is on the left while the corresponding ROI plots are on the right. Gas Index 16 is marked in the plots to highlight the returns for freon-114.

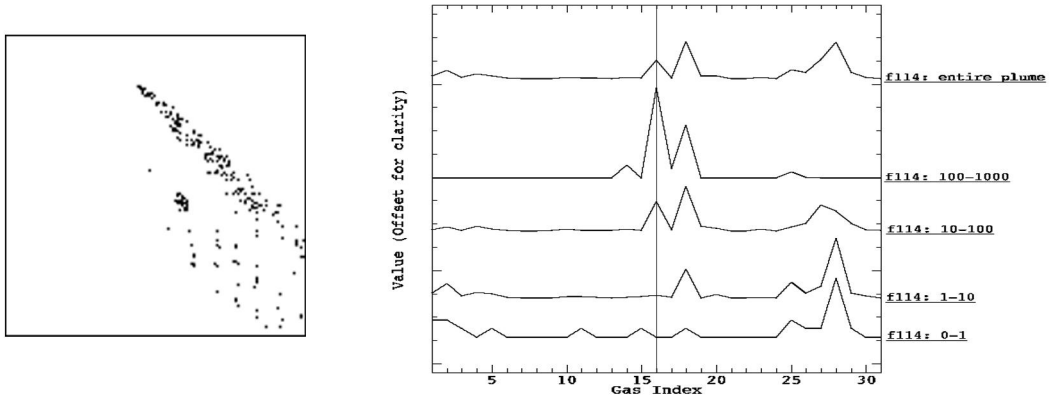

Figure 3. Results for the freon-114 plume of Case 2. The detection map for gas index 16 is on the left while the corresponding ROI plots are on the right. Gas Index 16 is marked in the plots to highlight the returns for freon-114.

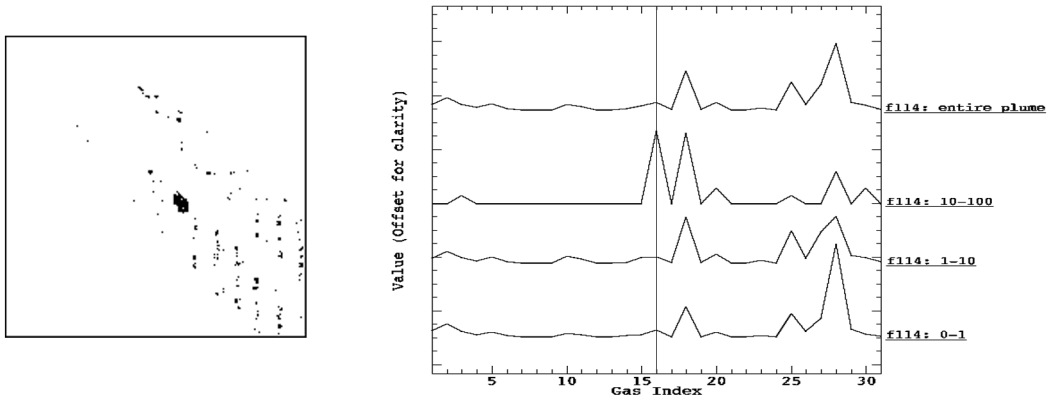

Figure 4. Results for the freon-114 plume of Case 3. The detection map for gas index 16 is on the left while the corresponding ROI plots are on the right. Gas Index 16 is marked in the plots to highlight the returns for freon-114.

