A Systems Level Characterization and Tradespace Evaluation of a Simulated Airborne Fourier Transform Infrared Spectrometer for Gas Detection

by

Aaron Weiner

M.S., Wright State University, 2003

A dissertation submitted in partial fulfillment of the

requirements for the degree of Doctor of Philosophy

in the Chester F. Carlson Center for Imaging Science,

College of Science,

Rochester Institute of Technology

July 1, 2010

Signature of Author _____

Accepted by _____

Coordinator, Ph.D. Degree Program Date

CHESTER F. CARLSON CENTER FOR IMAGING SCIENCE

ROCHESTER INSTITUTE OF TECHNOLOGY

ROCHESTER, NEW YORK

CERTIFICATE OF APPROVAL

Ph.D. DEGREE DISSERTATION

The Ph.D. Degree Dissertation of Aaron Weiner
has been examined and approved by the
dissertation committee as satisfactory for the
dissertation required for the
Ph.D. Degree in Imaging Science

_____

Dr. David Messinger, Dissertation Advisor

_____

Dr. Christina Collison

_____

Dr. Carl Salvaggio

_____

Dr. John Schott

_____

Date

DISSERTATION RELEASE PERMISSION

ROCHESTER INSTITUTE OF TECHNOLOGY

CHESTER F. CARLSON CENTER FOR IMAGING SCIENCE

Title of Dissertation:

A Systems Level Characterization and Tradespace Evaluation of a Simulated Airborne Fourier Transform Infrared Spectrometer for Gas Detection

I, Aaron Weiner, 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

A Systems Level Characterization and Tradespace Evaluation of a Simulated Airborne Fourier Transform Infrared Spectrometer for Gas Detection

by

Aaron Weiner

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

Abstract

The remote sensing gas detection problem is one with no straightforward solution. While success has been achieved in detecting and identifying gases released from industrial stacks and other large plumes, the fugituve gas detection problem is far more complex. Fugitive gas represents a far smaller target and may be generated by leaking pipes, vents, or small scale chemical production. The nature of fugitive gas emission is such that one has no foreknowledge of the location, quantity, or transient rate of the targeted effluent which requires one to cover a broad area with high sensitivity. In such a scenario, a mobile airborne platform would be a likely candidate. Further, the spectrometer used for gas detection should be capable of rapid scan rates to prevent spatial and spectral smearing, while maintaining high resolution to aid in species identification. Often, insufficient signal to noise (SNR) prevents spectrometers from delivering useful results under such conditions. While common dispersive element spectrometers (DES) suffer from decreasing SNR with increasing spectral dispersion, Fourier Transform Spectrometers (FTS) generally do not and would seemingly be an ideal choice for such an application.

FTS are ubiquitous in chemical laboratories and in use as ground based spectrometers, but have not become as pervasive in mobile applications. While FTS spectrometers would otherwise be ideal for high resolution rapid scanning in search of gaseous effluents, when conducted via a mobile platform the process of optical interferogram formation to form spectra is corrupted when the input signal is temporally unstable.

This work seeks to explore the tradespace of an airborne Michelson based FTS in terms of modeling and characterizing the performance degradation over a variety of environmental and optical parameters. The major variables modeled and examined include: maximum optical path distance (resolution), scan rate, platform velocity, altitude, atmospheric and background emissivity variability, gas target parameters such as temperature, concentration-pathlength, confuser gas

presence, and optical effects including apodization effects, single and double-sided interferograms, internal mirror positional accuracy errors, and primary mirror jitter effects. It is through an understanding of how each of the aforementioned variables impacts the gas detection performance that one can constrain design parameters in developing and engineering an FTS suitable to the airborne environment.

The instrument model was compared to output from ground-based FTS instruments as well as airborne data taken from the Airborne Hyperspectral Imager (AHI) and found to be in good agreement. Monte Carlo studies were used to map the impact of the performance variables and unique detection algorithms, based on common detection scores, were used to quantify performance degradation. Scene-based scenarios were employed to evaluate performance of a scanning FTS under variable and complex conditions. It was found that despite critical sampling errors and rapidly varying radiance signals, while losing the ability to reproduce a radiometrically accurate spectrum, an FTS offered the unique ability to reproduce spectral evidence of a gas in scenarios where a dispersive element spectrometer (DES) might not.

Acknowledgements

Lots of people helped me…they go here: Placeholder

Dedication

Dedication goes here

Contents

List of Figures

List of Tables

1 Introduction

Gaseous effluent species detection and identification has long been a major goal of the remote sensing community. Several benefits can be derived from knowing the species, and possibly the quantity, of gases in the environment. Specifically, valuable information can be extracted from tracking the sources of environmental pollutants such as "greenhouse" gases, oxides of sulfur and carbon, and other gases associated with industrial facilities. In addition to environmental monitoring, knowledge of the gases produced from an industrial facility may provide information about the kinds and quantities of materials produced at a given site, the efficiency of the various production processes, and the times and rates of activity across the site. This information is useful for those conducting treaty and compliance monitoring, or even commercial espionage, among other intelligence purposes.

Gas detection and identification as a process has long been achieved in the laboratory setting. Fourier Transform Infrared Spectrometers (FTS) provide high resolution, low noise spectra covering the shortwave through far infrared (two to twenty micrometers) in fewer than ten seconds (Griffiths & de Haseth, 2007). In addition, dispersive element spectrometers (DES), such as long focal length monochromaters using highly-ruled diffraction elements, can resolve the finest of spectral details. These instruments can provide nearly certain identification and reliable characterization of any gas. Unfortunately, the technology and methods applied in the laboratory do not readily translate to the outdoor remote sensing problem.

In the laboratory, several steps are taken to isolate the true spectrum from major sources of spectral contamination. The gas sample is isolated in a vessel and the temperature and pressure of the collection conditions are tightly controlled. Further, the background environment in the spectrometer can be easily removed via subtraction and/or inert gas purge. Finally, a controlled source emission is passed through the gas sample and monitored for absorption, or a known source is used as an emission stimulus, providing predictable results (Griffiths & de Haseth, 2007). This collection environment is clearly not possible when the remote sensing problem is moved outdoors where the gas vessel becomes the atmosphere and the background can be extremely varied. Of course, the key then becomes to either control, mitigate the ill effects of, or quantify as many of these parameters as possible.

One proven approach to the remote sensing of gas plumes, which mitigates the effects of many of the parameters that cannot be controlled, is in the detection of gaseous effluent from exhaust stacks of major industrial facilities (Chaffin, Marshall, & Kroutil, 1995). To minimize the effects of the background, one places a staring spectrometer on the ground and points up at the stack to collect through the vapor plume with the cold sky as a background. This step provides good thermal contrast, ensuring the plume will be in emission, while providing a relatively stable background that can be well modeled with minimal effort. In addition, using a static collection system to stare at a virtually constant stream of the gaseous target is nearly an ideal detection scenario. Providing the

near constant presence of the source target allows for substantial integration times, which significantly improves the signal quality. Finally, the likely large amount and high temperature of effluent leaving the stack provides a strong spectral emission from the target gas, further improving the signal quality. Clearly, this collection method is well poised to obtain good results in this scenario.

Of course, when one mounts the spectrometer to a moving platform, such as an aircraft, the scenario becomes increasingly complex, but certainly tractable. One still knows where to stare to obtain a target spectrum and assuming a reasonable pointing accuracy, the major problem then becomes the changing background and atmosphere, vibration, and pointing accuracy that may introduce additional spectral features into the target spectrum further complicating detection and identification. However, many of these problems can be compensated for in post-processing with accurate modeling of the collection environment.

While major strides have been made in the staring imaging spectrometer field, where high resolution tripod mounted imaging spectrometers can be left to autonomously collect and process detection images at video rates, the down-looking airborne problem leaves much to be desired (Chamberland, Villemaire, & Tremblay, 2004). The problem becomes increasingly complex if the collection requirement is not to stare at a fixed source such as an exhaust stack, but rather to locate a source, such as a leak in a pipe in a given region. In the case of the so-called "fugitive" gas emission, not only is the knowledge of the source position lost, but the predictability associated with the rate and thermal contrast of the chemical plume also becomes uncertain. Now the requirements placed on the sensor to address the new collection requirements become that of a "gas finder" and then a "species identifier" – essentially two different sensor systems.

The FTS is particularly well suited to accomplish both of these sensor roles due to its method of forming a spectrum. Adjustments to instrument resolution are only primarily limited by the space available in the optical path and the instrument throughput is significantly higher than traditional DES instruments. Whereas a DES will suffer substantial loss in signal to noise (SNR) with increased spectral dispersion to achieve fine spectral separation, an FTS retains much of the original SNR at high resolution. In terms of scan rate and resolution, a single FTS can be operated in both a mode optimal for finding gas leaks as well as being able to rapidly switch to a "species identifier" mode – a task similarly accomplished by two *different* DES instruments. While the flexibility offered by an FTS is unique among spectrometers, the method of spectrum formation also has drawbacks associated with temporally fluctuating input signals during integration.

1.1 Objectives

This work seeks to explore the tradespace of an airborne Michelson based FTS in terms of modeling and characterizing the performance degradation over a variety of environmental and optical parameters. The major variables modeled and examined include: maximum optical path distance (resolution), scan rate, platform velocity, altitude, atmospheric and background emissivity variability, gas target parameters such as temperature, concentration-pathlength, confuser gas presence, and optical errors including apodization effects, single and double-sided interferograms,

internal mirror positional accuracy, and primary mirror jitter effects. It is through an understanding of how each of the aforementioned variables impacts the gas detection performance that one can constrain design parameters in developing and engineering an FTS suitable to the airborne environment. This effort will be aided by the use of the Digital Imaging and Remote Sensing Image Generation (DIRSIG) model along with its associated Blackadar Gas Plume Generator. Using DIRSIG, complex scenes are developed with multiple background material emissivities and temperatures, variable atmosphere and an accurate plume model steeped in the latest advances in computational gas plume theory. This realistic modeling of the collection environment allows for testing of data for an instrument that is not operational and real data is not readily available.

In condensed format, the objectives of this work include:

- Identify and Map FTS instrument parameters that affect detection performance
- Construct and validate an airborne Michelson FTS-based computer model that employs each of the identified parameters
- Develop detection performance metrics independent of SNR that accurately quantify the impact of each parameter
- Independently evaluate the tradespace among sets of parameters using Monte Carlo random trials to determine likely operational envelopes and thresholds of reliable detection performance
- Construct simulated scenarios in which the operational envelopes are employed to determine the performance of the FTS system in given collection scenarios

Relevant background and theory information to the gaseous plume remote sensing problem will be presented in the subsequent section. The background is then followed by an approach section detailing the performance parameters and tradespace map, the approach and the inner workings of the FTS model, and the construction methods of the complex scene models used in the analysis. This is followed by an experimental section that covers all of the trials and studies used to map the tradespace as well as the design of the scenarios used to evaluate system performance. The next section will cover the results and their analysis in forming the foundation for the FTS performance characterization. Finally, the work and the most meaningful results are summarized in the conclusion section with recommendations for future work in taking the model to higher levels of fidelity.

2 Background

The background is written from the stand point of introducing concepts that are critical in the understanding of the enabling processes that allow one to conduct remote sensing in the hopes of detecting and identifying target gases in the Long Wave Infrared (LWIR). These topics cover: gaseous and solid phase spectroscopy, atmospheric influences on spectra collected in the LWIR, a mathematical model for sensor-reaching radiance, a primer on Fourier Transform Infrared Spectrometers (FTS), and a summary of prior work in the field.

2.1 Gaseous Spectroscopy

The foundational underpinning of this work relies on the understanding of the spectroscopic phenomena of gaseous chemical species that allows for the exploitation of collected spectra. Each atom and molecule can be exploited via its interaction with light to reveal its unique identity. To further understand how this is possible, one must understand, among other things, the nature of the electronic environment of each atom and molecule.

While the scientific discipline of spectroscopy encompasses multiple transition modes and sensing modalities, this discussion will be confined to valence electron transition associated with vibrational and rotational molecular motion. These transitions commonly occur in the infrared and can be sensed by FTS and myriad dispersive spectrometers.

