

Gas Plume Species Identification in LWIR Hyperspectral Imagery
by Regression Analyses

by

David Pogorzala

B.S. Rochester Institute of Technology, 2002

A thesis submitted in partial fulfillment of the
requirements for the degree of Master of Science
in the Chester F. Carlson Center for Imaging Science
Rochester Institute of Technology

April 13, 2005

Signature of the Author _____

Accepted by _____
Coordinator, M.S. Degree Program Date

CHESTER F. CARLSON CENTER FOR IMAGING SCIENCE
ROCHESTER INSTITUTE OF TECHNOLOGY
ROCHESTER, NEW YORK

CERTIFICATE OF APPROVAL

M.S. DEGREE THESIS

The M.S. Degree Thesis of David Pogorzala
has been examined and approved by the
thesis committee as satisfactory for the
thesis required for the
M.S. degree in Imaging Science

Carl Salvaggio, Ph.D., thesis Advisor

John Schott, Ph.D.

David Messinger, Ph.D.

Date

THESIS RELEASE PERMISSION
ROCHESTER INSTITUTE OF TECHNOLOGY
CHESTER F. CARLSON CENTER FOR IMAGING SCIENCE

Title of Thesis:

**Gas Plume Species Identification in LWIR Hyperspectral Imagery by
Regression Analyses**

I, David Pogorzala, hereby grant permission to Wallace Memorial Library of R.I.T. to reproduce my thesis in whole or in part. Any reproduction will not be for commercial use or profit.

Signature _____ Date _____

Gas Plume Species Identification in LWIR Hyperspectral Imagery by Regression Analyses

by
David Pogorzala

Submitted to the
Chester F. Carlson Center for Imaging Science
in partial fulfillment of the requirements
for the Master of Science Degree
at the Rochester Institute of Technology

Abstract

The goal of this research was to develop an algorithm for identifying the constituent gases in stack releases. At the heart of the algorithm is a stepwise linear regression technique that only includes a basis vector in the model if it contributes significantly to the fit. This significance is calculated by an F -statistic. Issues such as atmospheric compensation, gas absorption and emission, background modeling, and fitting a linear regression to a non-linear radiance model were addressed in order to generate the matrix of basis vectors. Synthetic imagery generated by the DIRISG model were used as test cases. Results show that the ability to correctly identify a gas diminishes as a function of decreasing concentration path-length of the plume. Results drawn from pixels near the stack are more likely to give an accurate identification of the gas present in the plume.

Acknowledgements

Thanks go out to...

Dave, for putting up with Erin and I even when we acted like 5 year-olds.

Carl, my Thesis Advisor, for all the advising.

Schott, for being the Man.

Erin, for making plumeification fun.

My folks, for thinking that I'm smart because I made this.

Leah, for just being great.

Contents

Certificate of Approval	i
Copyright Release	ii
Abstract	iii
Acknowledgements	iv
List Of Figures	xii
1 Introduction	1
1.1 RIT Gas Research	1
1.2 Motivation for the Gas Problem	2
1.2.1 Environmental Monitoring and Regulation	2
1.2.2 National Security Issues	2
1.3 Hyperspectral Imagery	3

1.4	Gas Spectroscopy	5
2	Previous Work	8
2.1	Stepwise Regression in the Reflective Region	8
2.2	Constituent Gas Identification	10
2.3	Differences from Previous Work	16
3	Methodology	21
3.1	Sources of Data	21
3.1.1	Real Imagery	21
3.1.2	Synthetic Imagery	23
3.1.3	Gas Spectra	27
3.2	Radiance Model	27
3.3	Regression Techniques	32
3.3.1	Matrix Regression	32
3.3.2	Stepwise Regression	34
3.3.3	Constraints on the Regression	37
3.4	Definition of Basis Vectors	40
3.4.1	Variables	40
3.4.2	Atmospheric Compensation	41
3.4.3	Background Characterization	43
3.4.4	Emission vs. Absorption	44

3.5	Formulation of the Identification Algorithm	46
4	Results	49
4.1	Test Cases	49
4.2	Algorithm Output	51
4.3	Analysis of Output	55
5	Analysis	59
5.1	Separate Plume Results	61
5.2	Mixed Gas Results	64
5.3	Blind Test Results	65
5.4	Effects of Atmospheric Compensation	66
5.5	Effects of the Change in the F -Statistic	67
6	Conclusions	69
6.1	Summary	69
6.1.1	Methods for Imposing Regression Constraints	70
6.1.2	Extending the Algorithm to Real Imagery	71
6.2	Recommendations for Future Work	73
6.2.1	Effect of F -Statistic Probability	73
6.2.2	Constraints On Concentration Path-Length	74
6.2.3	Atmospheric Compensation	74

6.2.4	Spectral Overlap	75
6.2.5	Effects of a Scene-Derived Background Estimate	76
6.2.6	Effect of Sensor Noise	77
6.2.7	Real Imagery	78
Bibliography		79
A Detection Maps From All Cases		83
B Graphical Results From All Cases		95

List of Figures

1.1	Absorption Spectrum of NH_3	7
2.1	Traditional vs. Stepwise (Per-Pixel) Unmixing [Gross and Schott, 1996]	9
2.2	Effect of Temperature Variability on SF_6 Absorption Spectrum	19
3.1	DIRSIG Plume Model Verification	24
3.2	Single-band image of Case 2 synthetic DIRSIG data at $10.73\mu\text{m}$	26
3.3	Comparison of self-emitted radiance with reflected downwelled radiance.	29
3.4	Workflow of the identification algorithm.	48
4.1	Example detection map	54
4.2	Snapshot of the ROI's used for Case 4 results.	57
4.3	Example scatterplot of the average return for each ROI	58
5.1	Absorption spectra for Freon-114 and freon-12	62
5.2	Absorption spectra for Freon-114 and CH_4	64

A.1	Detection maps for Case 2	84
A.2	Detection maps for Case 2 with atmospheric compensation	85
A.3	Detection maps for Case 2 with a 95% F -test	86
A.4	Detection maps for Case 3	87
A.5	Detection maps for Case 4	88
A.6	Detection maps for Case 4b	89
A.7	Detection maps for Case 4b with atmospheric compensation	90
A.8	Detection maps for Case 5	91
A.9	Detection maps for Case 5 with atmospheric compensation	92
A.10	Detection maps for Case 6	93
A.11	Detection maps for Case 7	94
B.1	Method A results for the NH_3 plume in Case 2.	96
B.2	Method A results for the Freon-114 plume in Case 2.	96
B.3	Method B results for the NH_3 plume in Case 2.	97
B.4	Method B results for the Freon-114 plume in Case 2.	97
B.5	Method A results for the NH_3 plume in Case 2. For this set of results, the F -test was calculated at a probability of 95%.	98
B.6	Method A results for the Freon-114 plume in Case 2. For this set of results, the F -test was calculated at a probability of 95%.	98

B.7 Method A results for the NH_3 plume in the atmospherically compensated Case 2 image.	99
B.8 Method A results for the freon-114 plume in the atmospherically compensated Case 2 image.	99
B.9 Method B results for the NH_3 plume in the atmospherically compensated Case 2 image.	100
B.10 Method B results for the freon-114 plume in the atmospherically compensated Case 2 image.	100
B.11 Method A results for the single plume in Case 3.	101
B.12 Method B results for the single plume in Case 3.	101
B.13 Method A results for the NH_3 plume in Case 4.	102
B.14 Method A results for the Freon-114 plume in Case 4.	102
B.15 Method B results for the NH_3 plume in Case 4.	103
B.16 Method B results for the Freon-114 plume in Case 4.	103
B.17 Method A results for the NH_3 plume in Case 4b.	104
B.18 Method A results for the Freon-114 plume in Case 4b.	104
B.19 Method B results for the NH_3 plume in Case 4b.	105
B.20 Method B results for the Freon-114 plume in Case 4b.	105
B.21 Method A results for the NH_3 plume in the atmospherically compensated Case 4b image.	106

B.22 Method A results for the Freon-114 plume in the atmospherically compensated Case 4b image.	106
B.23 Method B results for the NH_3 plume in the atmospherically compensated Case 4b image.	107
B.24 Method B results for the Freon-114 plume in the atmospherically compensated Case 4b image.	107
B.25 Method A results for the single plume in Case 5.	108
B.26 Method B results for the single plume in Case 5.	108
B.27 Method A results for the single plume in the atmospherically compensated Case 5 image.	109
B.28 Method B results for the single plume in the atmospherically compensated Case 5 image.	109
B.29 Method A results for the CH_4 plume in Case 6.	110
B.30 Method A results for the dclp12 plume in Case 6.	110
B.31 Method B results for the CH_4 plume in Case 6.	111
B.32 Method B results for the dclp12 plume in Case 6.	111
B.33 Method A results for the Unknown Gas 1 plume in Case 7.	112
B.34 Method A results for the Unknown Gas 2 plume in Case 7.	112
B.35 Method B results for the Unknown Gas 1 plume in Case 7.	113
B.36 Method B results for the Unknown Gas 2 plume in Case 7.	113

Chapter 1

Introduction

1.1 RIT Gas Research

The proposed thesis fits into a larger body of algorithm development research aimed at learning more about gas plumes in hyperspectral imagery. This research is threefold. Given a hyperspectral image: *detect* any and all gas plumes in the scene, *identify* the constituent gases present in the plumes, and *quantify* in both relative and absolute terms the amount of each gas. The ideal end product of this body of research would be a fully-automated algorithm that performs all of the above tasks. As stated above, this research will attempt to solve the *identification* problem by identifying each gas present in a *pre-detected* plume.

1.2 Motivation for the Gas Problem

Quantifying the amount of individual gases being released from factory stacks is of great importance to two main special interest groups; environmental and defense agencies.

1.2.1 Environmental Monitoring and Regulation

Effective monitoring of air pollution, in particular from large production facilities with numerous discharge locations, presents quite a challenge. These discharges are typically released from tall stacks where the gases quickly dissipate into the atmosphere. These conditions make accurate and reliable independent measurement methods difficult. Airborne remote sensing is a convenient platform that allows a user to survey an entire group of facilities at once. The development of algorithms that accurately identify and quantify the individual gases being released will greatly enhance the abilities of regulatory agencies to enforce legislation.

1.2.2 National Security Issues

The enforcement of environmental legislation can be accomplished by simply monitoring the release rates of specific chemicals and gases that have been flagged as hazardous. However, given to the right expert, this information can also be used to draw conclusions about the methods and materials involved in the manufacturing process within the plant. Defense agencies are interested in the use of hyperspectral imagery to see what

all the “bad guys” are making.

A second use for the gas problem is for tactical purposes. Given a real-time imager and a fast-acting algorithm, a field commander would be able to detect the release of a gaseous chemical weapon and mobilize their troops accordingly. In this, as well as other applications, a solution to the gas problem would be a great asset.

1.3 Hyperspectral Imagery

Airborne-imaging spectrometers give rise to three-dimensional data cubes. Each pixel can be considered as a spectrum sampled at J points, with J being the number of bands in the specific instrument. When J numbers in the tens to hundreds, the image is referred to as *hyperspectral*. In the 0.4–0.7 μm , the visible (VIS), 0.7–1.1 μm , the near infrared (NIR) and 1.1–3 μm , the shortwave infrared (SWIR) bandpass regions, reflected solar irradiance is the dominant source of radiance seen by the sensor. Collectively these three regions are referred to as the *reflective* region. Beyond 5 μm , the longwave infrared (LWIR) region, solar irradiance is dominated by self-emitted radiance, hence being referred to as the *emissive* region. Between 3–5 μm , the midwave infrared (MWIR) region, the reflective and emissive sources are present in approximately equal magnitude [Schott, 1997].

The altitude of the aircraft is the largest factor in determining the spatial resolving

power of an airborne imager. The ground sample distance (GSD), or the size of the pixel on the ground, is proportional to the altitude. Typical GSD's for aircraft-mounted imagers are 1 to 5 m. At this resolution, individual materials do not always fill an entire pixel. Namely, a single pixel may span the boundary between an asphalt parking lot and a grassy field. When a target fills an entire pixel, it is said to be *fully resolved*. Conversely, a *mixed pixel* is one that contains several material types. There are three general types of mixed pixels. An areal mixture is when there is very little spatial interaction between the materials within the pixel, such as the parking lot - grassy field boundary. An aggregate mixture arises when the material types have many boundaries between them. Given a pixel with a 50 m GSD, the random placement of automobiles in a parking lot is an example of an aggregate mixture. The third class of mixtures is intimate, where the two materials can not be spatially separated. An example of this is an optically thin plume of gas passing in front of a background material. In all mixture types, the pixel spectrum is a weighted combination of the spectrum from each material. *Unmixing* is the process of identifying the materials within a mixed pixel and determining the relative abundances of each. Interpreting the spectrum of a mixed pixel as a J -dimensional vector is the basis of most unmixing and target detection algorithms. In both areal and aggregate mixtures, the unmixing process is linear. Unmixing intimate mixtures is a nonlinear process [Gross, 1996].