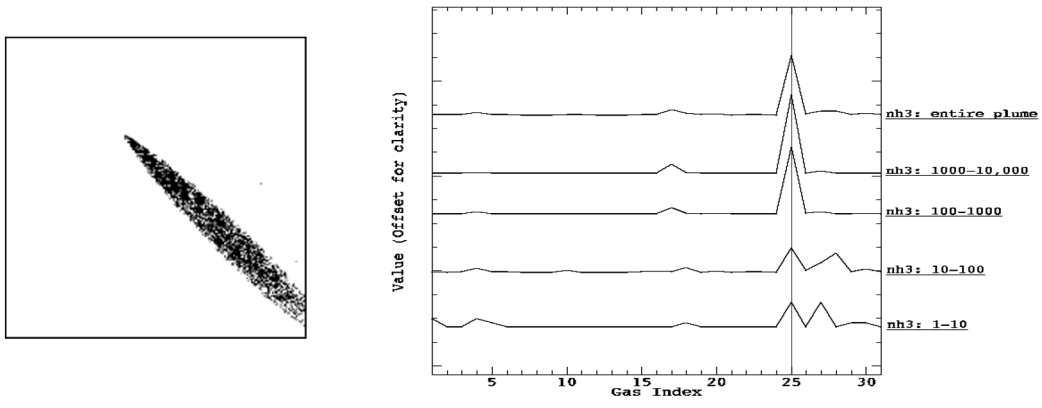

Figure 5. Results for the NH_3 plume of Case 1. The detection map for gas index 25 is on the left while the corresponding ROI plots are on the right. Gas Index 25 is marked in the plots to highlight the returns for NH_3.

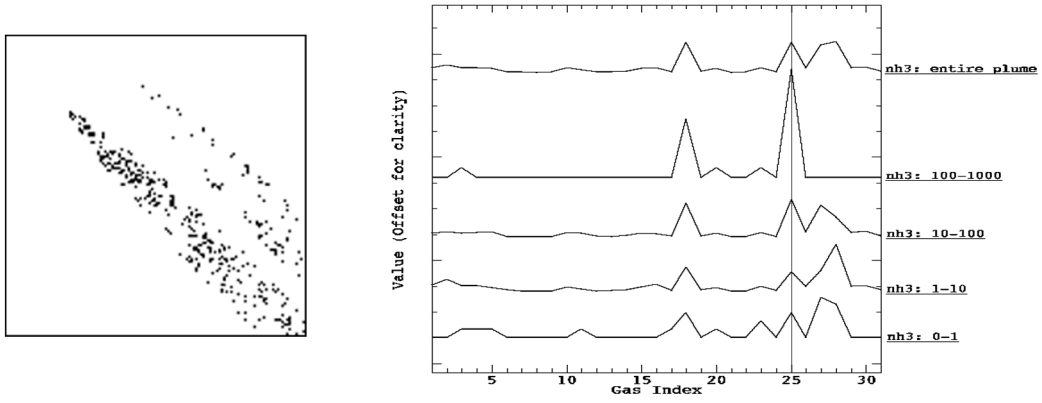

Figure 6. Results for the NH_3 plume of Case 2. The detection map for gas index 25 is on the left while the corresponding ROI plots are on the right. Gas Index 25 is marked in the plots to highlight the returns for NH_3.

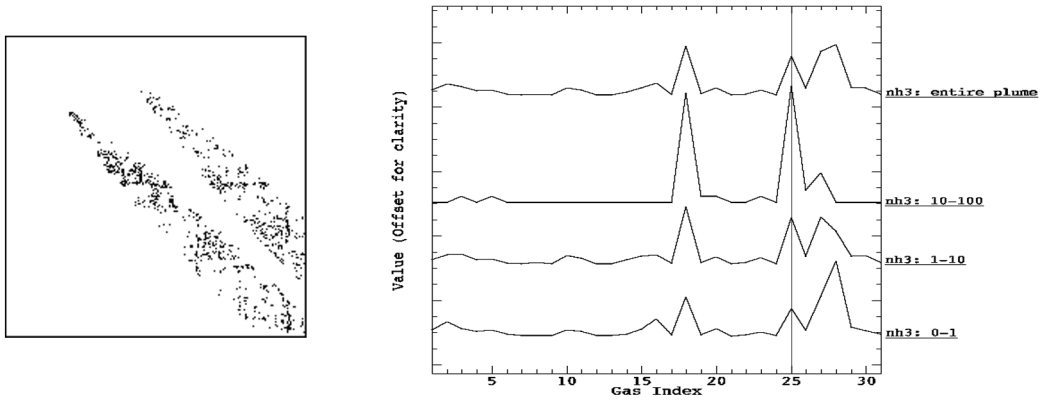

Figure 7. Results for the NH_3 plume of Case 3. The detection map for gas index 25 is on the left while the corresponding ROI plots are on the right. Gas Index 25 is marked in the plots to highlight the returns for NH_3.

each basis vector and subsequently collapsed to a per-species level. Regions of interest were defined based on DIRSIG truth outputs and used to analyze the algorithm output.