The descriptive model of the electronic environment in a molecule is comprised of various allowed quantum states that constitute a given molecular orbital. It is precisely these allowed quantum states that give rise to the discrete nature of spectra associated with gases. Molecular orbitals are formed when quantum states of two atoms overlap defining regions of space where an electron is likely to be found. The energy associated with each allowed quantum state is affected by such factors as distance from the nuclei and proximity of neighboring electrons. Interactions between each electron and the protons found in both the “home” atom nucleus and the new neighboring atom’s nucleus affect the motion, and therefore energy, of the electron through forces of electrostatic attraction. Repulsive forces from nearby electrons and nuclear shielding from electrons in lower quantum states placing them closer to the nucleus counteract these forces of attraction. Clearly, one can see that fundamental quantities such as the number of protons in the nucleus and number of electrons associated with an atom have great influence over the electronic environment and the resulting energy states. The distinctiveness of these quantities in each atom and molecule is the primary reason why these entities can be uniquely indentified via their spectrum.

While the electronic distribution of a molecule is an important consideration, the real story behind the generation of the observed spectral features is associated with the vibrations and rotations that

a molecule undergoes. Starting with a diatomic example, one may envision the electrostatic attraction and ensuing overlap of quantum states as a bond between the two atoms that keeps them in proximity of one another. This bond is often modeled as a spring with an associated restoring force undergoing simple harmonic motion. The potential energy of the restoring force models the energy bounds of the allowed vibrational quantum states (Laidler, Meiser, & Santuary, 2003). While the potential energy of a spring, as a function of displacement, can be modeled as a parabola, that of a molecule needs further modification. Assuming the molecule has a temperature it will vibrate. While undergoing vibration, the range of movement will be influenced from the forces of nuclear repulsion and electrostatic attraction resulting in anharmonic deviations from the parabolic approximation of the potential energy, particularly at high vibrational energy levels. As the two nuclei approach each other, the force of repulsion generated by the protons grows asymptotically with decreasing internuclear separation. Conversely, forces of attraction decrease gradually as the nuclei separate until the nuclei are so far apart that the momentum of the nucleus is sufficient to overcome the attraction and the molecule dissociates into two atomic fragments. At the point of dissociation, no further vibrational quantum states can exist, which serves to provide an envelope of energy associated with the vibration.

This envelope of energy is modeled by the Morse potential energy diagram and is shown below in Figure 2.1. The figure shows several possible transitions in both absorption and emission. Note how absorption events are initially limited to the ground vibrational state of the ground electronic state in this system. The fact that most absorption originates from the ground electronic state limits the locations of line positions that one would observe for a system in absorption. Whereas in emission, the process may originate from any level where an excited population may exist, in addition to internal relaxation processes and quenching that result in a reduction of energy before the photon is released with the remaining energy. This phenomenon generally results in several more line positions evident in an emission spectrum as compared to an absorption spectrum for a given molecule. Further, this concept gives rise to fluorescent and phosphorescent modes where short wavelength, high energy photons are absorbed, but some time later, after relaxation events or system crossings take place, a longer wavelength, lower energy photon emerges.

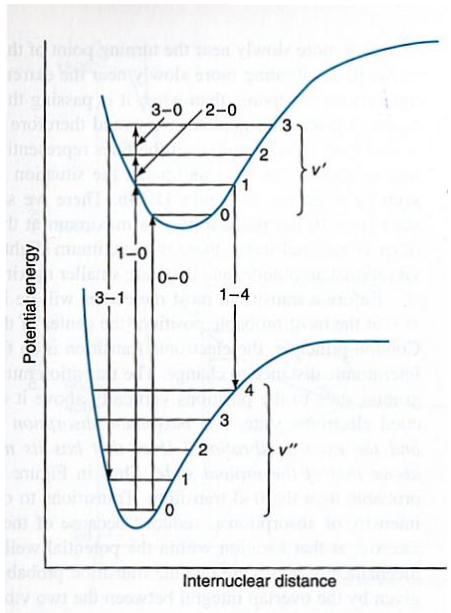

FIGURE 2.1 MORSE VIBRATIONAL POTENTIAL ENERGY DIAGRAM FROM LAIDLER

The vibrational quantum states associated with a given form of vibration begin relatively equally spaced, as shown in the Figure 2.1, and then become infinitesimally close as the dissociation level is approached. This narrowing of energy level spacing approaches a continuum representing the lack of constraining forces on the two nuclei as they essentially become separate, free entities (Laidler, Meiser, & Santuary, 2003). For practical purposes of spectra encountered in passive remote sensing, generally transitions between the ground and first and second excited vibrational states will give rise to the observable spectral features. The energy spacing for these transitions is nearly identical and thus overtone (or harmonic) transitions, those transitions that are multiples of energy from the ground to first excited state, can quickly be identified in the spectrum at line positions that correspond to nearly twice the energy of the transition between the ground and first excited state – the fundamental transition.

While the diatomic vibrational model serves as an easily grasped conceptual understanding of the origins of vibrational spectra, it certainly does not tell the whole story. Polyatomic molecules, by virtue of having more than two atoms, now have additional degrees of freedom that translate to multiple modes of vibration. Carbon dioxide is very illustrative of this point.

The total vibrational degrees of freedom (DoF) of a molecule can be found by multiplying the number of atoms in the molecule by three and subtracting five for a linear molecule and six for a bent molecule (Laidler, Meiser, & Santuary, 2003). In the case of carbon dioxide, a linear molecule, there are four DoF, whereas water, a bent molecule, would have three. These DoF can be visualized by placing a three-dimensional axis at the center of the molecule. We can then set the molecule in motion and see that the oxygen atoms can each vibrate in an equidistant manner from the central carbon atom, or alternate in an asymmetric fashion. Additionally, the oxygen atoms can bend within and out of a plane superimposed through the central axis. To see how these DoF translate to

observed spectral features, one must understand the energy associated with each vibration as well as the mechanism of coupling between electromagnetic radiation and the molecule.

From an energy standpoint, it turns out that comparably little energy is required to make a molecule bend – for carbon dioxide, this energy is on the order of 667 cm^{-1} (15 µm, 82.7 meV). It should follow that one should observe two spectral features for carbon dioxide for bending, one for each bending DoF, however only one is found. This is due to the fact that it takes the molecule the same amount of energy to bend regardless of what "direction" the bending takes place. The planes we superimpose on the molecule are simply a matter of recordkeeping and do not imply differences in energy. Additionally, the asymmetric stretching noted above will have a much higher energy associated with it due to the energy required to cause asymmetric displacement of the atoms in the bond. For carbon dioxide, this energy is about 2350 cm^{-1} (4.25 µm, 291.4 meV) and a spectral feature is in fact observed at this corresponding line position. Interestingly, the symmetric stretch, which corresponds to an energy of 1340 cm^{-1} (7.46 µm, 166.1 meV) has no observable spectral feature. This spectral feature is not evident because the symmetric mode of vibration in carbon dioxide does not allow for the coupling of electromagnetic radiation with the molecule.

In order for a molecule to interact with light to manifest vibrational-rotational spectral features, the molecule must possess a rhythmically transient net dipole (Laidler, Meiser, & Santuary, 2003). A dipole is formed when a separation of charge has occurred. In this case the highly electronegative oxygen and less electronegative carbon atoms would form a dipole moment between each other. In the case of an asymmetric stretch, for the side of the molecule with the farther displaced oxygen atom, this side will have a larger magnitude dipole than the other, minimally displaced oxygen. The result is a net dipole that oscillates as the individual dipoles change in magnitude through each vibration. From electrodynamics, it is known that an oscillating charge creates a field. It is this field that allows coupling with electromagnetic radiation (which itself has alternating electric and magnetic fields) with the appropriate energy to cause an electron transition between two vibrational states. It is now easy to see that the mode of symmetric stretching for carbon dioxide always poses two equal dipoles against each other creating no net dipole. This mode of vibration cannot couple with light and is therefore deemed infrared-inactive. In the bending vibrational mode, the bending of the molecule removes the direct opposition of the dipoles and allows for the summation of both creating a net dipole that weakly oscillates through the bending motion.

To further complicate matters, each vibrational mode will have its own vibrational spectral transition with associated harmonic transitions, in addition to the combination bands that arise from combined transitions between fundamental and harmonic transitions of two different vibrational modes. The identification of combination bands is not straightforward and is well beyond the scope of this discussion.

What has essentially been accomplished to this point would allow one to identify the center of the band position of a spectral feature that would arise from a vibrational transition. As was demonstrated with carbon dioxide, the amount of energy required to stimulate a vibrational mode of a molecule is on the order of tenths of an eV and most molecules in ambient atmospheric

conditions will experience sufficient thermal energy to excite the first few of these transitions. What have not yet been discussed are the rotational transitions that often accompany each observed vibrational transition in the infrared (IR).

Rotational transitions arise from the rotational DoF associated with a molecule, and if a three dimensional axis is superimposed, it can be seen that there are three rotational DoF. Of course, as was seen in the argument relating the multiple bending DoF in carbon dioxide to the singular vibrational transition, the same can be said for rotational DoF and the energy of rotation. Again, our recordkeeping does not constrain a molecule's rotation in three-space and it stands to reason that the same amount of energy is required to accomplish a similar rotation through any plane of an axis.

Just as very little energy is required to bend a molecule; even less is required to rotate a molecule. Pure rotational transitions occur in the microwave region of the electromagnetic spectrum where required energies are on the order of hundredths of an eV. In fact, in most ambient atmospheric conditions, the thermal energy alone is enough to elevate the most populous rotational energy level from the ground state (as is the case with vibrational levels) to the third or fourth rotational level (or higher). While rotational transitions can be selectively excited, they generally accompany observed vibrational transitions associated with IR energies. This creates a joint vibrational-rotational, or ro-vibrational, transition. This is of importance to the remote sensing community as this type of transition is responsible for the spectral band shape that is observed for gases.

Vibrational-rotational transitions are modeled by what are known as allowed transitions, or selection rules. These rules were derived from the limits of symmetry in the molecule with consideration given for favorable overlap of the wavefunctions describing the energy levels of the system. The rules dictate that the only allowed transitions are between rotational states that are separated by one energy level. If, for example, the originating level is from rotational level five in the vibrational ground state and it transfers to a level in the excited vibrational state, it must transition to an excited state in rotational level four or six. Thus, if the transition is from five to four, the change in levels is one less and is a P-type transition. If the transition is from five to six, the change in levels is one more and is an R-type transition. If the rotational energy levels were all equally spaced, as is nearly true with the first few vibrational levels, these rules would be of no consequence, but the necessity for this convention arose due to the non-linear increase in energy in each successive rotational level. It can be seen that if each transition is between two differing rotational levels of what is otherwise two equally spaced vibrational levels, one will see spectral line positions occur on either side of what would be the line position of a purely vibrational transition. The P transitions are less energetic than the R transitions and form a manifold, or branch, of spectral lines at lower energy line positions from the central vibrational line position, whereas the R branch forms on the opposite side. Figure 2.2 is a modified Morse potential energy diagram showing the vibrational energy envelope limiting the vibrational energy as a function of internuclear separation, with an inset showing an exploded view of notional rotational transitions of the P and R branches from the ground to first excited vibrational states. Note the non-linear increase in rotational energy level spacing with each successive rotational level.