1.4 Gas Spectroscopy

Since hyperspectral imagery captures a visible/infrared spectrum of a target, it is important to understand where these spectra come from and what determines their shape. A molecule at equilibrium contains several types of energy. Two examples include vibrational energy due to the stretching and shrinking of inter-atomic bonds, and the electronic energy of electrons orbiting the nuclei. These energies are present in discrete, quantized amounts. If the molecule were to emit this energy or absorb more energy, it too would be quantized. Consider two adjacent energy levels to be E_1 and E_2 , where $\Delta E = E_2 - E_1$. Emission and absorption of energy by a molecule can be translated into a spectral dimension by

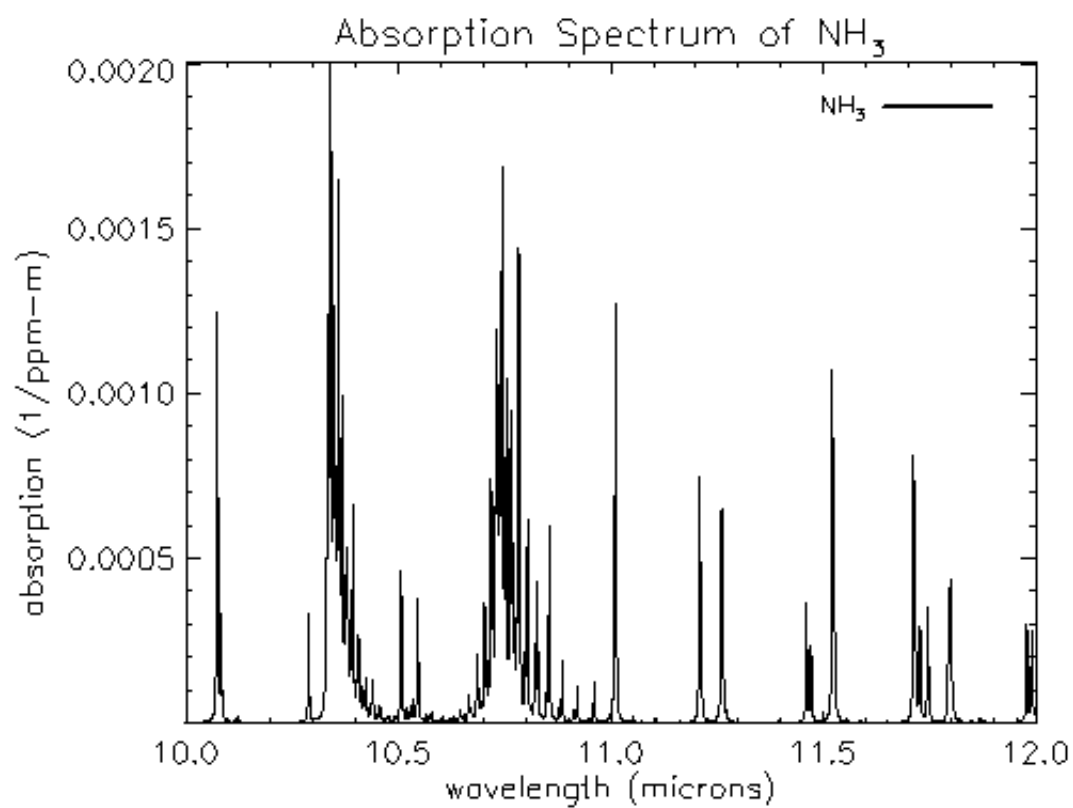
$$\Delta E = \frac{hc}{\lambda}, \quad (1.1)$$

where $h = 6.626 \times 10^{-34} \text{ J} \cdot \text{s}$ is Planck's constant, $c = 3.0 \times 10^8 \text{ m/s}$ is the speed of light, and λ is the wavelength associated with ΔE . Wavelength is typically expressed in μm in the LWIR, but here it is given in meters to balance the equation. A positive value for ΔE indicates that energy is being emitted, while a negative value indicates absorption.

The different molecular energy types described above each correspond to a different region in the electromagnetic (EM) spectrum. High energies in the X-ray region correspond to electrons in the inner shells of an atom changing energy levels. Microwave

energies correspond to changes in rotational energies of the entire molecule. Features in the infrared region, the region in which current hyperspectral imagery is being collected for the gas problem, are the result of molecular vibration [Banwell, 1972]. In these transitions the electric dipole of the molecule undergoes a change in direction. This occurs for two reasons. First, the bond length between two atoms is not fixed, but rather oscillates around a mean distance. As the bond length increases or decreases, the dipole of the entire molecule can change as well. Second, the angle between multiple bonds oscillates. These vibrations change the overall shape of the molecule, again changing the direction of the dipole.

An example gas spectrum for NH_3 is shown in Figure 1.1. This molecule exhibits several significant absorption features, most notably the two at 10.4 and 10.75 μm . Using equation (1.1), we see that these features correspond to energy transitions of 1.911×10^{-20} J and 1.849×10^{-20} J, respectively.

Figure 1.1: Absorption Spectrum of NH_3

Chapter 2

Previous Work

2.1 Stepwise Regression in the Reflective Region

The reflective region of hyperspectral imagery is heavily exploited for the purposes of material identification and target detection. Most algorithms designed for either purpose work by unmixing each pixel in the image. Standard unmixing is performed using the matrix-regression techniques discussed in Sec. 3.3.1. A stepwise approach is employed by Gross and Schott (1996). The theory and implementation of their algorithm is described in detail in Sec. 3.3.2. and is therefore not discussed here. However, they do investigate the event where the model produces unrealistic results for any given pixel. Such results are defined as large positive or negative mixing ratios. When these are detected, the most abundant material found by the model is removed from the library, and the stepwise algorithm is re-run on the pixel. The output from this second




Figure 2.1: Traditional vs. Stepwise (Per-Pixel) Unmixing [Gross and Schott, 1996]

regression is found to be significantly different from the first, and often produces sensible results. While this self-check could be implemented in gas identification, it is more appropriate for gas quantification. Material quantification was the goal of the study in Gross and Schott (1996). An unusually large abundance fraction was of more concern to them. If their only goal was to identify the materials, as is the case here, the actual value of the mixing ratios would have been of little importance.

Gross and Schott tested the stepwise approach against two standard-unmixing processes. The results are recreated in Figure 2.1. The test image consisted of 10 end-members, of which only 4 (grass, trees, dirt and water) were present in the majority of the pixels. Both the first and second unmixing algorithms used a fixed number of

endmembers in each pixel. The first used all 10 endmembers while the second incorporated only the 4 predominant materials. Both methods cause the model to “overfit” the endmembers to the pixel, resulting in a high error term. The last method is the stepwise approach that uses a different endmember set for each pixel. As seen in Figure 2.1, this approach produces the smallest error term.

2.2 Constituent Gas Identification

Much work has been done in the field of identifying gas releases via remotely sensed data, however, most of these studies are either a sensor test-case or are performed using a side-looking geometry. Marinelli, *et al.* (2000) test their Fabry-Perot spectroradiometer by identifying dimethyl methylphosphonate (DMMP) and SF_6 both inside a controlled chamber and from a small release in front of a warm building. Their detection scheme (Spectral Matched Filter) requires the user to define the wavelengths of interest based on features of the gas in question. In a sense they are detecting a pre-identified gas. Lisowski and Cook (1996) use singular value decomposition (SVD) and band differencing to detect SO_2 in overhead imagery. Again, a user is required to select a spectral subset of the image corresponding to the SO_2 features. This is required for the SVD method since the spectral subset is reduced to a set of eigenvectors upon which the entire image is projected, and is required for the differencing technique since the background is normalized to the gas absorption peak before subtraction.

Harig, *et al.* (2002) tested a Fourier-transform infrared (FTIR) spectrometer by analyzing a gas cloud in a horizontal geometry. Although their geometry, and therefore radiance model, does not apply to this problem, their identification algorithm is worth investigation. They begin by converting the pixel radiance spectrum to a brightness temperature spectrum using

$$T_{br}(\lambda) = \frac{h c}{\lambda k \ln \left[\frac{2 \pi h c^2}{\lambda^5 L(\lambda)} + 1 \right]} \quad (2.1)$$

where $k = 1.38 \times 10^{-23}$ J/K is the Boltzmann gas constant. Equation 2.1 is derived by inverting the standard Planck equation for a blackbody radiator. They scale a single library absorption spectrum as well as atmospheric gas spectra to this brightness temperature spectrum. The atmospheric gas spectra are subtracted from the pixel spectrum to create a difference spectrum. The correlation coefficient is computed between this difference spectrum and the library absorption spectrum in specified spectral windows. If this coefficient is above a threshold, then that particular species is said to be present in the gas cloud. The threshold is different for each gas that is tested.

This notion of using a specific set of spectral windows for each material (*i.e.*, gas species) was established by Clark, *et al.* (1990). Clark demonstrated an algorithm that produced material maps of minerals by locating spectral features unique to each mineral and measuring the band depth at these wavelengths. This is an attractive possibility for the gas problem since absorption features of gases are more pronounced than those of minerals. The work by Harig, *et al.* reinforces this idea.

Perhaps the most complete investigation into the gas problem was done by Young (2002). He performed gas detection using a spectral matched filter and gas identification using both constrained and unconstrained linear regression. His data set is described in Section 3.1.1.

The radiance model used by Young to describe the gas problem is similar to that derived here in Section 3.2. Neglecting atmospheric effects and utilizing Kirchoff's Law as described in Section 3.2, both models begin with the same preliminary expression for the radiance from a plume pixel as,

$$L(\lambda) = \epsilon_s(\lambda) B_s(\lambda, T_s) \left[1 - \sum_{i=1}^D c_i k_i(\lambda) \right] + \sum_{i=1}^D c_i k_i(\lambda) B(\lambda, T_p), \quad (2.2)$$

where

$\epsilon_s(\lambda)$ = surface emissivity spectrum,

$B_s(\lambda, T_s)$ = Planck function evaluated at the surface temperature, T_s ,

c_i = column density of gas i [ppm-m],

$k_i(\lambda)$ = absorption spectrum of gas i [1/ppm-m], and

$B_p(\lambda, T_p)$ = Planck function evaluated at the plume temperature, T_p .

At this point the two models begin to differ. The assumptions used in this study are discussed in Section 3.4 while Young's are discussed here.

Young rearranges (2.2) and defines a *thermal contrast spectrum* for the i^{th} gas as,

$$C_i(\lambda) = c_i [B_p(\lambda, T_p) - \epsilon_s(\lambda) B_s(\lambda, T_s)], \quad (2.3)$$

He then makes the assumption that $C_i(\lambda)$ varies slowly compared to $k_i(\lambda)$, and consequently can be approximated as a constant,

$$C_i(\lambda) \approx C_i = c_i [B_p(\bar{\lambda}_i, T_p) - \epsilon_s(\bar{\lambda}_i) B_s(\bar{\lambda}_i, T_s)], \quad (2.4)$$

where $\bar{\lambda}_i$ is calculated for each gas species. Incorporating the thermal contrast factor, and assuming atmospheric compensation, Young writes his radiance model as,

$$L(\lambda) = \sum_{i=1}^D C_i k_i(\lambda) + \epsilon_s(\lambda) B_s(\lambda, T_s). \quad (2.5)$$

The first term in (2.5) represents the plume radiance due to the temperature difference between the plume and the background. He defines the second term in (2.5) as the background term, $L_0(\lambda)$. In order to characterize the background, he expands $L_0(\lambda)$ onto a set of basis vectors, $B_j(\lambda)$,

$$L_0(\lambda) = \sum_{j=1}^C \beta_j B_j(\lambda), \quad (2.6)$$

where β_j are the coefficients for each basis vector. Substituting this into (2.5) he arrives at

$$L(\lambda) = \sum_{i=1}^D C_i k_i(\lambda) + \sum_{j=1}^C \beta_j B_j(\lambda), \quad (2.7)$$

as his final model.

Young attempts three methods for defining the basis vectors, $B_j(\lambda)$: direct extraction of image pixels, clustering of the background pixels, and using the eigenvectors of the scene covariance matrix. He determined that, for simple gas identification purposes, each of the methods produced comparable results. It is only when he attempted to quantify the gases that he discovered that directly-extracted pixels produce the best results.

The regression model employed by Young is similar to that described in Section 3.3.1. Rather than solving for the mixing ratios using (3.8), he incorporates a lower limit that will prohibit the regression error, ϵ , from being smaller than the pixel noise covariance, \mathbf{C}_x . This turns the solution to the regression equation into,

$$\mathbf{f} = (\mathbf{A}' \mathbf{C}_x \mathbf{A})^{-1} \mathbf{A}' \mathbf{C}_x \mathbf{x}. \quad (2.8)$$

This formulation is used to solve for C_i and β_j in (2.7). In the unconstrained case, both the C_i and β_j coefficients are allowed to take on any sign. In the constrained case, the β_j coefficients are required to be positive and all of the C_i are required to be the same sign, meaning that all gases are simultaneously in absorption or emission within the same pixel.

The latter constraint has been shown by Foy, *et al.* (2002) to be invalid in some cases where the temperature contrast between the plume and the background is near zero. They describe the scenario in which Freon 113 is seen in both emission and absorption

against a “common paint” within the same pixel. The relative intensities of the Freon 113 absorption spectrum and the paint emissivity spectrum are what dictate how the gas is seen. If this case were to occur, any attempts at regression on the pixel would fail. The likelihood of detecting this phenomenon in a typical hyperspectral scene is low. Consider that most plumes in hyperspectral imagery span hundreds of pixels. The zero-contrast region is seen in, at most, a few dozen pixels. Of these, very few, if any pixels, will contain the correct combination of gases and background materials such that their spectral features overlap to produce this effect. With this in mind, these extreme cases can be ignored and we can assume, as does Young, that all gases will be seen in either emission or absorption within a pixel.