In two of the three cases the algorithm was able to correctly identify both freon-114 and NH_3 in their respective plumes. The ability to correctly identify gas species in single-gas plumes is directly related to the concentration of the gas. Highlights of the algorithm are that the background is completely image-derived, regions in which the gas is in both emission and absorption are built into the regression, no prior knowledge of gas species or plume temperature are needed, and the algorithm is completely automated.

When applying this algorithm to real imagery, the lack of truth maps prevents the definition of ROI's based on exact plume column densities. However since many of the correct identifications were made in only the ROI's corresponding to the largest concentrations, we would like to be able to retain our ability to monitor the relative magnitude of the returns between ROI's. One manner in which this could be achieved is to manually draw ROI's based on the plume detection maps. A user would be able to draw, at a minimum, three regions of "low", "medium", and "high" concentrations. The highest-concentration ROI would be drawn nearest the stack and emanate downwind, with successive ROI's growing larger to encompass all plume pixels. This would then allow the user to analyze the algorithm output in the same manner as shown in this work.

ACKNOWLEDGMENTS

This work was funded under the Office of Naval Research Multi-disciplinary University Research Initiative "Model-based Hyperspectral Exploitation Algorithm Development" #N00014-01-1-0867. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Office of Naval Research. The authors would like to thank Herb Fry, Kevin Mitchell and Brian McVey for their helpful correspondence during the course of this research.

REFERENCES

1. W. Marinelli, C. Gittins, A. Gelb, and B. Green, "A tunable fabry-perot etalon-based long-wavelength infrared imaging specroradiometer," *Applied Optics* **38**(12), pp. 2594–2604, 2000.
2. J. Lisowski and C. Cook, "A svd method for spectral decomposition and classification of ares data," in *SPIE Optical Science, Engineering and Instrumentation Annual Meeting*, *Proceedings of SPIE* **2821**, pp. 14–29, 1996.
3. R. Harig, G. Matz, and P. Rusch, "Scanning infrared remote sensing system for identification, visualization and quantification of airborne pollutants," in *Instrumentation for Air Pollution and Global Atmospheric Monitoring*, J. Jensen and R. Spellicy, eds., *Proceedings of SPIE* **4574**, pp. 83–94, 2002.
4. S. Young. Aerospace Corporation, RAND Communication, 2002.
5. R. Clark, A. Gallagher, and G. Swayze, "Material absorption band depth mapping of imaging spectrometer data using a complete band shape least-squares fit with library reference spectra," *Proceedings of the Second Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) Workshop* , pp. 176–186, 1990a.
6. H. Gross and J. Schott, "Application of spatial resolution enhacement and spectral mixture analysis to hyperspectral images," *SPIE Optical Science, Engineering, and Instrumentation Annual Meeting* **2821**, pp. 30–41, 1996.
7. J. Schott, S. Brown, R. Raqueno, H. Gross, and G. Robinson, "An advanced synthetic image generation model and its application to multi/hyperspectral algorithm development," *Canadian Journal of Remote Sensing* **25**(2), pp. 99–111, 1999.
8. S. Kuo, J. Schott, and C. Chang, "Synthetic image generation of chemical plumes for hyperspectral applications," *Optical Engineering* **39**(4), pp. 1047–1056, 2000.
9. J. Hackwell, D. Warren, R. Bongiovi, H. S.J., T. Hayhurst, D. Mabry, M. Sivjee, and J. Skinner, "Lwir/mwir imaging hyperspectral sensor for airborne and ground-based remote sensing," in *Imaging Spectrometry II*, M. Descour and J. Mooney, eds., *Proceedings of SPIE* **2819**, pp. 102–107, 1996.
10. D. Gu, A. Gillespie, A. Kahle, and F. Palluconi, "Autonomous atmospheric compensation (aac) of high resolution hyperspectral thermal infrared remote-sensing imagery," *IEEE Transactions on Geoscience and Remote Sensing* **38**(6), pp. 2557–2570, 2000.

11. P. Bajorski, E. Ientilucci, and J. Schott, "Comparison of basis-vector selection methods for target and background subspaces as applied to subpixel target detection," in *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery X*, S. Shen and P. Lewis, eds., *Proceedings of SPIE* **5425**, pp. 97–108, 2004.
12. N. Draper and H. Smith, *Applied Regression Analysis, 2^{nd} ed.*, John Wiley & Sons, New York, 1981.