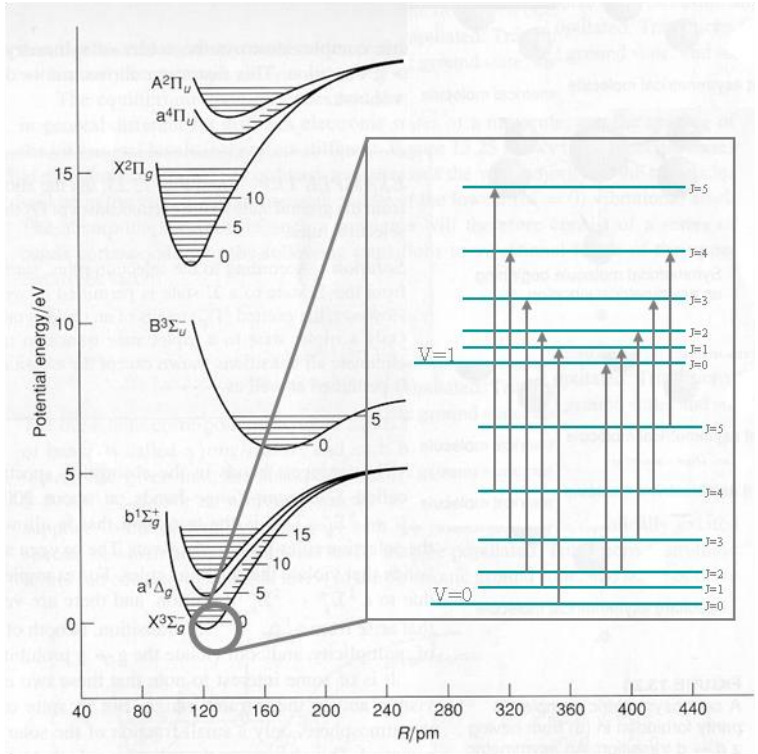

FIGURE 2.2 MORSE POTENTIAL ENERGY DIAGRAM WITH ROTATIONAL TRANSITION INSET (LAIDLER, MEISER, & SANTUARY, 2003)

The transition that has not yet been addressed, that between the same rotational energy levels between the ground and excited vibrational levels, is known as a Q branch transition and is considered forbidden. It is unfortunate that the name "forbidden" was chosen for this transition type – it is not really forbidden at all, it is just unlikely given a molecule's symmetry. All that is required for this transition to occur is an accompanying change in orbital angular momentum. This change will depend on the molecular geometry and can be seen in various modes of vibration. In fact, in a simple IR spectrum of a single gas, it is not uncommon to see vibrational-rotational bands where some bands have Q branch spectral features and other bands from differing modes that do not (as is the case with carbon dioxide).

The band shapes correspond to the population of molecules that undergo a given transition. The shape is humped because the majority of the population resides well above the ground state due to the ambient thermal energy. For the R branch, since each successive transition requires additional energy over the previous transition, the available population quickly curtails and the transition series comes to an abrupt end. Similarly, the P branch transitions begin with a hump and quickly curtail as well as each successive transition becomes smaller and smaller and eventually becomes self-limiting. The Q branch appears as more of a spike, or Lorentzian-like line shape, as these transitions are tightly clustered due to the approximately equal spaced energy transitions and would occur near the purely vibrational transition line position. Figure 2.3, from the NIST Chemistry Web Book, shows a LWIR spectrum of dichloromethane with nearly distinct P, Q and R branches. Note the wide-band character of the P and R branches and the Lorentzian character of the Q branch.

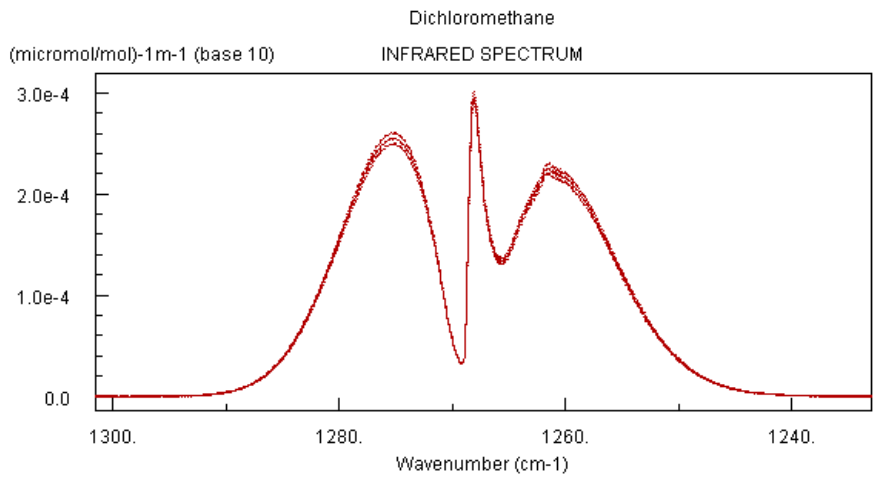

FIGURE 2.3 LWIR SPECTRUM OF DICHLOROMETHANE WITH P, Q, AND R BRANCHES, RES. 0.125 CM^{-1}

Minor changes to the band shape will occur as the thermal energy available to the gas changes. A lower temperature will manifest itself as a contraction of the bands towards the central vibrational transition, with a slight bias towards the P branch. Conversely, heating the gas will promote the population to higher rotational levels, broadening the bands and lowering the intensity of each individual line as the population becomes spread thin.

Factors that affect the position of each successive rotational line include centrifugal distortion and the rotational inertia of the molecule. The effects of centrifugal distortion become increasingly important at high rotational energies. The faster the molecule spins, the greater affect "centrifugal force" will have in causing the internuclear separation to increase beyond that associated with the given vibrational energy for a given vibrational mode. When the internuclear separation increases, there is a resultant increase in the rotational inertia causing an overall decrease in spacing between adjacent rotational lines. Additionally, the reduced mass, which is a component of the rotational inertia, is essentially a ratio of the masses of the atoms in the molecule. The reduced mass influences the energy required to increase the rotation of a molecule and manifests itself as the primary driver of separation between rotational lines. The rotational sensitivity to mass is so great that if one had sufficient spectral resolution, the line positions between two isotopes of one atom in a molecule (with a reasonable contrast in mass) could be resolved, even though the mass difference is only that of one or two neutrons (1.675×10^{-24} g/neutron). This phenomenon manifests itself as rotational peak "doublets", or pairs, with intensities that correspond to the abundance of the isotopic species in the sample.

It is generally the case that in remote sensing of gaseous targets, not only is the spectral resolution required to resolve individual rotational lines in a ro-vibrational band unnecessary for gas identification, it is unreasonable to try to achieve such resolution outside of the laboratory environment. Collected spectra in the LWIR generally show detail at the band level where one may be able to resolve the P, R, and Q branches (if present) of a band as distinct entities, perhaps with some indication of the rotational line detail underlying the band in each branch. The positions of a

few bands representing a few different vibrational modes are generally sufficient to identify a gas. Furthermore, little can be obtained about the rotational, or vibrational, temperature of the gas from the band shape at this resolution as the data are only suited to the coarsest of estimates. The figure below illustrates many of the points from this section including the distinctness of the P and R branches, the appearance of rotational lines in a ro-vibrational band that are partially discernable, and isotopic splitting of rotational lines from the two isotopes of chlorine.

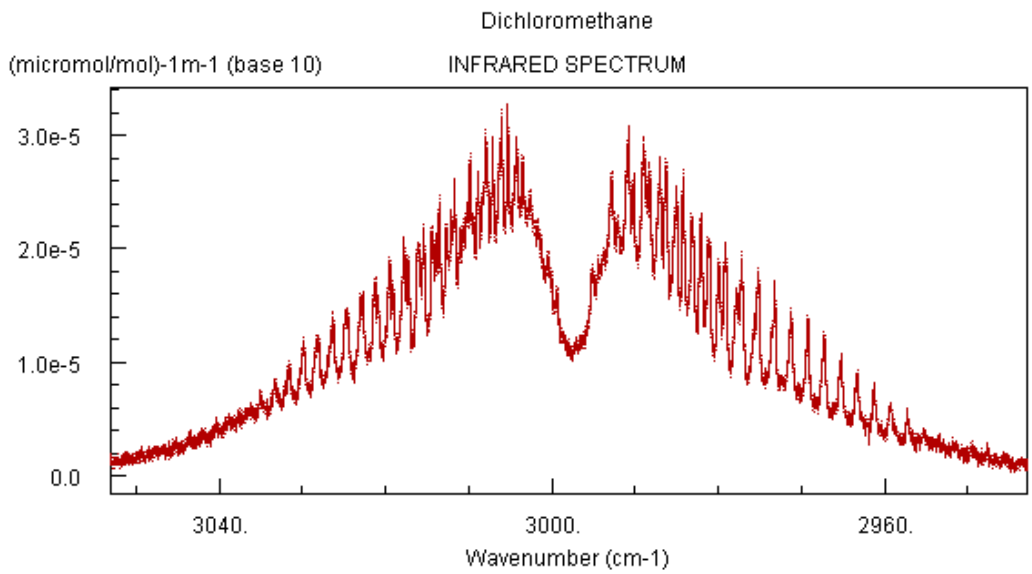

FIGURE 2.4 INFRARED SPECTRUM OF DICHLOROMETHANE, RES. 0.125 CM^{-1}

Note that both of the figures from this section showing the spectrum of dichloromethane illustrate the point made where the particular mode of vibration either allows for, or prevents the Q branch transitions from occurring. Figure 2.4 indicates a stretching mode, while Figure 2.3 represents a type of bend known as a deformation, which changes the molecular symmetry allowing for a change in orbital angular momentum.

2.2 Solid Phase Spectroscopy

While the purpose of this work is to identify chemical species in the gas phase, a major portion of the observed spectrum will be due to solid phase spectral phenomena. This portion of a spectrum generally comprises the background (minus atmosphere) of whatever the down-looking sensor "sees" behind/through the target gas.

Atoms and molecules found in the condensed state behave quite differently than when essentially unconstrained in the vapor phase. Solids are generally represented by a lattice network of atoms in close formation of one type, or a combination of regular, repeating orders (Huheey, Keiter, & Keiter, 1993). Atoms constrained to this state continue to vibrate, but the electronic environment of each is substantially perturbed and influenced by that of the surrounding atoms, which directly affects the energy of the system as a whole. In addition to photon interaction, phonons are introduced, which represents energy transferred to an entity through the vibrational motion of the bulk lattice

(Patterson & Bailey, 2007). All of these factors combine to essentially smear out what was once the discrete spectral nature of entities in the gas phase. In fact, the emission spectrum of a solid is generally well modeled by a continuous function known as the Planck Black Body Equation, equation 2.1, with units of spectral radiance

$$L_\lambda = \frac{2hc^2}{\lambda^5\left(e^{\frac{hc}{\lambda kT}}-1\right)} \quad \left[\frac{W}{m^2\ sr\ \mu m}\right] \tag{2.1}$$

where the constants and parameters include Planck's constant (*h*) [J s], the speed of light (*c*) [m/s], Boltzmann's constant (*k*) [J/K], the temperature of the radiator (*T*) [K], and the given wavelength of emission (*λ*) [micron]. It is clear that the equation is dependent on the Boltzmann population of the system at a given temperature. This allows one to model the limits of the energy available to each transition of a given wavelength, resulting in the characteristic short wavelength-side drop-off associated with black body radiation.

The construct of a black body is such that an optically opaque, reflectionless physical environment exists in which an emitted photon is absorbed and re-emitted multiple times before it exits the system. As such, the black body model is well suited to describe solids, but also models the general emission from liquids and even extremely dense gases such as the sun.

While the task of background removal would be straightforward if all one needed were the temperature and a Planck black body model, real backgrounds are not so cooperative. Virtually all background materials need to be modulated by a term known as emissivity to bring the predicted black body radiance in line with observed measurements. When energy in the form of a photon impinges on a solid, there are multiple possibilities as to what may happen. Depending on the material properties of the solid, the energy may be redirected in the form of reflection, transmitted through the solid, or absorbed by the solid. Conservation of energy suggests that the sum of the amount of energy reflected, transmitted, and absorbed should equal the total amount of energy incident on the solid. Further, Kirchoff's law states that when it comes to emission, assuming the solid is in thermal equilibrium, a true black body will re-emit whatever is absorbed. This factor, which relates an amount of expected emission from the absorption, is known as emissivity. As can be seen from this example, even if one assumes a dense, optically opaque solid, some amount of energy will be lost to reflection, lowering the amount that can be absorbed, and thus directly lowering the emissivity. So to model a real solid, one must modulate the Planck predicted radiance by the emissivity, a value ranging from zero to one, resulting in what is known as a gray body.