In order to test the unconstrained and constrained methods, Young establishes a test case against an off-plume subtraction. By hand-selecting a non-plume pixel, he is confident in the accuracy of the solved-for coefficients. The candidate gas signatures included in the model were N_2O , H_2O , CO_2 , NO_2 , and NH_3 . These gases were known to be the effluents from the refining plant in the scene. In both regression cases, the background basis set was selected by the background clustering method. The unconstrained case yielded unrealistic results. Two of the gas coefficients were negative, and none of the five were similar to those arrived at by the off-plume subtraction method. The output from the constrained case was much more consistent with the subtraction results. As per the regression constraints, all coefficients were the same

sign (positive), and their magnitudes were similar to those of the off-plume method.

As mentioned above, Young used a spectral matched filter to detect the plume in the scene. The formulation was similar to the regression solution of (2.8), with three differences. First, rather than using the covariance of only the pixel of interest, the matrix was computed for the entire scene. Second, the basis vector set, \mathbf{A} , consisted of only the single absorption spectrum of the gas being detected. Last, the pixel spectrum, \mathbf{x} , was mean-subtracted, $\mathbf{x} - \bar{\mathbf{x}}$. As mentioned above, only a few gases were tested using the matched filter as Young had prior knowledge of which species were likely to be in the plume.

2.3 Differences from Previous Work

The works noted in the previous section provide a working knowledge of the gas problem, approaches to solving the problem, and techniques that are used in this thesis. They do not, however, form a complete outline of the path of this research. It is important to illustrate the differences between the work done above and the manner in which it was applied to this research.

The stepwise-regression technique employed by Gross and Schott (1996) was performed on VIS/NIR image data, where the library of spectra were reflectance rather than absorption spectra. The regression is purely a statistical technique. Its success does not

depend upon assumptions to the radiance model that are only valid in the VIS/NIR. With this in mind, it is expected that stepwise regression will perform well using absorption spectra. One goal of this research was to determine if the stepwise technique can be extended to LWIR image data.

Several comparisons can be made between the methods that are employed in this research and those done by Young. The first issue is the thermal contrast spectrum defined in (2.3). As shown above, Young assumes that the spectral dimension of this term is not necessary and redefines it as a constant. For comparison purposes, it is shown again,

$$C_i(\lambda) \approx C_i = c_i [B_p(\bar{\lambda}_i, T_p) - \epsilon_s(\bar{\lambda}_i) B_s(\bar{\lambda}_i, T_s)] , \quad (2.9)$$

along with his final radiance model,

$$L(\lambda) = \sum_{i=1}^D C_i k_i(\lambda) + \epsilon_s(\lambda) B_s(\lambda, T_s). \quad (2.10)$$


This assumption is suspect because spectral information may be lost. The analogous model in this study is shown to be,

$$L'(\lambda) = \sum_{i=1}^D c_i k_i(\lambda) [B(\lambda, \widehat{T_s} \pm \Delta T) - \bar{B}(\lambda, T_s)] + \bar{B}(\lambda, T_s). \quad (2.11)$$

The derivation leading to this equation is found in Chapter 3. Ignoring for the moment the background terms in each model, two important differences can be noted here. The first is that the equivalent contrast term in (2.11), $[B(\lambda, \widehat{T_s} \pm \Delta T) - \bar{B}(\lambda, T_s)]$, is still a function of wavelength. The spectral information contained in the background

emissivity will still be present. The second difference lies in which terms are part of the regression basis vectors and which are solved for as the coefficients of those vectors. Young's basis vectors are comprised of the absorption spectra, $k_i(\lambda)$. The plume temperature, T_p , and more importantly the temperature contrast, ΔT , are solved for in the regression rather than influence it. In (2.11), the basis vector set is comprised of the absorption spectra multiplied by the contrast term. By defining both positive and negative values for ΔT , the regression model will account for gases in either emission or absorption. The only term that is solved for in the regression is c_i .

Young does not account for the effect of the temperature of the gas on its absorption spectrum. He states that, "With the exception of carbon dioxide, water vapor and nitrogen dioxide, the spectral absorption coefficients are independent of temperature in the range relevant to this study (~ 15 to $\sim 80^\circ\text{C}$).". He claims that temperature affects these gases in magnitude only. A preliminary study done here contradicts this assumption. It has shown that temperature variations have a visible effect on both the magnitude of spectral features and the overall shape of the spectrum. This phenomenon is depicted in Figure 2.2. Here it can be seen that the absorption of SF_6 at $10.53\ \mu\text{m}$ is larger when the gas is at 5° , while at $10.60\ \mu\text{m}$ the absorption is larger when the gas is at 50° . The result is that the overall shape of the absorption spectrum is different for the two extreme temperatures. In addition, the height of the dominant peak at $10.55\ \mu\text{m}$ is scaled as a function of temperature. Effects such as this are not accounted

Figure 2.2: Effect of Temperature Variability on SF_6 Absorption Spectrum

for in a single-temperature model.

Finally, all of the gas studies mentioned above require prior knowledge of which gases are present. The SVD method employed by Lisowski and Cook, the computation of a correlation coefficient over spectral features, and the spectral matched filter used by Marinelli and also by Young all perform a detection of gases that have been previously

identified. The scope of this research is significantly different in that the user will have no knowledge as to which gases are present in the scene. The algorithm developed here requires no user interaction. This goal can not be met by any of the previously-discussed works in this chapter.

Chapter 3

Methodology

3.1 Sources of Data

3.1.1 Real Imagery

The Spatially Enhanced Broadband Array Spectrograph System (SEBASS) has been used in the detection and identification of gaseous effluents in previous works. In particular, this sensor provided the data used in Young’s work discussed in Section 2.2. This data set is described in greater detail below. Unfortunately these images are not available for use in published work, and are therefore not used in this study. The merits of testing the identification algorithm derived in this thesis on this data set are discussed in Chapter 6. However, since the synthetic imagery generated in the DIRSIG environment is based on the SEBASS sensor, a brief discussion of the instrument is nec-

essary. The SEBASS sensor is a hyperspectral imaging spectrometer created, owned and operated by the Aerospace Corporation. SEBASS contains two spectrometers. The first operates in the MWIR from 2.9–5.2 μm with a spectral resolution of $\approx 0.025 \mu\text{m}$, while the second senses in the LWIR from 7.5–13.6 μm at $\approx 0.05 \mu\text{m}$ increments. The detector in each spectrometer is 128×128 pixels. The pixel spread in the across-track direction provides 128 samples while the along-track direction provides the 128-band spectral dimension. The image of the ground is dispersed in the spectral dimension by a prism. Young’s imagery utilized all but four of the channels in the LWIR spectrometer, a span of 7.67–13.42 μm in $J = 124$ spectral bands. [Young, 2002].

The sensor imaged a commercial refinery in southern California at an altitude of 6000 ft. This altitude, coupled with an angular resolution of 1.1 mrad per pixel, resulted in a GSD of ≈ 2.0 m. The total across-track width of the scene was $128 \text{ pixels} \times 2.0 \text{ m/pixel} = 256 \text{ m}$. Nothing was known about the manufacturing processes at the site and no ground truth data was acquired at the time of the overflight. The only auxiliary data collected was the wind speed and direction at three-hour intervals. Using a spectral matched filter, Young has shown that there is one SO_2 , one C_2H_4 and two NH_3 plumes in this scene [Young, 2002]. The procedures leading to these results were described in Chapter 2.

3.1.2 Synthetic Imagery

Synthetically-generated imagery was the main source of image data used in this project. Data generated by the Digital Imaging and Remote Sensing Image Generation (DIRSIG) model was used 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 accuracy of the model has been verified in a recent study [Hattenberger, *et al.*, 2003]. Field data obtained from a ground-based spectrometer was recreated in DIRSIG. Standing on the ground, the operator took two measurements, one on the plume and one off. The off-plume (background) spectrum was then subtracted from the on-plume (target) spectrum. The goal of the DIRSIG study was to match this difference spectrum. The result of the study, shown in Fig. 3.1, clearly shows the ability of the JPL model (described below) to simulate real-life plume phenomenology. The measured gas features seen between 800 and 1200 cm^{-1} are matched in location, shape and height by the simulated plume. The negative radiance values in the real spectrum below 800 cm^{-1} and beyond 1300 cm^{-1} are not gas absorption features, but rather noise artifacts.

The plume model within DIRSIG is the Jet Propulsion Laboratory (JPL) plume model. Spatially, the plume is constructed by modifying the gas concentration. In the downwind direction the concentration follows an exponential decay. In both the vertical and across-track direction the concentration follows a Gaussian distribution. The width of

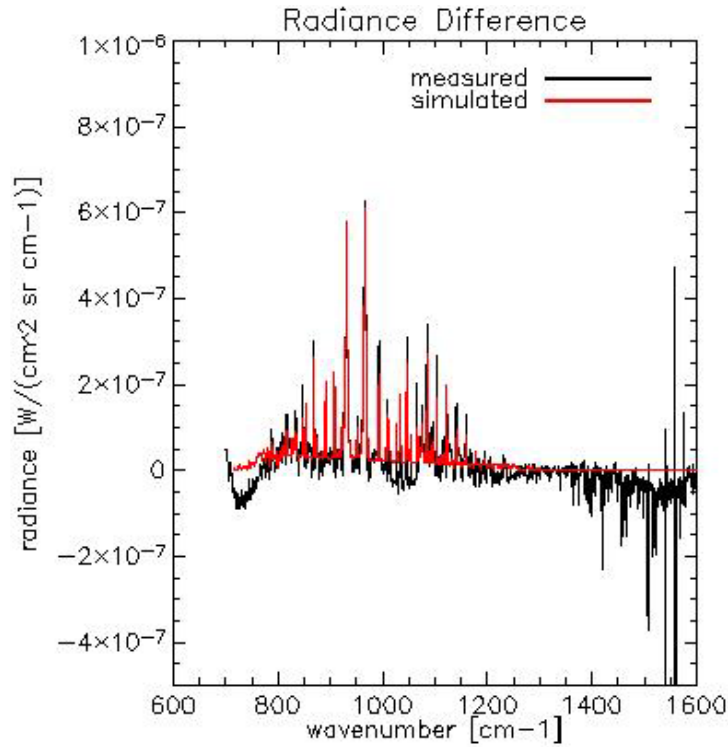


Figure 3.1: DIRSIG Plume Model Verification

the distribution increases as the plume travels further downwind. The result is a plume that is spatially continuous. There are no “clumps” or “gaps” but rather a gradual decay. It is the spectral, not spatial, aspect of the plume model that is of greater importance to this study.

Unlike other plume models that treat the gas as a single layer, the JPL model gives depth to the plume by defining it as a cloud of points. When DIRSIG traces a ray from the sensor down through the plume, it encounters one of these points at the top of the

plume. This point carries with it information about the temperature and concentration of each of the gases at that particular location. The ray is then traced down through the plume until it encounters a second plume point. The JPL model then uses the temperature and concentration of this point to integrate the radiative transfer along the path. This stepwise transfer is carried out until the ray has passed through the plume. The advantage of this approach is that it can reproduce phenomenology that a single-slab model can not. One such example is the case when the radiance emitted by the hot mid-section of the plume is absorbed by the cooler gases on the edges of the plume. In such a scenario all of the emission can be absorbed, and the plume rendered not visible.

The DIRSIG tool requires that the user specify a sensor model for each scene. This necessity also provides the ability for a user to simulate a scene as it would look to a particular sensor. This ability was utilized in this study in the generation of the synthetic data sets. The sensor model used in these simulations was constructed to model the SEBASS sensor. This sensor was chosen as a basis since it has been shown in the work done by Young that supervised gas identification can be performed on SEBASS data [Young, 2002]. The data sets, which are outlined in Section 4.1, were all constructed to be identical except for the number of plumes in the scene and the gas species they emit. The data generated by DIRSIG were image cubes of size (X, Y, J) , where X is the number of samples, Y is the number of lines, and J is the number of



Figure 3.2: Single-band image of Case 2 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.

bands. For all data cases, the number of bands is fixed at $J = 128$, while the values for X and Y vary across the data sets. A noise cube that emulates that of the SEBASS sensor was added to each of the test cases. A single band at $10.73\mu m$ from Case 2 is shown in Figure 3.2. The two plumes are clearly seen moving from the middle of the scene to the lower-right corner.

3.1.3 Gas Spectra

Spectra for thirty-one gases have been supplied, most of which have been measured at three temperatures, 5, 25 and 50° C. This resulted in a total of 88 absorption spectra for use in identification. All spectra were resampled to the wavelength scale in the DIRSIG imagery.

As expected, the absorption coefficient of most gases change in magnitude as a function of temperature. In addition, several gas species, such as SF₆, exhibit not only a change in height but also in shape due to temperature. This phenomenology has been discussed in Chapter 2 and can be seen in Figure 2.2. Because of this it is felt that it is important to include the absorption spectra from all three temperatures when populating the basis vector set. This is discussed in greater detail in Sec. 3.4.

3.2 Radiance Model

The radiance equation used to model the plume problem is derived here. If we assume that the mixing of gases in a plume can be described by a linear combination, then the gas problem is similar to that of a mixed pixel in the reflective region. Without compensating for the atmosphere and neglecting any reflected downwelled radiance, $L_d(\lambda)$, the model for a mixed pixel in the reflective region with C background materials

and D target materials is shown to be

$$L(\lambda) = \frac{E_0(\lambda) \cos \theta}{\pi} \left[\sum_{i=1}^D \alpha_i A_i(\lambda) + \sum_{j=1}^C \beta_j B_j(\lambda) \right] \tau_{atm}(\lambda) + L_u(\lambda), \quad (3.1)$$

where

$$L(\lambda) = \text{pixel radiance,}$$

$$E_0(\lambda) \cos \theta = \text{exoatmospheric irradiance at incident angle } \theta,$$

$$\alpha_i = \text{mixing coefficient for } i^{th} \text{ target material,}$$

$$A_i(\lambda) = \text{reflectance of target material } i,$$

$$\beta_j = \text{mixing coefficient for } j^{th} \text{ background material,}$$

$$B_j(\lambda) = \text{reflectance of background material } j,$$

$$\tau_{atm}(\lambda) = \text{atmospheric transmission, and}$$

$$L_u(\lambda) = \text{upwelled radiance.}$$

Extending this model into the LWIR region changes the relevant parameters in (3.1).