To further complicate this model, emissivity, just as is true with absorption and reflection, is a wavelength dependent property. This fact can substantially change the model of spectral character of a solid showing, in some cases, extreme deviation from the Planck-predicted radiance. Regions in the infrared where there can be extreme deviation from the Planck model are in what are known as Restrahlen bands. Restrahlen bands manifest themselves as major portions of black-body-modeled radiance that are missing due to intense photon absorption from the resonance of the previously mentioned phonons. This is a phenomenon unique to the infrared region and occurs where the

frequencies associated with optical phonons – the vibrational harmonics of the solid lattice, are in resonance and causes intense absorption. Restrahlen features, much like the ro-vibrational bands of gases, are often telltale indicators of the identity of component elements in a solid (Elachi & van Zyl, 2006). Figure 2.5 demonstrates the deviation from the Planck predicted radiance in the LWIR for concrete at 300K. The smooth curve in the chart is the Planck predicted blackbody radiance while the red trace immediately beneath it is the same 300K blackbody curve modulated by the emissivity spectrum of concrete, which is shown as the green trace that corresponds to the right vertical axis.

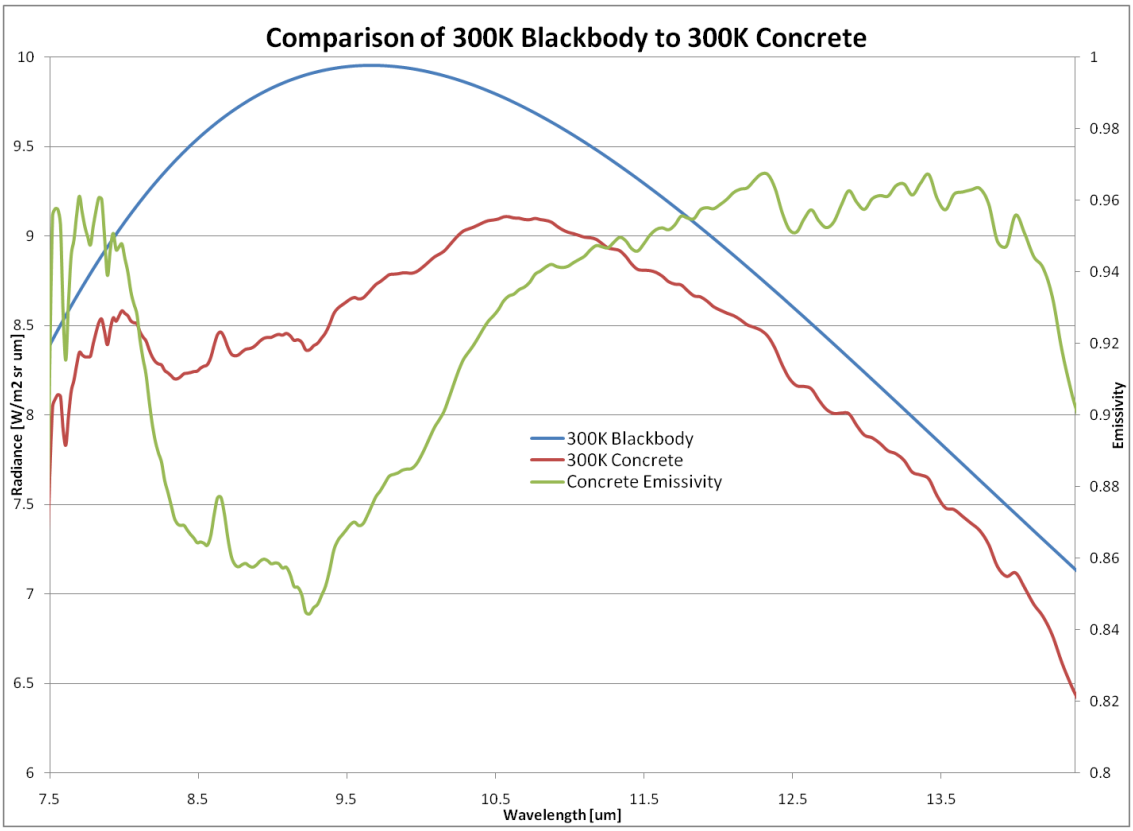

FIGURE 2.5 300K BLACKBODY COMPARISON TO 300K CONCRETE WITH EMISSIVITY OVERLAY

Clearly, background removal or modeling, which is a seemingly straightforward task, can now be seen to be exceedingly difficult depending on the composition of the solid. Fortunately, the discrete spectral features associated with the target gas one may be trying to identify in the collected spectrum are generally finer than those of the solids in the background. Thus, the gas-associated spectral features may appear to "ride on top" of the more slowly spectrally varying background features, still allowing for isolation or identification.

2.3 Atmospheric Effects in the LWIR

The atmosphere is primarily composed of nitrogen and oxygen, leaving approximately one percent to an assortment of other constituents. On the surface, one may observe that nitrogen and oxygen,

being homonuclear diatomics, possess no rhythmically transient net dipole and therefore, will be IR-inactive. Assuming for the moment that this premise is true, we will focus on the other one percent of the atmosphere that will have substantial influence on the LWIR remote sensing problem.

The following table, taken from the Handbook of Chemistry and Physics, illustrates a typical composition of the atmosphere at sea level.

Species	Fraction	Density (n/cm^{-3})
N_2	78.084%	1.99×10^{19}
O_2	20.946%	5.33×10^{18}
Ar	0.934%	2.38×10^{17}
CO_2	0.033%	8.4×10^{15}
CH_4	1.7×10^{-40}%	4.33×10^{13}
N_2O	3.2×10^{-50}%	8.15×10^{12}
CO	1.5×10^{-50}%	3.82×10^{12}

CRC Handbook of Chemistry and Physics, 73rd ed., CRC Publishing Co.,1992.

TABLE 2.1 ATMOSPHERIC COMPOSITION AT SEA LEVEL (CHEMICAL RUBBER COMPANY, 1992)

Note that the noble gas argon is shown as taking up nearly the entire bulk of the remaining one percent of atmospheric constituents, yet argon has no vibrational or rotational modes and is only subject to electronic transitions. Thus, argon is not of concern for the purposes of this work. This leaves an exceedingly small amount of gases left that are responsible for the overwhelming atmospheric absorption features found throughout the infrared portion of the electromagnetic spectrum.

It is a combination of three gases that are responsible for the bulk of these atmospheric absorption features, particularly in the LWIR, yet some were not included in the table above. One of these gases that was notably left off of the composition table shown above is water. Water is extremely variable throughout the atmosphere and is highly dependent on geographic location and time of day, so it was not appropriate to include it in the list above. An additional gas left off of the table is ozone. This makes sense in that ozone is primarily a stratospheric gas and is generally found in negligible concentrations near the surface, outside of dense urban areas. Ozone concentrations are highly correlated to altitude and this impact to airborne remote sensing will be explored shortly. Finally, carbon dioxide, as shown in the table, is the prominent IR-active “mainstay” gas throughout the atmosphere and a primary atmospheric absorption culprit.

Fortunately for the sake of remote sensing, the ro-vibrational features of the aforementioned gases do not entirely mask transmission through the atmosphere. There are three primary windows in the IR that encompass the shortwave, midwave, and LWIR infrared regions. For the purposes of this

work, a subset of the LWIR window from approximately 8-12.5 microns (800-1250 cm^{-1}) will be used. This window is flanked by absorptions from the bending mode of water near 8 microns, accompanied by the deformation mode of methane and stretching harmonic overtone of N_2O. At approximately 9.8 microns, ozone punches a hole in the window with a strong dual stretching mode, P and R branch band absorption feature. Finally, the bending mode of carbon dioxide fills in the long wavelength side of the window with an intense absorption. The LWIR window coincidently covers the spectral range where a majority of the gases of interest to remote sensing have significant ro-vibrational features, many of which occur coincidentally with ozone. These gases will be reviewed in detail in the next chapter.

Radiative transport in this work was accomplished via an atmospheric model called MODTRAN. MODTRAN (Moderate Resolution Atmospheric Transmission) was developed by the Air Force Geophysical Laboratory at Hanscom Air Force Base twenty years ago as an order of magnitude resolution improvement to the then current LOWTRAN 7 code (Low Resolution Atmospheric Transmission) (Berk, Bernstein, & Robertson, MODTRAN: A Moderate Resolution Model for LOWTRAN 7, 1989). MODTRAN generates transmission values by employing pre-determined column densities of various component atmospheric gases, such as the 1976 US standard atmosphere, and sums their absorption across the spectral bands of interest. These terms are then applied to the path the radiance is to travel in both height and angle through the atmosphere. Finally, scattering is considered based on visibility/aerosol loading and environment selected (urban/rural) accounting for radiance scattered in and out of the signal path through multiple scattering events. These absorption and scattering events are built into layers that keep track of the radiance attenuation through the entire path length. Finally, radiance produced by the sun based on time of day and date is considered as well as adverse weather conditions the user may select such as rain, ice, clouds, or differing amounts of component gases and aerosols. These radiance values are then assembled, sampled by the instrument resolution provided, and output into large tables detailing the individual contributions of the various effects. The careful consideration and improvements that have gone into the construction of MODTRAN over the last several decades of its existence have made it the industry standard in moderate resolution atmospheric modeling for remote sensing use.

Figure 2.6 illustrates the substantial transmission losses through the atmosphere in the LWIR due to the named atmospheric gases above at a vertical path length of 1, 5, and 20 kilometers as generated by the MODTRAN model. This data was generated using MODTRAN v4r3, for a mid-latitude summer geographic region, at a temperature of 278K, with 23 km visibility in a rural scattering environment, at 0.05 µm (~5 cm^{-1} at 10 µm) resolution.

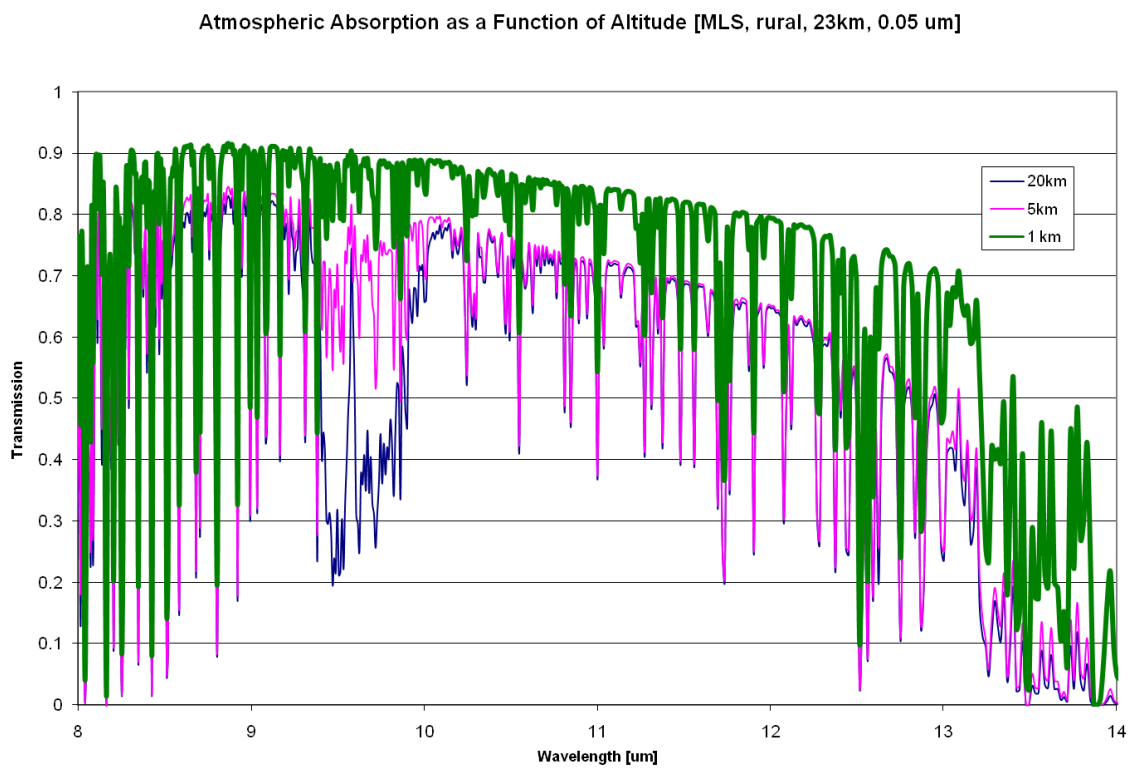

FIGURE 2.6 ATMOSPHERIC ABSORPTION PROFILE BY ALTITUDE IN THE LWIR

Several important points can be made from Figure 2.6. Clearly, individual rotational lines, primarily due to water, can be seen throughout the spectrum. This range of rotational lines, which extends well through the far infrared to the microwave is known as the water rotational continuum (Wozniak & Dera, 2007). Rotational lines for water have a large rotational constant (energy spacing) due to the decreased reduced mass and therefore, decreased rotational inertia, thus it is not unreasonable to be able to resolve individual rotational lines of water. Conversely, carbon dioxide and ozone have comparatively smaller rotational constants due to a much larger rotational inertia resulting in closely spaced rotational lines that form the more recognizable band shapes seen in earlier example spectra.

The most important take away from Figure 2.6 is the difference in absorption with altitude. When comparing the two lower traces correlating to the five and twenty kilometer altitudes, there is minimal differences in nearly all absorption features through the two path lengths with the major exception of the ozone absorption feature. This demonstrates that firstly, essentially all of the absorption due to carbon dioxide and water vapor will occur in the first five kilometers of path length above the ground and second, transmission through ozone is improved by nearly a factor of three when dropping down to 5 km from 20 km. The final trace at a 1 km vertical path shows that ozone is virtually removed from the spectrum and that there is an appreciable improvement in transmission with an expected decrease in the total column concentration of the remaining gases, particularly for water. A final point best illustrated by the left to right slope of the two lower traces

shows that both Rayleigh (molecular) scattering, which is most prominent among shorter wavelengths in the visible, and aerosol scattering have minimal impact to the transmission in the LWIR. Should scattering have had a greater impact, a slope of reduced transmission proportional to λ^{-4} would have been more evident.

As many gases have prominent features in the 9-10 micron region of the LWIR, and due to the significant sensitivity and variability of ozone absorption with altitude, it seems prudent to restrict the operational altitude of a given aerial collection vehicle to approximately one, and at most, five kilometers for the purposes of this study.

As previously demonstrated, the majority of absorption due to water vapor in the atmosphere will occur in the first 5 kilometers of altitude however, the extent of that absorption is dependent on current humidity conditions. Figure 2.7 shows an example of a mid-latitude summer (MLS) standard atmosphere model from MODTRAN compared to a tropical model at 5 km altitude. Note the increased absorption and overall reduced transmission due to nearly double the total column water vapor content. Both of these atmospheric models are employed in this work.

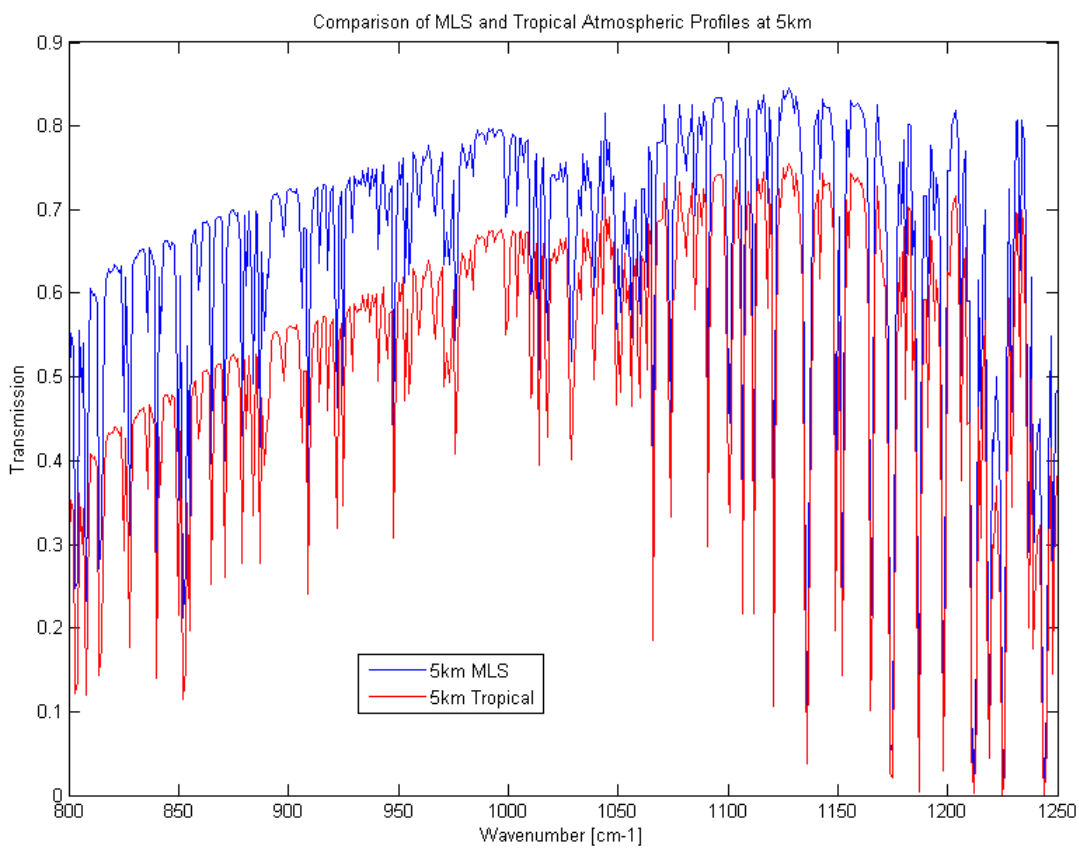

FIGURE 2.7 COMPARISON OF TROPICAL AND MLS ATMOSPHERIC MODELS AT 5KM

Of course, in addition to a cursory analysis of the atmospheric transmission spectrum, understanding the complete impact the atmosphere will have on the target signal requires an

analysis of both the up and down-welled radiance that accompanies the target radiance, among other sources. This analysis can be accomplished via radiometric modeling.

2.4 Mathematically Modeling the Sensor Reaching Spectral Radiance of a Gas

Developing a model to determine the sensor-reaching radiance for the gas detection problem is not a straightforward task and requires multiple assumptions. In the LWIR, the main drivers of signal in a collected spectrum will come from the target gas itself (assuming it is present in the scene), the background, and the atmosphere. Secondary effects from reflections and other sources of stray radiation can be shown to have minimal impact to the overall collected radiance for natural materials (Tonooka, 2001). The development of this sensor reaching radiance model illustrates the foundation of the radiative transport methods implemented in the FTS simulation code used in this work.

2.4.1 Model Assumptions and Mathematical Development

Beginning with the background, which can be modeled as a modified black body radiator as previously established, a temperature must be set to govern the strength of the emission. This temperature can be extracted from brightness or apparent temperature studies of the scene, through a priori knowledge of the conditions in the scene, or through an educated guess (Boonmee, 2007). For the sake of simplicity, the latter method will be used. Next, the emissivity at a given wavelength must be determined. Spectral emissivity estimation can involve methods that are quite complex, and again, for the purposes of simplicity, it can be shown that in the absence of a priori emissivity knowledge, a flat estimate of 0.9 across the LWIR spectral range is a reasonable assumption, particularly from nadir observations (Tonooka, 2001). For this work, actual spectral collects of materials such as concrete, asphalt, steel, and sand taken from Trona, California by members of the Digital Imaging and Remote Sensing (DIRS) group at the Center for Imaging Science (CIS) at the Rochester Institute of Technology (RIT) will be used as background sources. These collects were analyzed by the group to produce emissivity curves in the LWIR and can be imported directly to any model requiring no emissivity or temperature estimates. Finally, in a direct line-of-sight (LOS) to the sensor, the only other impact to the background radiance is the contribution from the atmosphere on a per wavelength basis, which is aptly modeled by the MODTRAN software.

With the background firmly in hand, atmospheric path effects need to be considered. With a minimally reflective surface in the LWIR (modeled at an emissivity of 0.9 for example) there is little worry of reflected radiance as a major source of energy for natural materials. Should the scene contain metals that are highly reflective, there is the possibility of increased spectral contributions from the atmospheric spectrum. This contribution often manifests itself as reduced atmospheric absorption with the possible inclusion of an ozone band. As previously mentioned, strengthening of the ozone band often results in an overlap of key gas spectral features in the LWIR and can impact detection performance, or further, result in false alarm detects. Figure 2.8 illustrates the effects of

up and down-welled radiance derived from MODTRAN using a 1 km tropical atmosphere and a 0.9 constant emissivity gray body surface.

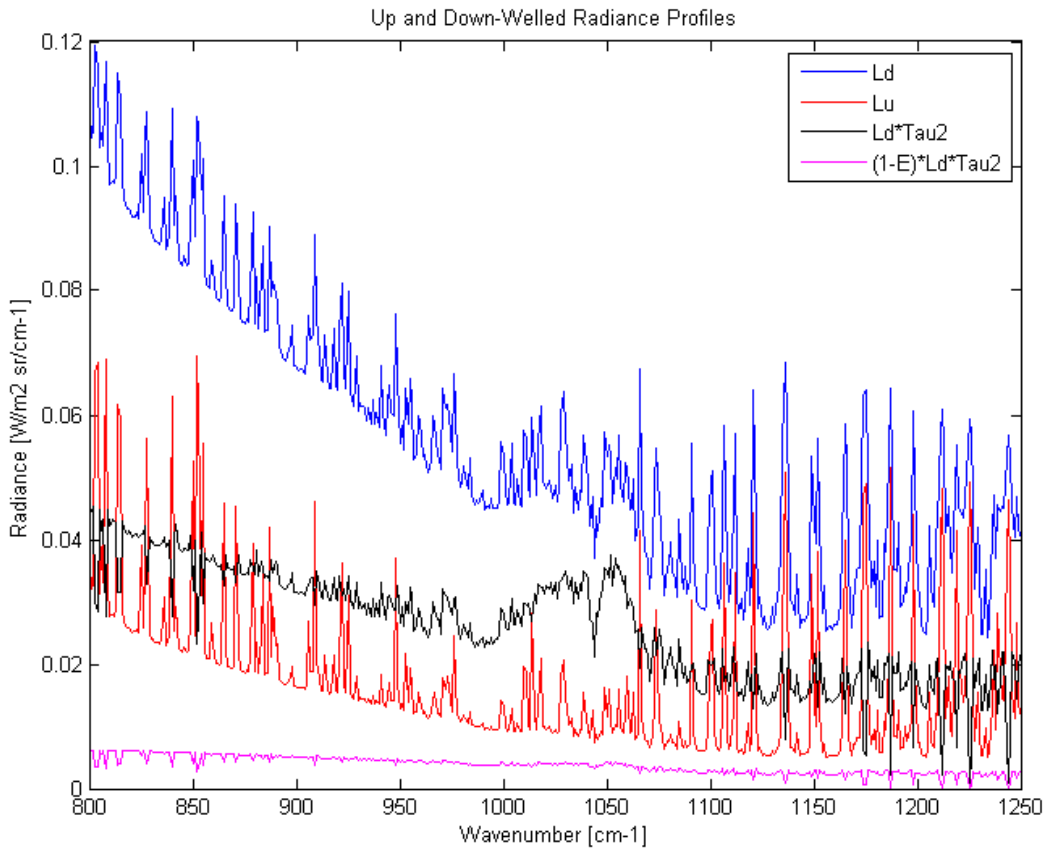

FIGURE 2.8 MODELED UP AND DOWN-WELLED RADIANCE

The red trace, second from the bottom, in the figure shows the contribution from the up-welled radiance, which can be seen as the most influential source to the sensor reaching radiance, effectively muting a portion of the absorption feature effects from atmospheric transmission. In this example, the up-welled radiance would comprise 30% of the total sensor reaching radiance. The black trace, third from the bottom, represents the down-welled radiance spectrum at the sensor for a perfectly reflective surface, while the blue trace is the surface arriving down-welled spectrum. Note the large spectral feature due to ozone from 1000-1060 cm^{-1} present in the down-welled radiance at the sensor. However, the amount of this source radiance that reaches the sensor is dependent on the reflectivity of the surface, which in this example is modeled by a gray body at a constant 0.9 emissivity. The resulting sensor reaching reflected down-welled spectrum is the maroon trace at the bottom of the chart and it shows a near-zero contribution from this source of radiance. Clearly, the magnitude of the contribution from these sources is dependent on not only the reflectivity of the surface material, but also the path distance to the sensor, which is to say the amount of atmosphere that radiates, and the temperature of the radiator as well. The following

panels show the differences and similarities between the up and down-welled radiance profiles for each atmospheric model used in this work.