As discussed in Sec. 1.3, the direct solar term, $E_0(\lambda) \cos \theta$, is dominated by self-emitted radiance and can thus be ignored. As a result, the temperature and emissivity of a material become more important than its reflectivity. In addition (as will be discussed in detail later in Sec. 3.4.3), many common materials have a high emissivity, and therefore a low reflectivity, in the LWIR. The result of this is that neither the downwelled radiance nor the downward self-emitted plume radiance will, upon reflection off of the background, contribute significantly to the total radiance seen by the sensor. Figure

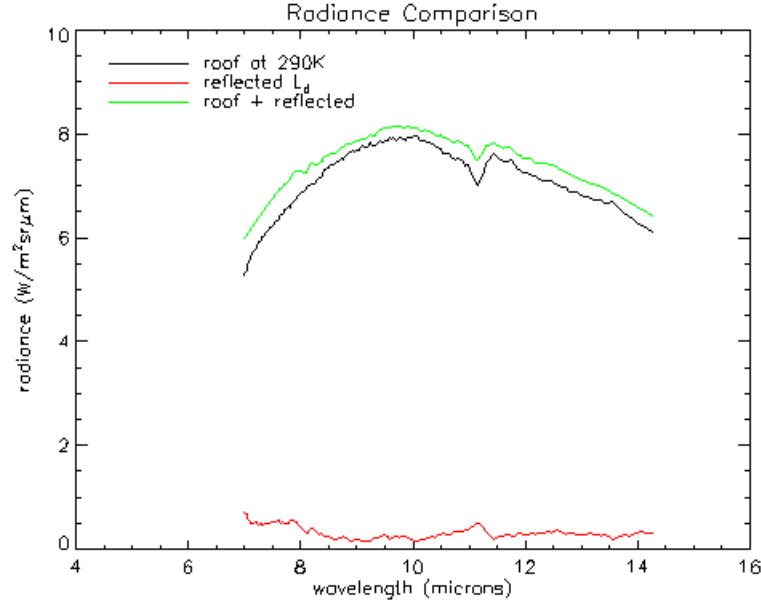


Figure 3.3: Comparison of self-emitted radiance with reflected downwelled radiance.

3.3 illustrates this claim. The black curve in this plot was generated by multiplying the Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) emissivity spectrum, $\epsilon(\lambda)$, for “gravel roof” by the Planck function evaluated at 290K. The red curve is the result of multiplying the MODTRAN-generated (mid-latitude summer) L_d by $1 - \epsilon(\lambda)$. The green curve is the sum of the black and red. The magnitude of the self-emitted radiance is changed little by the addition of the reflected downwelled term. Any changes in spectral shape are small and do not greatly impact the identification scheme. This allows us to express the radiance of pixels that do not contain a plume

(*i.e.*, a “background” pixel) as

$$L_s(\lambda) = \left(\sum_{j=1}^C \beta_j \epsilon_j(\lambda) B(\lambda, T_s) \right) \tau_{atm}(\lambda) + L_u(\lambda), \quad (3.2)$$

where

β_j = mixing coefficient for j^{th} background material,

$\epsilon_j(\lambda)$ = emissivity spectrum for the j^{th} background material, and

$B(\lambda, T_s)$ = Planck function evaluated at the surface temperature, T_s .

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’s self-radiation does not significantly reflect off of the ground and to the sensor. Also, 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

$$L_p(\lambda) = \left(\left[\sum_{j=1}^C \beta_j \epsilon_j(\lambda) B(\lambda, T_s) \right] \tau_p(\lambda) + \epsilon_p(\lambda) B(\lambda, T_p) \right) \tau_{atm}(\lambda) + L_u(\lambda), \quad (3.3)$$

where

T_p = plume temperature,

$\tau_p(\lambda)$ = transmission of the plume, and

$\epsilon_p(\lambda)$ = plume emissivity.

Equation (3.3) demonstrates 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 attenuated by the atmosphere. Conservation of energy says that if the gas is optically thin, we can set $\tau = 1 - \epsilon$ and arrive at

$$L_p(\lambda) = \left(\left[\sum_{j=1}^C \beta_j \epsilon_j(\lambda) B(\lambda, T_s) \right] [1 - \epsilon_p(\lambda)] + \epsilon_p(\lambda) B(\lambda, T_p) \right) \tau_{atm}(\lambda) + L_u(\lambda). \quad (3.4)$$

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

$$\epsilon_p(\lambda) = \sum_{i=1}^D c_i k_i(\lambda) \quad (3.5)$$

where

c_i = column density of gas i [ppm-m], and

$k_i(\lambda)$ = absorption spectrum of gas i [1/ppm-m].

Substituting this into (3.4) results in the final model for a pixel containing multiple background materials and multiple gases as

$$L(\lambda) = \left(\left[\sum_{j=1}^C \beta_j \epsilon_j(\lambda) B(\lambda, T_s) \right] \left[1 - \sum_{i=1}^D c_i k_i(\lambda) \right] + \left[\sum_{i=1}^D c_i k_i(\lambda) \right] B(\lambda, T_p) \right) \tau_{atm}(\lambda) + L_u(\lambda). \quad (3.6)$$

Note that this model still accounts for the effects of the atmosphere. The question of whether or not to compensate for these effects is addressed later in this thesis.

3.3 Regression Techniques

As discussed in Chapter 2, stepwise regression has been shown to be able to unmix hyperspectral pixels with a greater degree of accuracy than standard matrix regression. This is because only those basis vectors that contribute to the fit are included in the model. In this section, a thorough explanation of stepwise regression will be given. However, before we discuss the stepwise technique, we must first discuss standard matrix regression in order to illustrate the differences between the two methods.

3.3.1 Matrix Regression

The standard model for regression in matrix form is [Draper and Smith, 1981]

$$\mathbf{x} = \mathbf{A}\mathbf{f} + \boldsymbol{\epsilon}, \quad (3.7)$$

where

$$\mathbf{x} = (J \times 1) \text{ pixel vector,}$$

$$\mathbf{A} = (J \times N) \text{ matrix of } N \text{ basis vectors,}$$

$$\mathbf{f} = (N \times 1) \text{ vector of basis vector abundances, and}$$

$$\boldsymbol{\epsilon} = (J \times 1) \text{ error term.}$$

The pixel vector \mathbf{x} is modeled by a linear combination of the basis vectors present in \mathbf{A} . The aim of regression is to solve for the coefficients, f_i , where $i=0,1,\dots, N$, that are

present within the vector \mathbf{f} . These coefficients are solved for as

$$\mathbf{f} = \mathbf{A}^+ \mathbf{x},$$

where

$$\mathbf{A}^+ = (\mathbf{A}'\mathbf{A})^{-1} \mathbf{A}'$$

denotes the pseudo-inverse of \mathbf{A} . Substituting this into (3.7) we arrive at

$$\mathbf{f} = (\mathbf{A}'\mathbf{A})^{-1} \mathbf{A}' \mathbf{x}, \quad (3.8)$$

as the final solution for the basis vector abundances. In (3.8) a prime denotes the transpose of a matrix or vector. The error term, ϵ , is necessary since \mathbf{x} is rarely completely modeled by the vectors in \mathbf{A} . Mathematically, these residuals are expressed as

$$\epsilon = [\mathbf{I} - \mathbf{A}(\mathbf{A}'\mathbf{A})^{-1}\mathbf{A}'] \mathbf{x}, \quad (3.9)$$

where \mathbf{I} is the identity matrix,

$$\mathbf{I} = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{bmatrix}. \quad (3.10)$$

Matrix regression is efficient and is easy to implement; however, it does have its trade-offs. The model will force a solution for each basis vector in \mathbf{A} regardless of whether or

not the corresponding gas is present in the plume. This causes the regression to place too much emphasis on fitting the given basis vectors to the pixel, rather than determining which basis vectors are actually needed to model the pixel. This fact also means that the algorithm will require user intervention rather than being fully automated. An interactive step will likely be necessary in which the user reviews the abundances of each basis vector. They then must establish a threshold that will keep only those gases with significant abundances.

3.3.2 Stepwise Regression

Stepwise regression has been shown to be successful at selecting the appropriate basis vectors (endmembers) in unmixing hyperspectral pixels [Gross and Schott, 1996]. The process begins by populating the matrix \mathbf{A} with each of the M candidate basis vectors in turn. A statistical test (F -test) is used to determine which candidate vector best fits the data. The $M - 1$ remaining vectors are then added one at a time 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. 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), $\mathbf{x}'\mathbf{x} - \mathbf{f}'\mathbf{A}'\mathbf{x}$ [Draper and Smith, 1981]. The ANOVA structure is depicted in Table 3.1. The mean squared (MS) column is defined by dividing the SS by the corresponding degrees of freedom (dof). 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

$$\mathbf{x} = \mathbf{A}_{N-1}\mathbf{f}_{N-1} + \boldsymbol{\epsilon}, \quad (3.11)$$

where

$$\mathbf{x} = (J \times 1) \text{ pixel vector,}$$

$$\mathbf{A}_{N-1} = (J \times [N - 1]) \text{ matrix of } N - 1 \text{ basis vectors,}$$

$$\mathbf{f}_{N-1} = ([N - 1] \times 1) \text{ vector of basis vector abundances, and}$$

$$\boldsymbol{\epsilon} = (J \times 1) \text{ error term,}$$

then we can see how this fits into the ANOVA analysis in Table 3.1.

The MS terms from Table 3.1 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

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 3.1: Sum of Squares ANOVA Table

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 dof abbreviated as χ_J^2 [Draper and Smith, 1981]. Thus we can 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.

$$H_0 : \mathbf{f}_i = 0$$

$$H_1 : \mathbf{f}_i \neq 0$$

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

$$\frac{MS_{extra}}{MSE} \quad (3.12)$$

follows an $F_{1,J-N}$ distribution [Gross and Schott, 1996]. When attempting to add a basis vector to the model, the ratio in (3.12) 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 the vector is added to the model. Substituting the equations from Table 3.1 into (3.12), we arrive at the final expression for the F-statistic,

$$F_{1,J-N} = \left[\frac{\mathbf{f}'\mathbf{A}'\mathbf{x} - \mathbf{f}'_{N-1}\mathbf{A}'_{N-1}\mathbf{x}}{1} \right] \left[\frac{J - N}{\mathbf{x}'\mathbf{x} - \mathbf{f}'\mathbf{A}'\mathbf{x}} \right] \quad (3.13)$$

where $\mathbf{f} = \mathbf{A}^\dagger \mathbf{x}$.

Once the algorithm is complete the abundance coefficients are calculated for the final regression model. The algorithm also includes a check that all coefficients are significant, meaning that they must be larger than 1.0×10^{-5} .

The attractiveness of stepwise regression is that only those gases truly in the plume should be present in the final regression model.

3.3.3 Constraints on the Regression

Within each iteration of the stepwise regression process, the vector of abundance coefficients, \mathbf{f} , is recalculated. This solution eventually leads to the basis vector in question either being included in or excluded from the model. However we need to ensure that the coefficients do not allow the regression to “overfit” the pixel spectrum. Some degree of error is expected when performing such a regression, but allowing the model to

take on any coefficient necessary to minimize this error leads to unrealistic results. A specific example lies in the case where the predominant spectral feature in a candidate basis vector may occur at the same wavelength of the primary feature in the pixel spectrum, but at a much smaller magnitude. These differences can be as large as several orders of magnitude. With no restraint on the values of \mathbf{f} , this basis vector may be fit to the pixel with an unrealistically large coefficient. With this in mind, we require a new approach to solving for the fit coefficients.

In order achieve this safeguard, rather than using the unconstrained pseudo-inverse solution of (3.8), we instead impose either a partial or full constraint on the coefficients with each calculation of \mathbf{f} . In partially-constrained regression, we force the solution of all coefficients in \mathbf{f} to be positive. In fully-constrained regression, not only are the coefficients positive but they are also scaled such that they sum to 1.0. The value of 1.0 is chosen here not because it represents any realistic limit but rather for numeric simplicity. To ensure that these constraints are necessary, results were generated using unconstrained regression in which we do not enforce any constraints on the coefficients.

When solving for the coefficients in a constrained fashion, we can no longer use the standard pseudo-inverse solution, $\mathbf{f} = \mathbf{A}^\dagger \mathbf{x}$. This equation does minimize the error term, ϵ , but it allows the elements of \mathbf{f} to take on any value necessary to arrive at the best solution. Instead, a solution to \mathbf{f} is found through a complex series of matrix equa-

tions. An in-depth analysis into these equations can be found in the local presentation given by Rhody, 2004.

Stepwise regression can be thought of as a two-step process. In step one, the matrix of appropriate basis vectors is constructed based on their contribution to the fit. In step two, the abundances of each basis vector is solved for to create the final model. This interpretation of the stepwise method presents us with options as to how we implement the regression constraints. In this research, two sets of results were generated. The first involved using the partial constraint for the creation of the final basis vector model while using the full constraint to solve for their abundances. This method is referred to as Method A in all results. Since the solution for \mathbf{f} is only forced to be positive in each iteration of the stepwise code, the algorithm runtime for Method A is on the order of 1–3 hours. The second method, Method B, involved using the full constraint for both the creation of the final model and the solution of the coefficients. Since each stepwise iteration requires a more complex solution for the basis vector abundances, the runtime for Method B is on the order of 10–13 hours. In addition, the use of the partial constraint for both steps was implemented for one data set, but the results were inferior to the other two techniques and was therefore not continued.

3.4 Definition of Basis Vectors

Both regression methods described above work by fitting basis vectors in \mathbf{A} to the pixel spectrum, \mathbf{x} . The most daunting task in this research has been the definition of these basis vectors. Several problems had to be addressed in order to populate \mathbf{A} with vectors that both accurately represent the plume phenomenology and lead to meaningful results. The success of the algorithm was most dependent upon the solutions to these problems and the extent of the approximations made in the process.

3.4.1 Variables

Analysis of (3.6) will show that there are, neglecting the atmospheric effects, $2C+2D+2$ unknown terms to be solved for. They are

β_j = mixing coefficient for background material j of C ,

$\epsilon_j(\lambda)$ = emissivity spectrum for background material j of C ,

c_i = column density of gas i of D ,

$k_i(\lambda)$ = absorption spectrum of gas i of D ,

T_s = surface temperature, and

T_p = plume temperature.

The goal of this study was to identify $k_i(\lambda)$ by solving for their relative coefficients, c_i . In a perfect situation, the absorption spectra would populate \mathbf{A} while the coefficients would be the fractions in \mathbf{f} . This would, however, leave four other variables unaccounted for. The question remains which variables can we safely assume a value for and which we cannot.

The following sections detail how we avoid having to solve for β_j , $\epsilon_j(\lambda)$, and T_s explicitly. Solving for the plume temperature is not necessary as it is the surface-plume temperature contrast, rather than the plume temperature itself, that is of importance.

3.4.2 Atmospheric Compensation

Removing the atmosphere prior to gas identification would completely eliminate the last two terms in (3.6), and would thus simplify the process; however, many of the target gases we are looking for are also present in the atmosphere. Examples of such gases are CO_2 , N_2O , and NH_3 . Water vapor is not of concern here given that it has no absorption features in the wavelength region in question. In addition, diatomic molecules such as O_2 and N_2 do not exhibit features in the LWIR due to their molecular structure. Overcompensating for the atmosphere would then diminish the likelihood of detecting these atmospheric gases in the plumes. The standard method for atmospheric compensation in the visible and SWIR region is to remove $E_0(\lambda) \cos \theta$, $\tau_{atm}(\lambda)$, and $L_u(\lambda)$ in (3.1). Given a standard model for a non-plume pixel in the LWIR, neglecting

downwelled radiance, as

$$L(\lambda) = \epsilon(\lambda) B(\lambda, T) \tau_{atm}(\lambda) + L_u(\lambda),$$

we arrive at the standard model of a compensated pixel,

$$L'(\lambda) = \epsilon(\lambda) B(\lambda, T) = \frac{[L(\lambda) - L_u(\lambda)]}{\tau_{atm}(\lambda)} \quad (3.14)$$

where a prime denotes quantities that have been compensated for atmospheric distortion. With these considerations, (3.6) becomes

$$L'(\lambda) = \left[\sum_{j=1}^C \beta_j \epsilon_j(\lambda) B(\lambda, T_s) \right] \left[1 - \sum_{i=1}^D c_i k_i(\lambda) \right] + B(\lambda, T_p) \left[\sum_{i=1}^D c_i k_i(\lambda) \right]. \quad (3.15)$$

Implementing (3.14) requires estimates of the parameters $\tau_{atm}(\lambda)$ and $L_u(\lambda)$. Algorithms such as Autonomous Atmospheric Compensation (AAC), Canonical Correlation Analysis (CCA) and In-Scene Atmospheric Compensation (ISAC) exist for the purpose of estimating these parameters. AAC and CCA are both model-based approaches whereas ISAC is scene-derived. Each technique provides an estimate of $\tau_{atm}(\lambda)$ and $L_u(\lambda)$ to a different spatial resolution. CCA does so on a pixel-by-pixel basis, AAC segments the image into blocks of many pixels and solves for each block, and ISAC provides a single estimate for the entire image. In addition, the image processing software ENVI contains an atmospheric compensation algorithm that is very similar to ISAC [Young, *et al.*, 2002]. This technique was the lone method used in this research.

3.4.3 Background Characterization

The at-sensor radiance, $L'(\lambda)$, in (3.15) is a result of a product of several unknown terms, β_j , $\epsilon_j(\lambda)$, T_s and c_i . This demonstrates the non-linearity of the radiance model, which theoretically prohibits the identification of gases via linear regression.

If we can eliminate the need to solve for each β_j and each $\epsilon_j(\lambda)$ by clustering the background terms into an *effective* self-emitted, surface-leaving radiance, $\bar{B}(\lambda, T_s)$, we can transform the radiance model into a linear equation. Clustering the C background materials into $\bar{B}(\lambda, T_s)$ can be done in a number of ways. One method is to take the mean of all non-plume pixels. A second technique would require the user to manually mask out the regions surrounding the plume pixels and take their average. Both methods have obvious detriments. The first has no way to take into account the variation of surface materials across the entire extent of the plume. The second method, requiring user interaction, means that the algorithm could no longer operate autonomously.

The solution to this problem used in this study was to model the background on a per-pixel basis. The Maximum Distance Method (Max-D), an in-house endmember extraction algorithm, was run on all non-plume pixels [Bajorski, *et al.*, 2004]. 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 stepwise regression,

Method A. The goal here is not for the endmembers to model the pixel very accurately. Significant residuals are not only expected but also desired in this case. Allowing the coefficients to take on any value eliminated the residuals too effectively, thus prompting the use of constraints on the regression. These values of \mathbf{f} 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.

By deriving an estimate of the surface defining an effective surface-leaving radiance in (3.15) we can rewrite the radiance model as

$$L'(\lambda) = \bar{B}_{(x,y)}(\lambda, T_s) \left[1 - \sum_{i=1}^D c_i k_i(\lambda) \right] + B(\lambda, T_p) \left[\sum_{i=1}^D c_i k_i(\lambda) \right]. \quad (3.16)$$

Close examination of (3.16) will show that it too is a non-linear problem. It attempts to solve for both T_p within the Planck equation and $c_i k_i(\lambda)$. This issue will be addressed in the next section where a substitution for T_p , the plume temperature, will be developed.

3.4.4 Emission vs. Absorption

To this point we have been referring to $k_i(\lambda)$ as the absorption spectrum of the i^{th} gas, however, as displayed in (1.1), energy can be both absorbed and emitted. The spectra discussed in Section 3.1.3 relate to the former case. If the plume were to be seen in emission, the gas features would be inverted. In order to identify gases in these conditions, the basis vector set must account for both absorption and emission features.

The plume is seen by the sensor as being in absorption whenever the effective plume temperature, $T_{p, eff}$, is less than the effective surface temperature, $T_{s, eff}$ and in emission when $T_{p, eff} > T_{s, eff}$. Defining a temperature difference as $\Delta T = T_p - T_s$, we can rewrite the plume temperature as $T_p = T_s \pm \Delta T$. Using this formulation, the gas is seen in emission when ΔT is positive and in absorption when it is negative. Assuming, for the moment, that there is only a single gas in the plume, we can rewrite (3.16) as

$$L'(\lambda) = \bar{B}_{(x,y)}(\lambda, T_s) - \bar{B}_{(x,y)}(\lambda, T_s) c k(\lambda) + B(\lambda, T_p) c k(\lambda). \quad (3.17)$$

Grouping similar terms yields,

$$L'(\lambda) = c k(\lambda) \left[B(\lambda, T_p) - \bar{B}_{(x,y)}(\lambda, T_s) \right] + \bar{B}_{(x,y)}(\lambda, T_s). \quad (3.18)$$

Before substituting the new expression for T_p , we must first derive a method for determining T_s . We obtain this value by inverting the Planck function to solve for a brightness temperature spectrum, $T_{s, br}(\lambda)$ as shown in (2.1). The solution to the Planck function in this case is the background estimate, $\bar{B}_{(x,y)}(\lambda, T_s)$. This value is used to allow for the estimation of T_s on a per-pixel basis. At this point, there are options for extracting the surface temperature estimate, $\widehat{T_s}$, from the brightness temperature spectrum. The maximum, minimum or average of $T_{s, br}(\lambda)$ are all viable options. For this research, the maximum of $T_{s, br}(\lambda)$ was chosen since the mean surface emissivity is likely to be less than 1.0 and therefore decreases the effective surface temperature. Now that we have arrived at a suitable estimate for $\widehat{T_s}$, and once again taking the case

of multiple gases into consideration, we can rewrite 3.18 as

$$L'(\lambda) = \sum_{i=1}^D c_i k_i(\lambda) \left[B(\lambda, \widehat{T}_s \pm \Delta T) - \bar{B}_{(x,y)}(\lambda, T_s) \right] + \bar{B}_{(x,y)}(\lambda, T_s). \quad (3.19)$$

By defining several positive and negative values for ΔT , we account for both the emission and absorption of the gas. This expression also circumvents the problem of solving for T_p . For this research, ΔT was fixed at 5 values: -10, -5, 0, 5, and 10° C. Multiplying this by the 88 absorption spectra, we arrive at a total of $M = 440$ candidate basis vectors to start with in the stepwise routine. This number of basis vectors would be prohibitive using standard linear regression. Stepwise regression allows us to deal with this amount of basis vectors and therefore formulate an answer from a more robust set of possibilities.

3.5 Formulation of the Identification Algorithm

Through several assumptions and substitutions, we have transformed the complete radiance model of (3.6),

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

into the final model of (3.19),

$$L'(\lambda) = \sum_{i=1}^D c_i k_i(\lambda) \left[B(\lambda, \widehat{T}_s \pm \Delta T) - \bar{B}_{(x,y)}(\lambda, T_s) \right] + \bar{B}_{(x,y)}(\lambda, T_s),$$

a form that resembles the regression equation of (3.7),

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

where

$$\mathbf{x} \Rightarrow L'(\lambda) - \bar{B}_{(x,y)}(\lambda, T_s)$$

$$\mathbf{A} \Rightarrow k_i(\lambda) \left[B(\lambda, \widehat{T_s} \pm \Delta T) - \bar{B}_{(x,y)}(\lambda, T_s) \right]$$

$$\mathbf{f} \Rightarrow c_i.$$

The steps involved in this constitute the bulk of the identification algorithm. Highlights of the process include the fact that the background to the plume is modeled on a per-pixel basis from image-derived endmembers, the fact that both emission and absorption spectra were built into the basis vector library, and that no prior knowledge as to which gases are present in the scene is necessary for identification. Figure 3.4 shows a pictorial representation of the algorithm to provide an overview of the workflow.

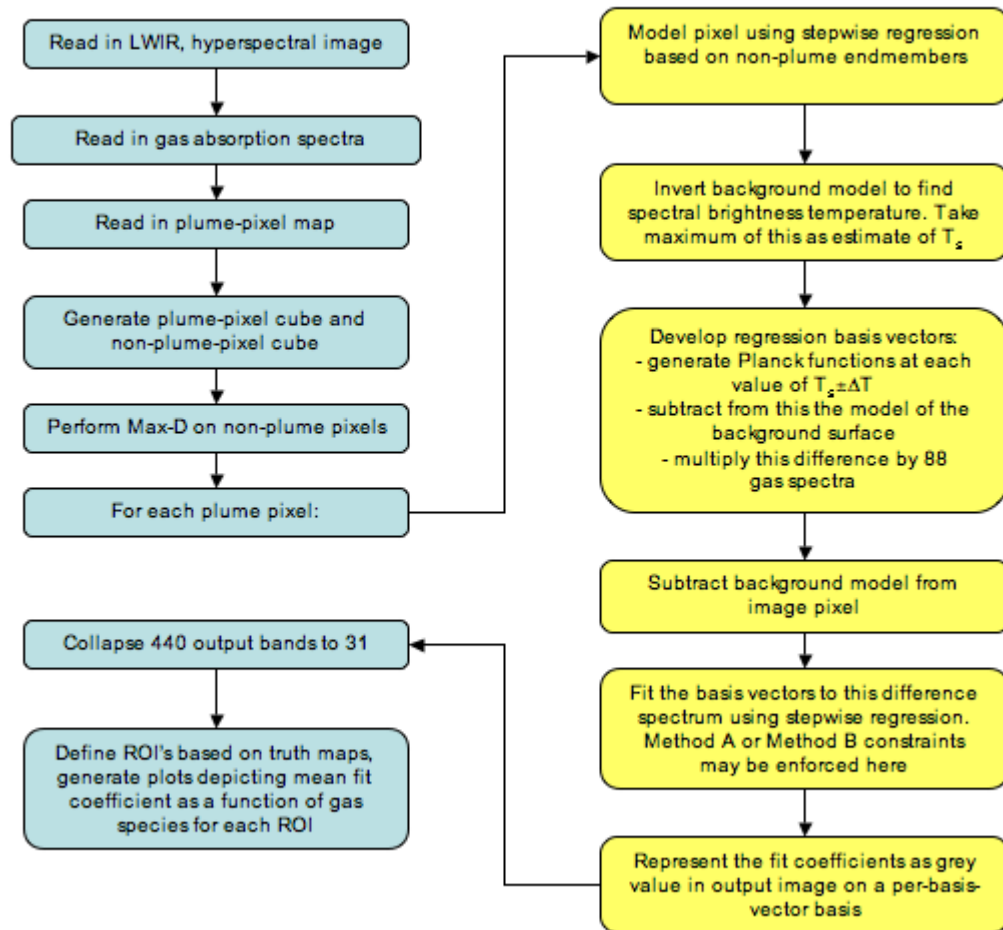


Figure 3.4: Workflow of the identification algorithm.