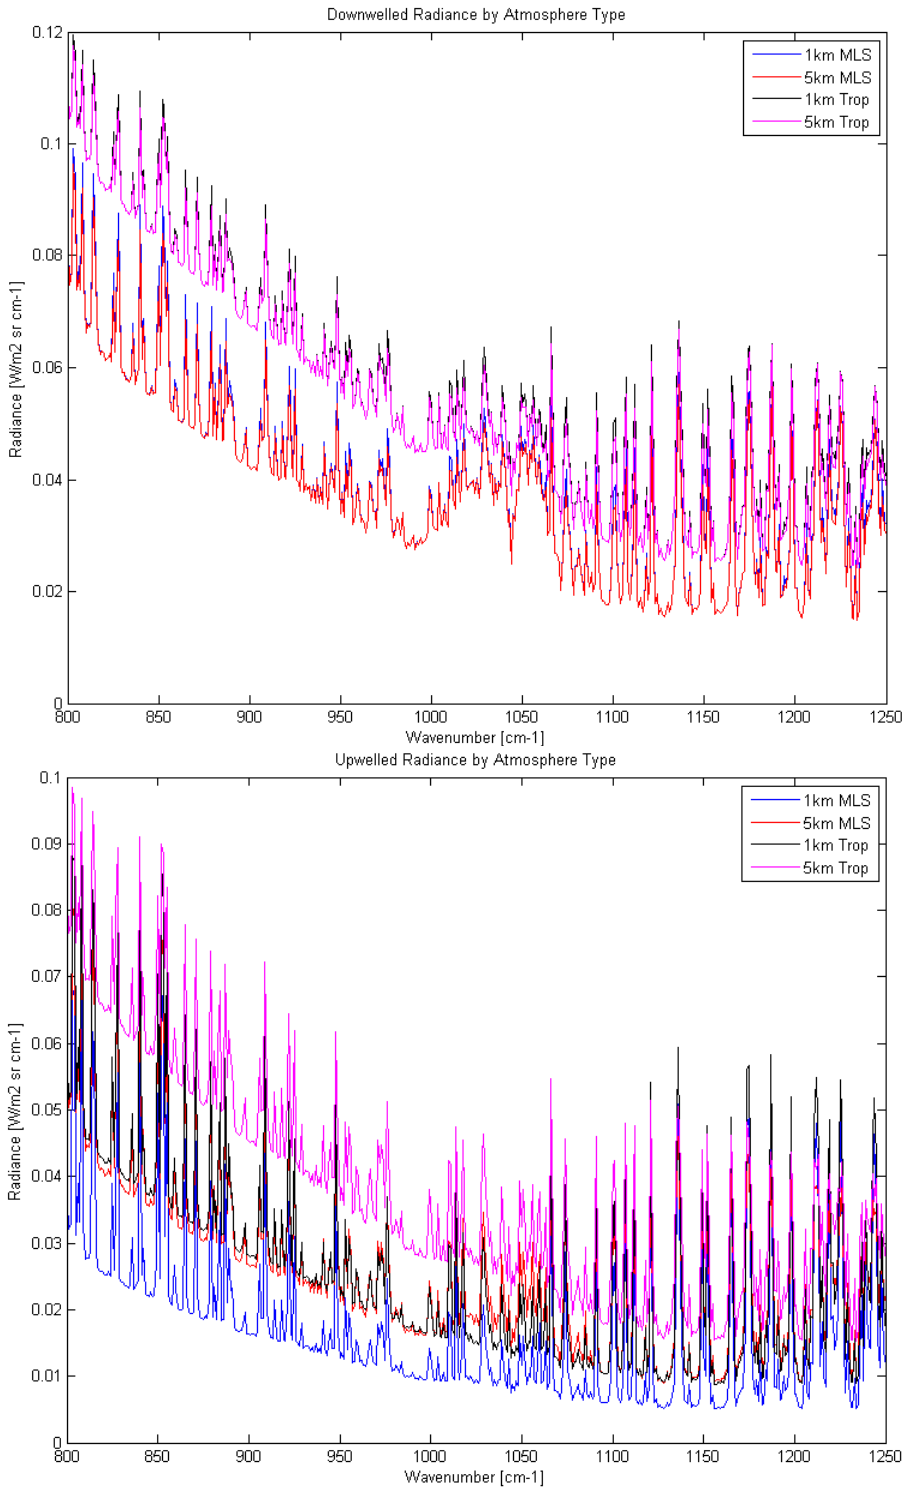

FIGURE 2.9 UP (BOTTOM) AND DOWN-WELLED (TOP) RADIANCE BY ATMOSPHERE TYPE

Figure 2.9 shows the spectral nature of each profile type. Both profiles are in emission as they are generated by the self-emission of the atmosphere. The downwelled radiance shown in the top panel represents the entire atmosphere's emission impinging on the surface. Note the substantial contribution from ozone. Also note however, that the strength of the emission is more dependent on atmospheric profile than altitude. This is because the altitude given is for the sensor height, whereas the downwelled comes from the contribution of the entire atmosphere. Therefore, the only influence to the total radiance comes from the amount of radiators in the atmosphere. With the tropical profile containing twice the water content of the MLS, there is a significant difference in the amount of radiating species. The upwelled radiance, which considers the self-emission of the atmosphere into the path of the collected radiance, is dependent on how much atmosphere is contained between the ground and sensor. As expected, the higher the altitude, or larger number of emitting species, the higher the contribution to the total radiance. Of interest is the proximity of the 5km MLS upwelled radiance profile to the 1km tropical profile. The additional 4km in altitude of the MLS profile is approximated by the increase in water vapor of the tropical profile.

To aid in the development of a sensor reaching radiance model, several assumptions can be made. For the moment, assume that if primarily natural surface materials are considered, the assumption of no reflected radiance serves to eliminate many second order terms including down-welled radiance, that when taken in bulk, have minimal overall effect. However, up-welled radiance, generated in the atmosphere along the sensor line-of-sight, becomes an important consideration and will contribute to the total sensor-reaching radiance. This quantity varies with wavelength at a given altitude across the LWIR window chosen for this work and can be modeled as a wavelength dependent additive term. To this point the sensor reaching radiance is modeled by equation 2.2,

$$L_\lambda = \varepsilon_\lambda B_T \tau_\lambda + L_{u,\lambda} \tag{2.2}$$

where λ is the wavelength of interest, T is the temperature of the background, B, which is modulated by a wavelength dependent emissivity term, ε, and the atmospheric transmission, τ, and the last term represents the up-welled radiance with units: [W m^{-2} sr^{-1} μm^{-1}]. The magnitude of both the atmospheric transmission and up-welled radiance is of course dependent on the altitude for which the model is implemented. To this point the model simply illustrates the sensor reaching radiance due to the self-emission of a surface material through the atmosphere with an additive up-welled radiance.

Finally, the target gas itself can be modeled. The most troublesome assumption here is that we can model the gas as a black body radiator. At first glance, this may seem counterintuitive in that the conditions associated with a black body do not seem to apply to a dispersed gas plume. While black body conditions may be true for photons generated via energy dissipated from nuclear fusion in the sun, as internal photons work their way through the comparatively dense chronosphere, this does not seem to hold true for fugitive gas emissions. One would typically associate optically thin, low concentrations of dispersed gas with plumes generated from fugitive-type emissions. The approach here is to assume local thermal equilibrium (LTE) for a well established/thermalized plume and use this temperature for the black body model, then treat the emissivity value as a modulator for a

selective radiator. Essentially, if one can link the value of the emissivity via the absorptivity of the gas, the black body model will provide the estimate of thermally available energy at each line position, which is then modulated via the absorptivity of the gas. It turns out that establishing the link from emissivity to absorptivity is rather straightforward.

First, we start with Beer's law, which states that the absorbance at a given wavelength is directly proportional to the absorptivity of the given species, the path length of the measurement through the species, and the concentration of the species. This relationship is shown in equation 2.3

$$\alpha_\lambda = \kappa_\lambda l c \tag{2.3}$$

where absorbance is on a per wavelength basis and is unitless, kappa is the absorptivity on a per wavelength basis with units of reciprocal concentration and reciprocal distance [ppm^{-1} m^{-1}] , *l* [m] is the path length in units of distance, and *c* is concentration [ppm], often in units of parts per million. While epsilon is historically used to represent absorptivity, kappa has been substituted to prevent confusion with the symbol for emissivity. Figure 2.12 shows an example of an absorptivity spectrum of a gas in the LWIR.

Recalling the conservation of energy for photon interaction with a solid mentioned in section 2.2, if we assume that the gas is optically thin and has essentially no reflection, the only terms left are transmission and absorption. Starting from this assumption, we next assume that the system is in LTE and can then say that Kirchhoff's law of emissivity being equal to absorbance applies. Then, it can be shown that the transmission through the gas is one minus the emissivity, which is equal to the expression for absorbance shown above. This development is shown in equation 2.4 as

$$\alpha + \tau + \rho = 1 \ \ (\text{where } \rho = 0\ )$$

$$\varepsilon = \alpha = \kappa l c \ \ (\text{at LTE})$$

$$\tau_\lambda = 1 - \kappa_\lambda l c \tag{2.4}$$

where the last step shows the equation has been solved for transmission. These equations allow for a direct replacement of the emissivity term with the absorptivity and concentration-path length value for the modulation of the gas emission. The transmission equation allows the background to be modulated by the absorptivity of the gas as it passes through the plume as prescribed by Beer's law. Beer's law has been shown to remain a linear relationship for low concentrations up to 250,000 ppm (Skoog, Holler, & Nieman, 1998).

This relationship can now be added to equation 2.2 with one important caveat. It must also be considered that the sensor will collect background radiance that has propagated through the plume and the atmosphere. Further, the radiance generated by the plume itself also travels through the atmosphere with an anticipated attenuation. Combining all of these concepts results in the equation below

$$L_\lambda = \left[\varepsilon_{\lambda,S} B_{T,S}(1-\kappa_\lambda lc) + \kappa_\lambda lc B_{T,P}\right]\tau_{\lambda,a} + L_{u,\lambda} \qquad (2.5)$$

where the new *S* and *P* subscripts indicate a parameter belonging to the surface or the gas plume, respectively. The subscript *a* attributes the transmission attenuation to the atmosphere. Note that the background radiance is attenuated by the transmission through the gas plume, whereas the sum of both the background and plume radiance is attenuated by the atmosphere. One may ask where the term is that represents the up-welled radiance generated from the atmosphere below the plume and is attenuated by the plume as it travels toward the sensor. This term may be neglected if it is assumed that the plume is sufficiently low to the ground that the space between it and the surface is negligible, resulting in minimal up-welled radiance. All gas plumes in this work are considered to be gas leaks and will therefore be essentially at ground level, supporting the assumption. This equation represented the SRR model for the early portion of this work, where surface materials only had up to a maximum of 15% reflectivity and resulted in minimal contributions from down-welled radiance.

Due to the fact that not all surface materials encountered in the real world will be near perfect black bodies, removing the assumption of neglecting the down-welled radiance is in order and results in the following modification to equation 2.5

$$L_\lambda = \left\{\left[\varepsilon_{\lambda,S} B_{T,S} + \left(1-\varepsilon_{\lambda,S}\right) L_{d,\lambda}\right](1-\kappa_\lambda lc) + \kappa_\lambda lc B_{T,P}\right\}\tau_{\lambda,a} + L_{u,\lambda} \qquad (2.6)$$

where the new radiance term, L_d, is the down-welled radiance, modulated by the surface reflectance and the gas plume, then transported to the sensor via the atmospheric transmission. This equation represents the radiative transport model used in this work for a single gas, atmosphere, and surface material and as handled by the DIRSIG SRR model covered in section 3.1.7.