Chapter 4

Results

4.1 Test Cases

The algorithm derived in this research was not a fixed workflow. Rather, there are several variables that were changed to generate different sets of results. Fully exploring the effects of every variable in the algorithm would be beyond the scope of this research. With this in mind, in every case the background was characterized with basis vectors solved for using partially constrained stepwise regression and the probability for the F-statistic was set at 99%. The only change in the algorithm that was fully explored was the manner in which the regression was constrained throughout the stepwise routine. As mentioned in Section 3.3.3, these different approaches have been named Method A and Method B. The F-statistic variable was changed for one particular set of results to gauge the impact it would have on the result.

Data Set	Number of Plumes	Gas Species	Release Rate
Case 2	2	NH ₃ , F114	50 g/s
Case 3	1	F114, TCE-1122	50 g/s
Case 4	2	NH ₃ , F114	2.5 g/s
Case 4b	2	NH ₃ , F114	0.25 g/s
Case 5	1	NH ₃ , TCE-1122	0.25 g/s
Case 6	2	CH ₄ , DCLP12	0.25 g/s
Case 7	2	unknown gases 1 and 2	gas 1: 0.25 g/s gas 2: 0.5 g/s

Table 4.1: DIRSIG Data Sets.

Aside from the variables mentioned above, the algorithm was run on several incarnations of the same DIRSIG scene. The number of plumes, number of gases within a plume, gas species contained in a plume, and gas concentration within the plume were all varied across several image sets. In addition, several of these images were atmospherically compensated using the standard routine in ENVI, further adding to the number of available data sets. Table 4.1 depicts the data sets used in the study along with their plume configurations. In this table, in any instance in which there is only one plume in the scene, it is assumed that all gases listed coexist within the single plume. For data sets in which there are two separate plumes in the scene, each plume contains one unique gas and the plumes do not interact spatially at any point in the image.

In all, there were seven data cases, each of which was tested using both Method A and Method B of the algorithm. Three of these data sets, Case 2, Case 4b and Case 5, had their atmosphere removed using the ENVI compensation routine. Each of these were also tested on both algorithm methods. Finally, the Case 2 image with no atmospheric compensation was run using Method A of the algorithm with a 95% significance level in the calculation of the F-statistic. There were 21 sets of results generated that will be discussed in this chapter. As mentioned in Section 3.3.3, the partial constraint was enforced for both the creation of the final basis vector model as well as the solution of the coefficients. Since these results were judged to be inferior to either Method A or B, they are not included in this discussion. From these 21 results, the best combination of variable settings can be determined and the best algorithm configuration is selected.

4.2 Algorithm Output

The raw output from the algorithm is an image cube of size (X, Y, M) , corresponding to the size of the input image. Each of the M “bands” in the output detection maps correspond to one of the candidate basis vectors fed into the regression. For pixel (x, y) , the value for each band is 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 is left at 0.0.

This image cube is then refined in a secondary routine that “collapses” the output down to 31 bands, one band per gas species used in this study. The value assigned to pixel (x, y) in the “flattened cube” in, for example, the Freon-114 band, is the average of the 15 basis vector bands in the raw output image that correspond to Freon-114. The fact that there are 15 basis vectors, and therefore 15 bands in the raw output, is the result of there being three absorption spectra for Freon-114 in the gas library (one curve for each gas temperature of 5, 25 and 50° C) multiplied by the five values used for ΔT , $[-10, -5, 0, 5, 10]$. As mentioned in Section 3.1.3, not all gases have three absorption spectra in the library. For these cases, fewer than 15 bands were averaged to arrive at the final value in the “flattened” image. These averages were then normalized by the per-pixel total abundance coefficient for all 440 bands for that location. This was done to scale all results between 0.0 and 1.0 for ease of analysis. Once the output has been reduced from 440 bands, the “spectral” dimension was changed to a *gas index*. Each “band” now represents the returns from one particular gas. Table 4.2 summarizes the gas-species/gas-index correspondence. The resulting detection maps for all cases are shown in Appendix A. Only the bands corresponding to the gases actually in the plume are shown for simplicity. All maps have been inverted such that a black pixel indicates a strong return. An example detection map from Case 4 is shown in Figure 4.1.

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 ₆ H ₅ F	Fluorobenzene	19	F134a	Freon-134a
4	C ₆ H ₆	Benzene	20	F218	Freon-218
5	CCl ₄	Carbon Tetrachloride	21	H ₂ O	Water
6	CH ₃ Cl	Methyl chloride	22	HCHO	Formaldehyde
7	CH ₃ SH	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	TCE-1122	Tetrachloroethane
14	EDB	1,–Dibromoethane	30	TCE	Trichloroethylene
15	EDC	1,2-Dichloroethane	31	VCL	Vinyl chloride
16	F114	Freon-114			

Table 4.2: 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.




Figure 4.1: Example detection map showing the Method A results from the Freon-114 plume in Case 4.

4.3 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 cube described above. Looking at a single band of the output image gives no indication of how the other bands look. 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 spatial extent 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 logarithmic 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. Figure 4.2 depicts the ROI's used in the analysis of results from Case 4 imagery.

As expected, the column density is highest near the stack in the center of the image and tapers off in all directions. The wind is blowing from the upper-left to the lower-right of the scene, which extends the higher column densities downwind. 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. An example of this series of scatterplots is shown in Figure 4.3, which was generated using the ROI's shown in Figure 4.2. Comparing the two figures shows that there is one curve in the plot for each ROI. For any given test case, the number of curves is dependent upon the range of column densities in the original DIRSIG image.




Figure 4.2: Snapshot of the ROI's used for Case 4 results. Red pixels indicate a column density range of 0–1, green 1–10, blue 10–100, and orange 100–1000. All column densities are reported in units of ppm-m. In addition, each plume had an ROI consisting of the entire plume which is not shown.




Figure 4.3: Example scatterplot of the average return for each ROI as a function of gas index. This particular plot shows the Method A results from the Freon-114 plume in Case 4.

Chapter 5

Analysis

Before we begin analyzing each result and its corresponding series of ROI plots, we need to understand exactly what the plots indicate and what conclusions can be drawn from them. To begin, we need to identify what the plot for a “perfect” identification would look like. Keeping in mind that the fit coefficients were normalized during the “flattening” stage, we can imagine that the ideal scatterplot for the Freon-114 plume in a Case 2 result would have a value of 1.0 at gas index 16 and 0.0 for all other gas indices. Ideally a response of 1.0 would occur for each ROI curve while all other gas indices would have a value of 0.0.

This ideal case occurs. There are pixels in which the incorrect gas species has a larger fit coefficient than the correct gas. If there are enough of these pixels within an ROI, that plotline will indicate that the wrong gas is present in the plume. As expected,

these false identifications are more likely to happen in ROI's corresponding to smaller concentration path-lengths. For example, if the return at gas index 4 was higher than at gas index 16 for the 10–100 ROI, benzene (gas 4) would be falsely identified over the actual species, Freon-114 (gas 16).

With this in mind, we must look at the results in each ROI and across the entire range of the plume. False identifications at low concentrations downwind might be balanced by a correct identification at higher concentrations upwind. Since false identifications are more likely to occur at lower gas concentrations, the results from high-concentration ROI's are “weighted” more than the others. If the highest-concentration ROI determined that gas A was present, while two lower-concentration ROI's indicated gas B, it is more likely that gas A is in fact the correct answer. This scenario, which did arise during the course of this study, is illustrated in Figure 4.3. In this result, from the Case 4 imagery, Freon-114 is correctly identified in only the largest of the five ROI's (100–1000 ppm-m range), yet is the correct answer. Even the plotline representing the entire plume indicates that either gas 18 (freon-12) or gas 28 (SO_2) is present rather than Freon-114 (gas 16).

This idea can also be seen in the detection maps in Appendix A. In the current example from Case 4, we can see in Figure 4.1 that the strong detections are concentrated near the origin of the plume, while downwind they are both weaker and more sparse. This reinforces the notion that the results from higher concentrations give a better indica-

tion of the correct gas species. This is the criteria that will be used in the analysis of the results. A successful identification will be declared if both the largest-concentration ROI indicates the correct gas and the detection map corresponding to that gas show strong returns across the majority of the plume pixels.

5.1 Separate Plume Results

The ability to correctly identify gas species in single-gas plumes is directly related to the concentration of the gas. Case 2 results (Figures B.1 – B.10), which contained the largest amounts of gas, were the most successful regardless of atmospheric compensation or change in F-statistic. Case 4 results (Figures B.13 – B.16) show a slight drop-off in lower-concentration ROI's. In Case 4b results (Figures B.17 – B.24), only the ROI with the largest concentration provided a correct identification (with the exception of the NH_3 plume using Method B). In some cases, such as the Freon-114 plume using Method B, the Freon-114 returns for other ROI's are nearly 0.0, well below the returns from other gases.

It should be noted, however, that in both Methods A and B for the Freon-114 plume (Figures B.18 and B.20), the returns for freon-12 are strong for all ROI's. In addition, they fall off in strength as one would expect. The 10–100 ROI has the strongest return while the 0–1 ROI exhibits the smallest. This indicates that freon-12 is very likely to be present in the plume. Since freon-12 is in the same family as Freon-114, it




Figure 5.1: Absorption spectra from the database for Freon-114 and freon-12. The spectra have been resampled to match the band centers of the SEBASS sensor.

has a similar molecular structure and therefore a similar absorption spectrum. The absorption spectra for the two gases are shown in Figure 5.1. The other gas with large returns across all ROI's in Figure B.18 is gas 28, SO_2 . 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 according to the correct identification criteria established. In addition SO_2 has no strong returns in the Method B results. In all, it is encouraging that at low gas concentrations, the algorithm was able to generically identify freon in the Freon-114 plume.

Another important observation that can be drawn from the Case 4b results is that, as the concentration of the gas decreases, the modeling of the background becomes more

important. This is seen most clearly in Figure A.6. Here, strong detections of Freon-114 downwind from the stack are very dependent upon the underlying surface structure. Specifically, the detections only occur over homes with asphalt-shingle roofing. This makes sense for two reasons. First, the spectral emissivity of shingles is relatively flat in the LWIR. Second, since there are many roofs in the scene, it is likely that these pixels became their own endmember. This made it easier for the algorithm to accurately model the background, which resulted in small residuals when it was subtracted from the image pixel. When the background is not modeled well, these residuals can have a significant impact on the shape of the difference spectrum and therefore result in incorrect identifications. In order to maximize the effectiveness of the background modeling, and therefore the effectiveness of the algorithm as a whole, the proper endmembers must be selected.

The only separate-plume gas that was not identified was CH_4 in Case 6. These results can be seen in Figures B.29 and B.31. In neither method does CH_4 have a strong return in any ROI. The inability of the algorithm to detect this gas is primarily due to the lack of strong features in the absorption spectrum of CH_4 . This is illustrated in Figure 5.2. The only features present in the CH_4 curve are a series of weak absorptions near 8 microns. By the time these features are seen by the sensor, they have been mixed with background features and propagated through the atmosphere. The resulting spectrum then contains virtually no traces of CH_4 , making it difficult to identify this gas. In




Figure 5.2: Absorption spectra from the database for Freon-114 and CH_4 . The spectra have been resampled to match the band centers of the SEBASS sensor.

addition, because the sensor is essentially seeing only background and atmosphere, no single gas is clearly identified by the algorithm.

5.2 Mixed Gas Results

In both cases containing mixed gases, Cases 3 and 5, the algorithm was not successful at identifying both gases present in the plume. In Case 3, as can be seen in Figures B.11 and B.12, Freon-114 was found convincingly using both methods. The detection maps in Figure A.4 also hint at a strong detection of Freon-114. TCE-1122, however, was not detected using both results. The ROI plots from Method A show two other gases, freon-12 and SF_6 , with higher returns than TCE-1122. The plots from Method

B do indicate TCE-1122, but not at the highest concentrations. When this was seen in Case 4b results, it was concluded that the gas was unlikely to be in the plume. Here, however, we can conclude it is present by analyzing its detection map in Figure A.4. The fact that the strong detections lie at the perimeter of the plume explains why the highest-concentration ROI's have smaller returns.

The algorithm did not perform as well on Case 5. This, however, was expected since the plume in Case 5 was being released at a much slower rate than in Case 3. The only definitive success in Case 5 was identifying NH_3 using Method B. Ammonia was found using Method A, but with less confidence than freon-12. Without prior knowledge as to which gas was in the plume, one would not be able to conclude that NH_3 was the correct gas over freon-12. TCE-1122 had the third-strongest return in the 10–100 ROI, but performed poorly in the other regions. Both methods failed to correctly identify either gas using either method in the atmospherically compensated image.

5.3 Blind Test Results

The purpose of the blind test case, Case 7, was twofold. First, is the algorithm able to give a clear indication of which single gas was in the plume? This is important because in all other cases, we were able to focus on a particular gas when analyzing the output. This is likely to have skewed the analysis in favor of the correct gases. Second, assuming that there is a clear solution, is that solution correct?