2.4.2 Characterizing Model Behavior

A close analysis of equation 2.6 shows that in spectral regions where the absorptivity is near zero, the gas term becomes negligible and the background term dominates. Conversely, when the plume is maximally absorbing, the transmission of the background radiance drops to zero. Aside from the concentration-path length product term, which dictates the "strength" of the given absorption, the temperature difference between the background and plume is a major driver in the "visibility" of the plume spectrum over the background. If the plume contains sufficient thermal energy, it will appear in emission and essentially ride on top of whatever the background radiance may be beneath it. If the plume lacks thermal energy, it will appear as an absorption feature due to the selective absorption of energy from the background radiance as it passes through the plume. This behavior is essentially identical to transmission of the background through another layer of atmosphere – albeit of differing composition. One can now see that in order to detect the presence of the plume in the observed spectrum, there needs to be sufficient thermal contrast between the plume and background.

Figure 2.10 demonstrates the concept of thermal contrast and was generated using both 290K and 310K CH_3Cl at 11,500 ppm-m with no atmosphere, over a flat 0.9 emissivity black body background at 300K giving a +/- 10K thermal contrast. Note that while the magnitude of the thermal contrast is

the same, the magnitude of the radiance values for a given absorption feature in emission compared to absorption are not. This is a result of the emissivity value being less than one. As can be seen from equation 2.6, at an emissivity of one and with a thermal contrast of zero, no spectral features should be visible, as expected. However, if the emissivity is reduced, even when the thermal contrast is still zero, the gas will appear in emission. In fact, depending on how low the emissivity becomes, the thermal contrast can be negative indicating absorption, while the spectrum will still be in emission. Said another way, there must be a difference in brightness temperature between the gas and surface for the plume to be visible in the recorded spectrum.

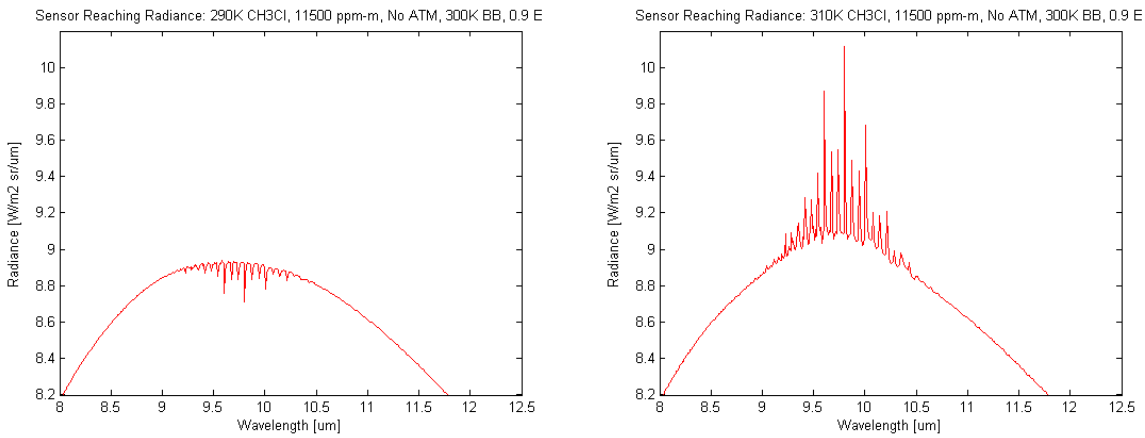

FIGURE 2.10 THERMAL CONTRAST COMPARISON FOR CH_3Cl IN ABSORPTION (LEFT) AND EMISSION (RIGHT)

From a detection standpoint, it can be seen that for a given thermal contrast, emission is going to be the easier target. This is due not only to the greater magnitude of the radiance at a given thermal contrast, but also the fact that as the plume begins to thermalize (approach an isothermal atmospheric temperature via molecular collisions) it will be visible longer in emission as the gas temperature approaches the surface temperature and then passes through it. In addition, most of the features in an observed spectrum will be due to atmospheric absorption and anything that is in emission stands out quite readily – unless of course the spectral feature in emission directly overlaps an atmospheric absorption feature, in which case the feature will be muted. To further illustrate this point, an example has been prepared for 295K ammonia gas over a 300K asphalt background shown inFigure 2.12fFigure 2.11. Ammonia was chosen because it has spectral features that extend across the entire LWIR band, while asphalt has a strong restrahlen feature that gives a highly variable emissivity over the band. The thermal contrast would suggest that the observed spectrum should be in absorption, but this is only found to be true where the emissivity of asphalt is above 0.95; anywhere below 0.95 and the ammonia spectrum appears in emission.

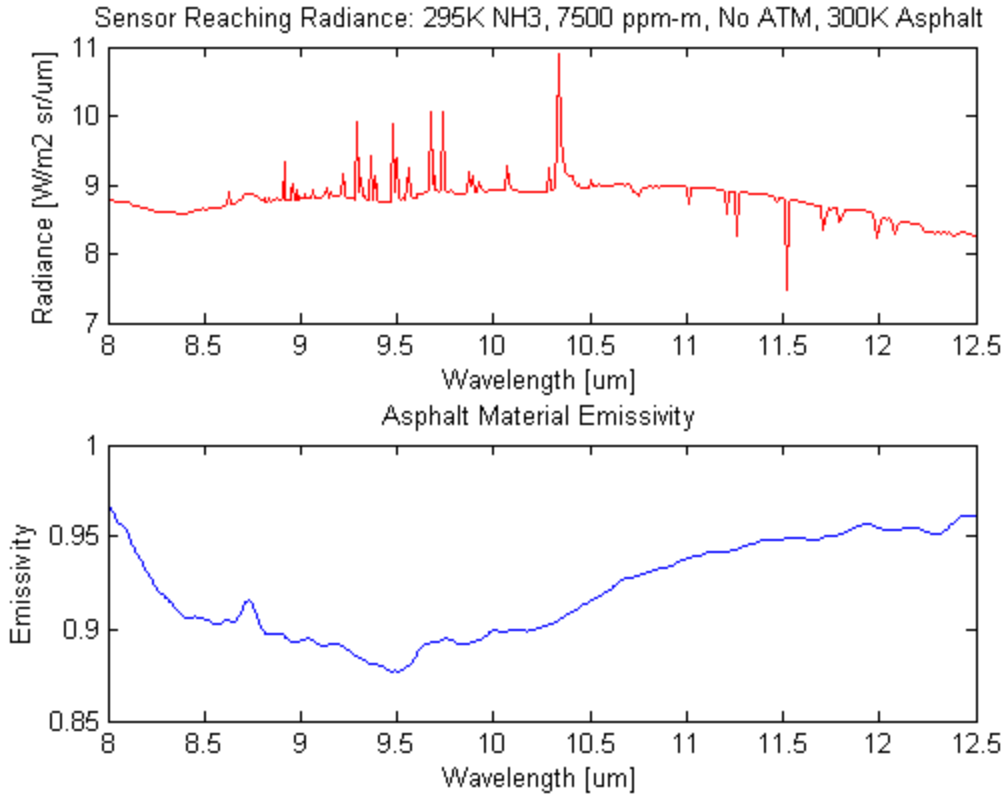

FIGURE 2.11 AMMONIA/ASPHALT IN BOTH EMISSION AND ABSORPTION (TOP) AND ASPHALT EMISSIVITY (BOTTOM)

Clearly, any attempt at species identification, or even detection, via a quantitative spectral scoring metric will require accuracy in both the absorption and emission regimes.

2.4.3 PNNL Gas Absorptivity Library

Absorptivity values for all gases used in this work were derived from the Pacific Northwest National Labs (PNNL) Quantitative IR database for Infrared Remote Sensing and Hyperspectral Imaging. The PNNL spectral gas database is a compilation of over 500 gas-phase spectra collected in pressurized vessels to simulate the anticipated atmospheric broadening one would observe under typical remote sensing conditions (Sharpe, Sams, & Johnson, 2002). The data are collected with high resolution (0.10 cm^{-1}, ≈1 nm at 10 µm) FTS systems at three different temperatures: 278, 298, and 323K. The approach of this work is to use the 298K data and change the black body temperature of the gas model within a bracketed range of this value. While further extrapolation of the absorptivity values between those provided by PNNL at the three baseline temperatures may result in slightly more accurate values, making an assumption as to the behavior of the absorptivity over these temperature ranges (i.e., linear) may introduce more error than benefit.

The absorptivity values are readily obtained from the PNNL spectra due to the nature of their format. Spectra, as shown in Figure 2.12, are reported under collection conditions of 1 ppm-m. This allows one to simply apply Beer's law in deriving the absorptivity from the spectrum by dividing the absorbance by the concentration and path length terms. This essentially means that the spectra are

reported in values of absorptivity, making applications to modeling efforts seamless. Figure 2.12 is an example of the PNNL output for an absorption feature of CH_3Cl gas at 298K in the LWIR.

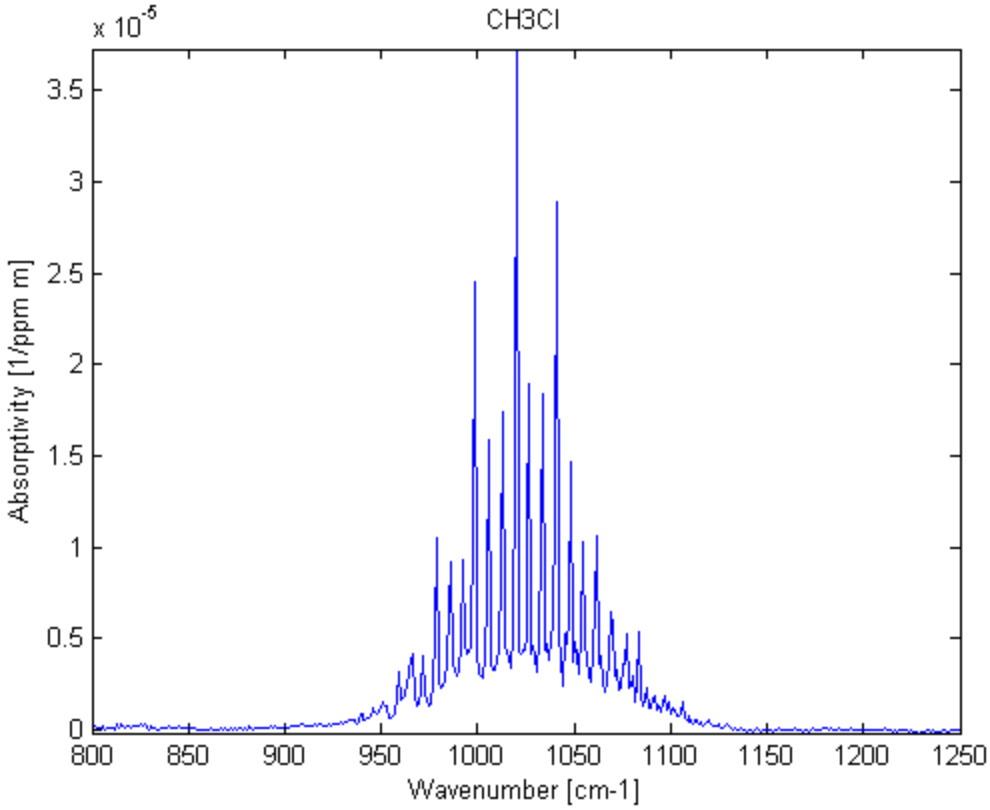

FIGURE 2.12 PNNL DATABASE ABSORPTIVITY SPECTRUM FOR CH_3Cl GAS IN THE LWIR

2.5 Fourier Transform Infrared Spectrometers

While dispersive element spectrometers are common-place and are conceptually straightforward to understand, FTS is a comparably newer technology and much less readily understood. The advent of the technology behind FTS dates back to A. A. Michelson and his Michelson Interferometer (Shepherd, 2002). Michelson designed a set of mirrors that split the path of light such that two paths are traveled by essentially equal magnitudes of half of the original source beam. One mirror at the end of one of the paths can be moved so that an optical path difference (OPD) is setup between the two paths. The beams are then recombined and output to a detector. Upon recombination, depending on the phase of the wavefront, the two beams will either constructively or destructively interfere. A constant sweep of the mirror will produce output of varying intensity as wavefronts move in and out of phase. Figure 2.13 below, from Beer, illustrates a conceptual Michelson interferometer as a spectrometer. It is important to note that in general, an additional optic known as a compensator is placed along the fixed optical path to account for the additional phase shift introduced by the beamsplitter as the light takes separate paths. The compensator adjusts this

phase shift so that the only phase difference between the two beams upon recombination is due to the difference in optical path traveled.