As Figures B.33 through B.36 clearly show, the algorithm was successful on both accounts. Regardless of which algorithm method was used, Gas 1 was identified as phosgene and Gas 2 was identified as C_6H_5F . Although Method A did perform better in the C_6H_5F plume, correct identification was possible using either method.

5.4 Effects of Atmospheric Compensation

Results from atmospherically native gas plumes indicate that compensation hurts our ability to correctly identify the gas. This effect is illustrated well in the differences between Figures B.17 and B.21 and Figures B.19 and B.23. Here, it can be seen that regardless of which algorithmic method was employed (Method A or B), the response for NH_3 was greatly diminished. In the non-compensated results, the NH_3 returns are strong enough to indicate that it is the only gas in the plume. In the compensated results, other gases are identified in the plume while the NH_3 returns resemble the “background noise” in the ROI plots.

When dealing with gases that are not found natively in the atmosphere, the opposite effect was found; compensation provided a slight improvement in the identification. As was the case with the native gases discussed above, this effect is depicted quite well in the results from Case 4b. The results from compensated imagery in Figures B.22 and B.24 show a significant improvement in Freon-114 returns over the non-compensated results of Figures B.18 and B.20. However, the increase in identification ability in

non-native gases is not sufficient enough to offset the degradation in performance with native gases. Freon-114 was identified adequately without compensation. Therefore, we must reason that if we were to have no prior knowledge as to which gases might be present in the plume, and therefore no indication whether the gas is native to the atmosphere or not, then it is best to not remove the atmosphere.

5.5 Effects of the Change in the F -Statistic

The net effect of changing the significance parameter in the F -statistic from 0.99 to 0.95 is seen largely in the change in the false identifications. The results for the Freon-114 and NH_3 plumes using a probability of 0.99 are seen in Figures B.2 and B.1. In these plots, several other gases have significant returns, creating a “background noise” effect to the ROI plots. The results generated using the other probability are shown in Figures B.6 and B.5. They clearly show that the returns for incorrect gases have been “consolidated”. Rather than there being peaks at several gas indices, there are only two or three other gases that have been identified. In the Freon-114 results, these peaks are also very diminished, while in the NH_3 results, they are merely consolidated. In both plumes it can be said that these results are preferred to those at the regular probability level.

The fact that a lower probability threshold improved the results can be explained. During the stepwise regression routine, a high probability threshold forces the regression

to take on extra basis vectors in order to minimize the residuals. A lower threshold would permit the regression to finish with a larger residual. Since most of these residuals stem from inconsistencies in the modeling of the background, we do not necessarily want to minimize them. We can also explain this effect by analyzing the change in F -statistic based on the change in probability. A high probability results in a lower F -statistic, which makes it easier to add new basis vectors to the model that perhaps should not be added. Assuming identical degrees of freedom, lowering the probability raises the F -value making it more difficult to add new basis vectors. With this in mind, we can conclude that in general, the basis vectors for the correct gases are among the first few included in the model. Allowing the regression to come to a solution sooner via a lower probability value prevents more vectors from being needlessly included in the model.

Chapter 6

Conclusions

6.1 Summary

Judging by the results analyzed in Chapter 5, we can conclude that the algorithm is successful at identifying gas species in separate plumes in synthetic imagery. In scenes in which the gas is being released at low rate (< 1.0 g/s), the ability to identify is limited to the region in the immediate vicinity of the stack. If it is possible that the gas being released is also native to the atmosphere, then the image should not undergo atmospheric compensation using the implementation of the ISAC algorithm found in ENVI. Spectral overlap, which was initially thought to be a major obstacle in this process, was found to have a minimal effect once the constraints were placed on the regression.

6.1.1 Methods for Imposing Regression Constraints

The question as to which type of constraints to impose on the regression, Method A or B, is still difficult to answer. This study provided us with 17 sets of results upon which we can base this decision. Of the 17 sets, the results had no significant difference between Methods A and B for eight of them. This leaves nine sets of results that did show a preference for one method or the other. Of these nine, Method B had more sets of superior results than Method A by a count of six to three. Based on this, we could conclude that Method B is the preferred set of constraints.

However, a second look at these nine sets shows that the method with better results was actually necessary for correct gas identification in only two cases. The first instance, seen in Figure B.12 was in Case 3 where Method B was necessary to identify both Freon-114 and TCE-1122 in the same plume. In this case, Method A was able to only identify Freon-114 (Figure B.11). The second instance, seen in Figure B.26, was Case 5 in which Method B was needed to identify NH_3 . Method A was not able to identify either NH_3 or TCE-1122 in Case 5 (Figure B.25). Both instances in which one method was necessary for correct identification involved plumes with two gases mixed together. In addition, Method B was the algorithm able to identify the gases in both cases.

This data indicates that Method B, fully constraining the solution for \mathbf{f} at each point in the stepwise routine, should be the preferred implementation of the algorithm. However, recall that the run-time for Method B was approximately three times longer than

for Method A. This extra time can be prohibitive, especially if this result was to then be passed off to another algorithm for gas quantification. If the results from Method B were shown to be necessary for correct gas identification across all test cases, then the extra run-time would be warranted. Considering that this was the case for only two of the 17 sets of results, we conclude that the manner in which constraints are imposed on the regression is irrelevant. Results may improve slightly if the solution to the vector of abundances is fully constrained within each iteration of the stepwise code, but said improvement is too little to warrant the extra run-time.

6.1.2 Extending the Algorithm to Real Imagery

Despite any success the algorithm may have had on synthetic imagery, the issue still remains that it is intended to be used on real imagery. Regardless of how accurate the DIRSIG model may be, it still can not encompass all aspects of real imagery. For instance there are sensor artifacts and noise that vary in time in real imagery. The atmosphere may have attributes that can not be fully modeled as well. Finally, there is the question of how well DIRSIG can emulate the image-wide statistics and background clutter of real imagery. These are issues that arise with any use of synthetic imagery. In addition to these, there are issues that are unique to this research.

One of these issues is the number of plume pixels available to run the algorithm on. In real imagery, this number is much smaller. This is because these pixels are being

found by algorithms that can not possibly detect every single pixel that a plume gas is present in. Reducing the number of plume pixels in the image does not change how the algorithm operates. The only effect is shorter run-times. When the algorithm is complete and the output is condensed to a 31-band image, there are fewer pixels to average for making the ROI plots depicted in Appendix B. This means that any pixels containing incorrect identifications will have a greater influence on the final plots, possibly preventing the correct gas index from being identified. However the pixels found by the detection scheme are likely to contain plume gases at strong concentration path-lengths, which would promote correct identifications. These two opposing factors merit investigation.

The final issue in extending this work to real imagery is the fact that the analysis was performed on ROI's generated by thresholds in the concentration path-length of the gas. This was possible because DIRSIG outputs truth maps that display the concentration on a per-pixel basis. This is obviously not possible in real imagery. Worse, many of the correct identifications were made in only the ROI's corresponding to the largest concentrations. Monitoring the relative magnitude of the returns between ROI's was an important ability in our analysis.

While we may not be able to define ROI's based on the exact gas concentrations, we are able to do so on an approximate level. Visual inspection of Figure 4.2 shows that each ROI has an elliptical shape extending downwind. This feature can be simulated

via the ROI tool in ENVI. If the user were to open the plume-mask image in ENVI, they would be able to interactively draw several ROI's. The highest-concentration ROI can be drawn nearest the stack, with successive ROI's growing larger. All ROI's can either be drawn by hand or generated using the "grow" function in ENVI. Three ROI's should be sufficient to assign all plume pixels to either a "high", "medium", or "low" concentration ROI. This would then allow the user to analyze the algorithm output in the same manner as shown in this work.

6.2 Recommendations for Future Work

The abilities and deficiencies of the algorithm constructed in this work can be better understood by more investigations. Many of these analyses went beyond the scope of this thesis and therefore were not carried out here. In this section we discuss what these studies are, what they can teach us about the algorithm, and how they can be used to solve the problem of gas quantification.

6.2.1 Effect of F -Statistic Probability

The first portion of the algorithm that deserves further study is the effects of changing the probability at which we calculate the F -statistic. As discussed in Section 5.5, changing this probability from 0.99 to 0.95 resulted in the "consolidation" of false identifications and an overall improvement in the ability to identify the correct gas.

This result, however, was only generated for the Case 2 image. It would be interesting to see what effects this, along with other, changes in probability has on plumes with a smaller release rate. Perhaps this would enhance our ability to identify gases at smaller concentration path-lengths. In addition, it may help identify both gases in Cases 3 and 5.

6.2.2 Constraints On Concentration Path-Length

Another aspect that requires more investigation is the values we allow for the gas concentration, c_i . In both Methods A and B, the final solution for these values in the vector \mathbf{f} were constrained between 0.0 and 1.0. In reality, the gas concentration path-length can take on values up to 1000 ppm-m. If we were to refine the basis vectors to allow for such values, and if we were to analyze the algorithm output on a per-basis-vector level, we could then get an indication of the temperature and concentration of the gas on a per-pixel basis. These real-world estimates of concentration path-length could then be used as initial estimates in gas quantification algorithms.

6.2.3 Atmospheric Compensation

In Section 5.4 we concluded that using the standard atmospheric compensation algorithm found in ENVI hinders our ability to identify atmospherically-native gases. The question remains, however, if all atmospheric compensation algorithms have this effect as well. As mentioned in Section 3.4.2, there exist at least three other atmospheric

compensation techniques that are suited to imagery in the LWIR. Two of these techniques are model-based while the other is completely scene-derived. Since the ENVI technique is scene-derived, it stands to reason that at least one of the model-based approaches should be tested to determine what effects that may have.

6.2.4 Spectral Overlap

As discussed briefly in Section 2.2, the notion of performing gas identification across species-specific spectral windows has been shown to be successful. This technique was to be used in dealing with overlapping spectral features between two or more gas species. This problem was minimized once constraints were placed on the regression, but some instances can still be found. Figures B.18 and B.20 show that rather than finding Freon-114 in Case 4b, the algorithm identified Freon-12. Figure 5.1 shows the absorption spectra for both of these gases as well. Since both gases are considered freon, they have similar molecular structures and therefore similar absorption features. The fact that this result occurred indicates that spectral overlap may still be a problem in discerning between two similar gas species that are not in the same family.

The problem with spectral overlap is that developing an automated method to deal with it is difficult. On a per-pixel basis, the initial output from the stepwise regression must be analyzed, two or more gases with similar absorption features must be identified, the wavelengths at which this overlap occurs must be masked out, and the

regression must be run again. However this entire problem may also be alleviated if the spectral resolution of both the imagery and the spectral library were higher. Two gas spectra that appear to be similar when sampled at an interval of $0.05\mu\text{m}$ may have a distinguishing features when sampled at $0.01\mu\text{m}$. Since this algorithm was tested on synthetic data, the possibility is there to create a dataset at a much higher spectral resolution.

6.2.5 Effects of a Scene-Derived Background Estimate

Another aspect of the algorithm that should be investigated further is the effect of subtracting the background estimate, $\bar{B}_{(x,y)}(\lambda, T_s)$, from both sides of the final radiance model, which is reproduced in Equation 6.1,

$$L'(\lambda) - \bar{B}_{(x,y)}(\lambda, T_s) = c_i k_i(\lambda) \left[B(\lambda, \widehat{T_s} \pm \Delta T) - \bar{B}_{(x,y)}(\lambda, T_s) \right]. \quad (6.1)$$

This estimate is constructed using image-derived basis vectors which include effects from the atmosphere. By subtracting the estimate from the pixel spectrum, $L'(\lambda)$, on the left-hand side of the equation, the algorithm may be performing a crude version of atmospheric compensation. If so, then when the algorithm is run on atmospherically compensated imagery, the vector being modeled by the stepwise routine may have been twice-compensated. This would help explain the poor performance of the algorithm on the NH_3 plume in atmospherically compensated data sets.

In addition, by subtracting the background estimate from Planck curves on the right-

hand side of the equation, we may be generating a basis vector set with “negative” atmospheric effects. This would create basis vectors that are not physically realistic. One possible solution to this would be to model the background estimate with a Planck curve and use this in the creation of basis vectors. Further studies into this issue could provide a marked improvement to the algorithm.

6.2.6 Effect of Sensor Noise

The impact of sensor noise on the ability of the algorithm to identify gases should also be investigated. A single test case, Case 4b with a 5x increase in sensor noise, was tested using Method A of the algorithm. The added noise produced results that were dramatically different. Freon-114, which was overshadowed in Case 4b by Freon-12, was undoubtedly identified in this case. Ammonia, which was solidly identified in Case 4b, is identified as the eighth most likely gas to be in the NH_3 plume.

These results demonstrate that the relative amount of sensor noise can have a significant impact on the ability to identify gases. This impact is completely dependent upon the nature of the spectral signature for each gas. The identification of a gas such as Freon-114, which has strong, broad absorption features, can be aided by added sensor noise. while gases with a weak, high-frequency signature such as NH_3 , can be hindered. An in-depth study should be performed to determine both which gases are impacted the most, and at what noise level does the shit hit the fan, so to speak.

6.2.7 Real Imagery

Finally, the algorithm should be tested on real imagery. Specifically, the SEBASS imagery discussed in Section 3.1.1 is appropriate for testing since it would allow the results from this algorithm to be compared to the results from Young's identification scheme. This would allow for a direct comparison of this work with an established identification algorithm. In addition, testing the algorithm on real imagery would allow us to address the issues discussed in Section 6.1.2. The steps proposed there for extending this work to real imagery should be effective, but this assumption should be verified.

Bibliography

- [Bajorski, *et al.*, 2004] Bajorski, P., Ientilucci, E., and Schott, J. “Comparison of Basis-Vector Selection Methods for Target and Background Subspaces as Applied to Subpixel Target Detection”, *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery X*, Shen, S. and Lewis, P., Editors, Proceedings of SPIE Vol. 5425 (2004), pp. 97–108
- [Banwell, 1972] Banwell, C. N., *Fundamentals of Molecular Spectroscopy*, 2nd ed., Berkshire: McGraw-Hill Book Company, 1972
- [Clark, *et al.*, 1990] Clark, R. N., Gallagher, A. J., and Swayze, G. A., “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*, JPL Publication, pp. 176–186, 1990a
- [Draper and Smith, 1981] Draper, N.R. and Smith, H., *Applied Regression Analysis*,

2nd ed., New York: John Wiley & Sons, 1981

- [Foy, *et al.*, 2002] Foy, B.R., Petrin, R.R., Quick, C.R., Shimada, T., and Tiee, J.J., “Comparisons Between Hyperspectral Passive and Multispectral Active Sensor Measurements”, *Chemical and Biological Sensing III*, Gardner, P.J., Editor, Proceedings of SPIE Vol. 4722 (2002), pp. 98–109
- [Gross, 1996] Gross, H.N. “An Image Fusion Algorithm for Spatially Enhancing Spectral Mixture Maps”, *Ph.D. Dissertation*, Rochester Institute of Technology, 1996
- [Gross and Schott, 1996] Gross, H.N. and Schott, J.R., “Application of Spatial Resolution Enhancement and Spectral Mixture Analysis to Hyperspectral Images”, *SPIE Optical Science, Engineering, and Instrumentation Annual Meeting*, Vol. 2821 (1996), pp. 30–41
- [Gu, *et al.*, 2000] Gu, D., Gillespie, A.R., Kahle, A.B., and Palluconi, F.D., “Autonomous Atmospheric Compensation (AAC) of High Resolution Hyperspectral Thermal Infrared Remote-Sensing Imagery”, *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 38, No. 6 (2000), pp. 2557–2570
- [Harig, *et al.*, 2002] Harig, R., Matz, G., Rusch, P., “Scanning Infrared Remote Sensing System for Identification, Visualization and Quantification of Airborne Pollutants”, *Instrumentation for Air Pollution and Global Atmospheric Monitoring*,

- Jensen J. O., Spellicy R. L., Editors, Proceedings of SPIE Vol. 4574 (2002), pp. 83–94
- [Hattenberger, *et al.*, 2003] Hattenberger, T., Messinger, D., Brown, S. and Schott, J., “Gaseous Effluent Modeling From Factory Stacks”, *Digital Imaging and Remote Sensing Laboratory Final Report to the Kelex Corporation*, 2003
- [Lisowski and Cook, 1996] Lisowski, J.J. and Cook, C.A., “A SVD Method for Spectral Decomposition and Classification of ARES Data”, *SPIE Optical Science, Engineering, and Instrumentation Annual Meeting*, Vol. 2821 (1996), pp. 14–29
- [Marinelli, *et al.*, 2000] Marinelli, W.J., Gittins, C.M., Gelb, A.H. and Green, B.D., “A Tunable Fabry-Perot Etalon-Based Long-Wavelength Infrared Imaging Spectroradiometer”, *Applied Optics*, Vol. 38 Issue 12, pp. 2594–2604 (2000)
- [Rhody, 2004] Rhody, H., “Sharpening of Hyperspectral Images”, Presentation to the LASS Sponsors, 2004
- [Schott, 1997] Schott, J.R., *Remote Sensing: The Image Chain Approach*, New York: Oxford University Press, 1997
- [Young, 2002] Young, S., “Detection and Quantification of Gases in Industrial-Stack Plumes Using Thermal-Infrared Hyperspectral Imaging”, Aerospace Report No. ATR-2002(8407)-1, 2002

- [Young, *et al.*, 2002] Young, S.J., Johnson, B.R., Hackwell, J.A., “An In-Scene Method for Atmospheric Compensation of Thermal Hyperspectral Data”, *Journal of Geophysical Research*, Vol. 107 Number D24, pp. 14-1 – 14-20 (2002)

Appendix A

Detection Maps From All Cases




Figure A.1: Results from Case 2. The upper row contains the results from Method A for the Freon-114 (left) and NH_3 (right) plumes. The lower row shows the same results using Method B. In the original image, the NH_3 plume is on the left while the Freon-114 plume is on the right.




Figure A.2: Results from the atmospherically compensated Case 2 image. The upper row contains the results from Method A for the Freon-114 (left) and NH_3 (right) plumes. The lower row shows the same results using Method B. In the original image, the NH_3 plume is on the left while the Freon-114 plume is on the right.




Figure A.3: Results from Case 2 using the 95% F-test values for the Freon-114 (left) and NH_3 (right) plumes. Method A was the only technique used. In the original image, the NH_3 plume is on the left while the Freon-114 plume is on the right.




Figure A.4: Results from Case 3. The upper row contains the results from Method A for Freon-114 (left) and tce1122 (right). The lower row shows the same results using Method B. In the original image, both gases are contained in the same plume.




Figure A.5: Results from Case 4. The upper row contains the results from Method A for the Freon-114 (left) and NH_3 (right) plumes. The lower row shows the same results using Method B. In the original image, the NH_3 plume is on the left while the Freon-114 plume is on the right.

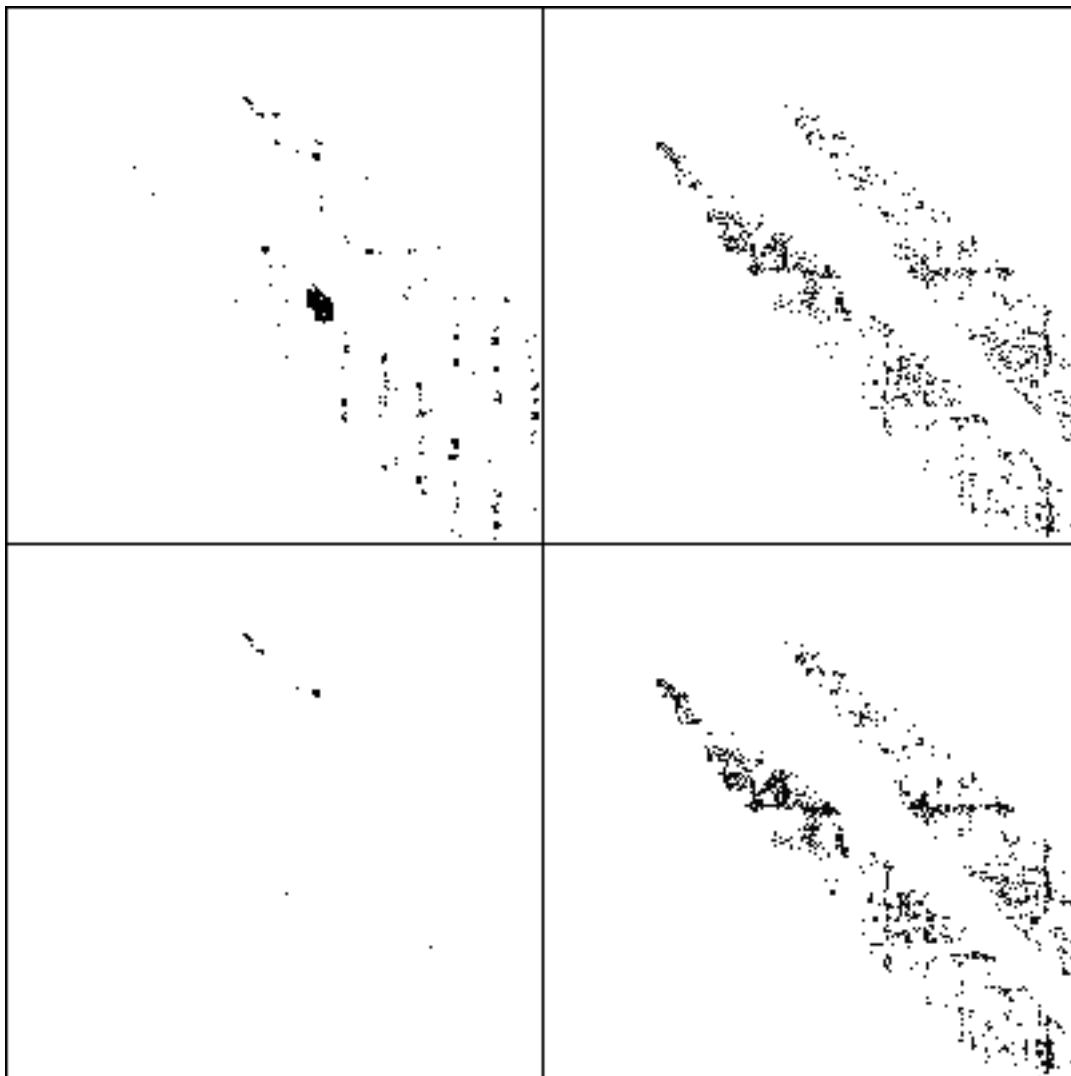


Figure A.6: Results from Case 4b. The upper row contains the results from Method A for the Freon-114 (left) and NH_3 (right) plumes. The lower row shows the same results using Method B. In the original image, the NH_3 plume is on the left while the Freon-114 plume is on the right.




Figure A.7: Results from the atmospherically compensated Case 4b image. The upper row contains the results from Method A for the Freon-114 (left) and NH_3 (right) plumes. The lower row shows the same results using Method B. In the original image, the NH_3 plume is on the left while the Freon-114 plume is on the right.




Figure A.8: Results from Case 5. The upper row contains the results from Method A for NH_3 (left) and tce1122 (right). The lower row shows the same results using Method B. In the original image, both gases are contained in the same plume.




Figure A.9: Results from the atmospherically compensated Case 5 image. The upper row contains the results from Method A for NH_3 (left) and tce1122 (right). The lower row shows the same results using Method B. In the original image, both gases are contained in the same plume.




Figure A.10: Results from Case 6. The upper row contains the results from Method A for CH_4 (left) and dclp12 (right). The lower row shows the same results using Method B. In the original image, the CH_4 plume is on the left while the dclp12 plume is on the right.




Figure A.11: Results from Case 7. The upper row contains the results from Method A for $\text{C}_6\text{H}_5\text{F}$ (left) and phosgene (right). The lower row shows the same results using Method B. In the original image, the phosgene plume (Unknown Gas 1) is on the left while the $\text{C}_6\text{H}_5\text{F}$ plume (Unknown Gas 2) is on the right.

Appendix B

Graphical Results From All Cases



Figure B.1: Method A results for the NH_3 plume in Case 2.

Figure B.2: Method A results for the Freon-114 plume in Case 2.



Figure B.3: Method B results for the NH_3 plume in Case 2.

Figure B.4: Method B results for the Freon-114 plume in Case 2.




Figure B.5: Method A results for the NH_3 plume in Case 2. For this set of results, the F-test was calculated at a probability of 95%.




Figure B.6: Method A results for the Freon-114 plume in Case 2. For this set of results, the F-test was calculated at a probability of 95%.




Figure B.7: Method A results for the NH_3 plume in the atmospherically compensated Case 2 image.




Figure B.8: Method A results for the freon-114 plume in the atmospherically compensated Case 2 image.




Figure B.9: Method B results for the NH_3 plume in the atmospherically compensated Case 2 image.




Figure B.10: Method B results for the freon-114 plume in the atmospherically compensated Case 2 image.




Figure B.11: Method A results for the single plume in Case 3.




Figure B.12: Method B results for the single plume in Case 3.



Figure B.13: Method A results for the NH_3 plume in Case 4.

Figure B.14: Method A results for the Freon-114 plume in Case 4.



Figure B.15: Method B results for the NH_3 plume in Case 4.

Figure B.16: Method B results for the Freon-114 plume in Case 4.



Figure B.17: Method A results for the NH_3 plume in Case 4b.

Figure B.18: Method A results for the Freon-114 plume in Case 4b.



Figure B.19: Method B results for the NH_3 plume in Case 4b.

Figure B.20: Method B results for the Freon-114 plume in Case 4b.




Figure B.21: Method A results for the NH_3 plume in the atmospherically compensated Case 4b image.




Figure B.22: Method A results for the Freon-114 plume in the atmospherically compensated Case 4b image.




Figure B.23: Method B results for the NH_3 plume in the atmospherically compensated Case 4b image.




Figure B.24: Method B results for the Freon-114 plume in the atmospherically compensated Case 4b image.




Figure B.25: Method A results for the single plume in Case 5.




Figure B.26: Method B results for the single plume in Case 5.




Figure B.27: Method A results for the single plume in the atmospherically compensated Case 5 image.




Figure B.28: Method B results for the single plume in the atmospherically compensated Case 5 image.





Figure B.29: Method A results for the CH₄ plume in Case 6.

Figure B.30: Method A results for the dclp12 plume in Case 6.

Figure B.31: Method B results for the CH_4 plume in Case 6.Figure B.32: Method B results for the dclp12 plume in Case 6.




Figure B.33: Method A results for the Unknown Gas 1 plume in Case 7.




Figure B.34: Method A results for the Unknown Gas 2 plume in Case 7.




Figure B.35: Method B results for the Unknown Gas 1 plume in Case 7.




Figure B.36: Method B results for the Unknown Gas 2 plume in Case 7.