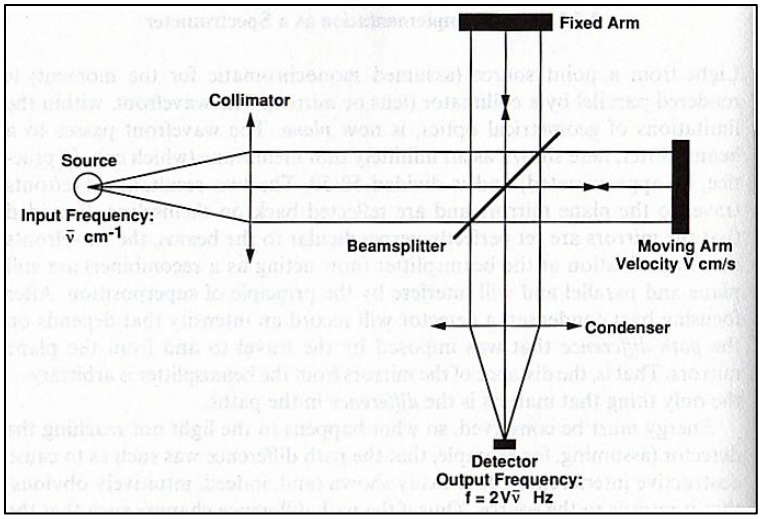

FIGURE 2.13 MICHELSON INTERFEROMETER AS A SPECTROMETER (BEER, 1992)

2.5.1 FTS Advantages

FTS are known to have several advantages over dispersive element spectrometers (DES) (Carter, Bennett, Fields, & Hernandez, 1993). One such advantage is called the Jacquinot advantage. The Jacquinot advantage describes the system etendue, which is a measure of the product of the maximum beam area and solid angle of the beam passing through the system. In spectrometers, the limiting optical element is typically the most expensive element to manufacture in a large size. For dispersive systems, this would be the grating and for FTS, the beam splitter. For comparably sized instruments, an FTS can be shown to have a system entendue over 150x greater than that of a dispersive system (Beer, 1992). While it is true that a much larger grating can be made to compensate for the system entendue as compared to a similar sized beam splitter, the FTS can maintain a substantially smaller and more compact instrument footprint while maintaining a similar entendue.

The Fellgett advantage describes the frequency-division multiplexing advantage of an FTS as compared to a DES. Initially, this advantage was made in comparison of the detection schemes of FTS and DES in which scanning a dispersive element across a single photodiode to collect an entire spectrum was compared to the simultaneous observation of all transmittable frequencies in a given spectrum in an FTS. With the advent of arrayed detectors, a DES can now collect multiple 'bands' across a given spectrum simultaneously. The problem here is that the size of the resulting spectral coverage is dependent on the amount of dispersion associated with the element, which ultimately becomes limited by the sensitivity of the detector for a given source intensity.

If an arrayed detector was introduced into an FTS, the focal plane could be arranged such that each detector element viewed a different spatial extent of the scene resulting in a spatially sampled interferogram, rather than a spatially integrated one. Conversely, the output of the DES must have

one dimension dedicated to the dispersed spectral information, while the other dimension would be a one-dimensional spatial extent. This is because a DES images the entrance aperture slit through the entire optical system, which only allows for finite spatial sampling. Thus, to build-up a spectrally resolved spatial scene would require a scan of the entire spatial scene, either optically or through the motion of the platform in a push broom fashion, to build-up individual spatially correlated spectra.

In other words, while a hypercube at a given resolution and number of bands must be built up via many scans of a DES, an FTS can build a hypercube of the same spatial scene in one scan. This is accomplished because the focal plane is defined spatially as having X and Y extent at a given time-position product of the OPD. A new spatially resolved interferogram is sampled from the focal plane at predetermined points in OPD building a time series of spatially sampled interferograms through a single travel of the moving mirror. When subjected to a Fourier transform, a single X-Y plane results in individual spectra associated with each spatial pixel position – a hypercube.

2.5.2 FTS Instrument Resolution Determining Factors

FTS systems are well known for providing a higher resolution output as compared to comparably-sized DES systems. There are several factors to consider in understanding how resolution is determined from an FTS. The first is instrument line shape. In typical DES systems, the entrance aperture is a narrow rectangular slit, while the exit is also a slit. Convolving these two rectangular functions ideally results in a triangular line shape that modulates the spectral output. Further, the ruling of the grating contributes to the spectral resolution as well.

Ostensibly, in an FTS, if one were able to build an infinitely long interferogram, a transform of this would result in a perfect reconstruction of the input. Unfortunately, the interferogram length is controlled by the total displacement of the mirror, and consequently, the maximum OPD. As a consequence of having a finite length with an abrupt start and end point, the interferogram appears to be modulated by a rectangular function. This function has the effect of a perfect notch filter; passing everything between two bounds, while extinguishing signal outside of the pass range. When a Fourier Transform is taken of a signal modulated by a rectangular function, because this is modeled as a linear system, the result is a transform of the interferogram convolved with the transform of the rectangular function. The end spectrum essentially consists of a series of line positions broadened by the sinc function. An example of a sinc function is shown in Figure 2.14.

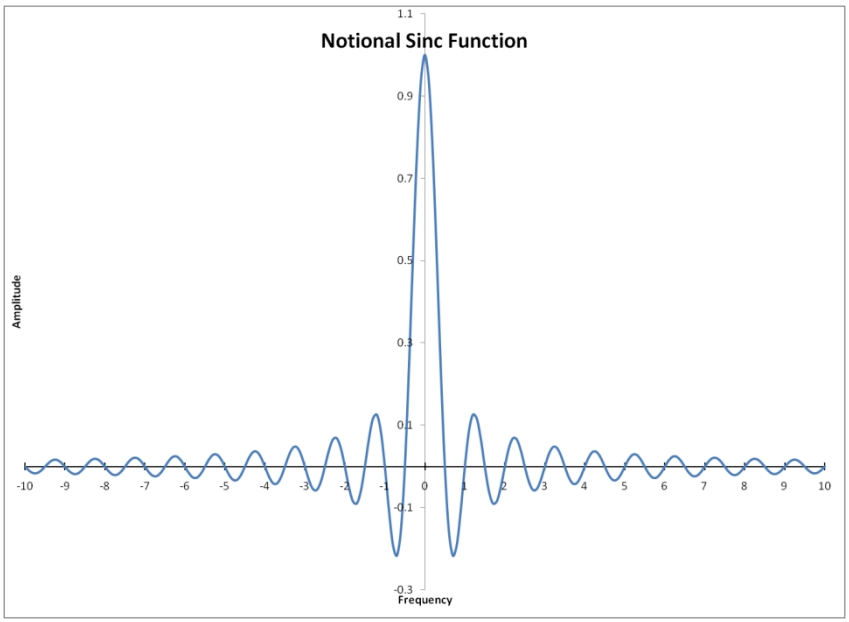

FIGURE 2.14 NOTIONAL SINC FUNCTION

A sinc function serves to broaden the shape of the spectral features contributing to reduced resolution, but a more troubling aspect is the effect of the side lobes. The negative and oscillating amplitude can induce spectral line shapes that are completely baseless in physical origin. While the magnitude of the effect is minimal, attempts at reducing the influence of the side lobes can be accomplished via apodization. Apodization, meaning "without feet", is the process of smoothly curtailing the ends of the sampled function such that it smoothly decays to zero, rather than abruptly ending. This results in a convolution in the transformed domain with a function that has only positive, minimal side lobes – essentially the square of the sinc function. While this step serves to remove the negative side lobes resulting in minimal spectral artifacts, the resulting function pushes that energy back to the main lobe and creates a further broadened spectral line shape. For most applications this is a reasonable trade-off to maintain greater spectral fidelity. Additional, more complex methods of apodization, such as the Hann and Hamming windows can be used, but are generally unnecessary for purposes of the remote sensing of gaseous targets. This is due to the nature of interferogram content in which most of the large amplitude components are found near the zero OPD point, reducing the affect of the amplitude behavior of the apodization function across most of the interferogram. As long as the interferogram is smoothly decayed to zero at the maximum OPD, the prior effects seem to be minimal – particularly since they represent the highest frequencies in the spectrum. Figure 2.15 compares the transform of the triangular apodized function with the original sinc function. Note the broadened central lobe, which is twice as large at the base as the original sinc function and the dramatic reduction in the side lobes.

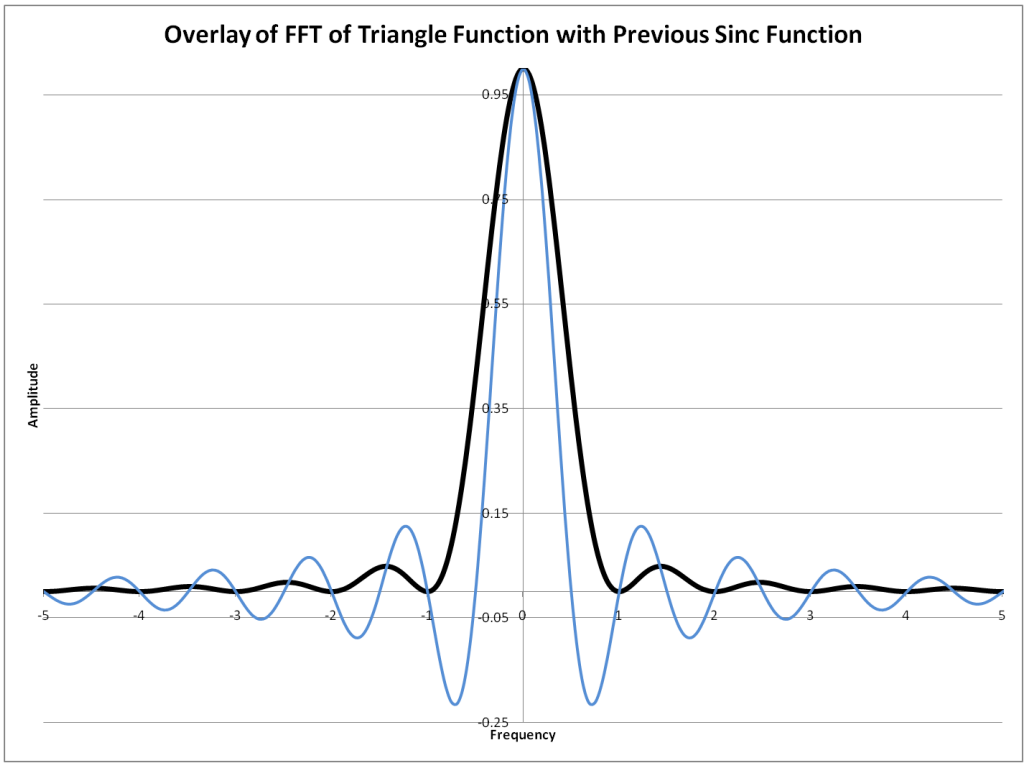

FIGURE 2.15 COMPARISON OF FFT OF TRIANGULAR APODIZING FUNCTION WITH ORIGINAL SINC FUNCTION

If we consider the instrument line shape without apodization, the major contributor to the breadth of the main lobe and resulting side lobes is inversely proportional to the maximum OPD. Thus, the greater the mirror travel, the finer the spectral line shape and consequently, the finer the spectral resolution. This relationship is shown in the equation below

$$f(\tilde{\upsilon}) = \frac{\sin(2\pi\tilde{\nu}L)}{2\pi\tilde{\nu}L} = \mathrm{sinc}(2\tilde{\nu}L) \qquad (2.7)$$

where $\tilde{\nu}$ is the frequency in reciprocal centimeters and L is the maximum OPD in centimeters. Figure 2.16 compares the same sinc function used in the figure above with one that has a maximum path difference 5x longer.

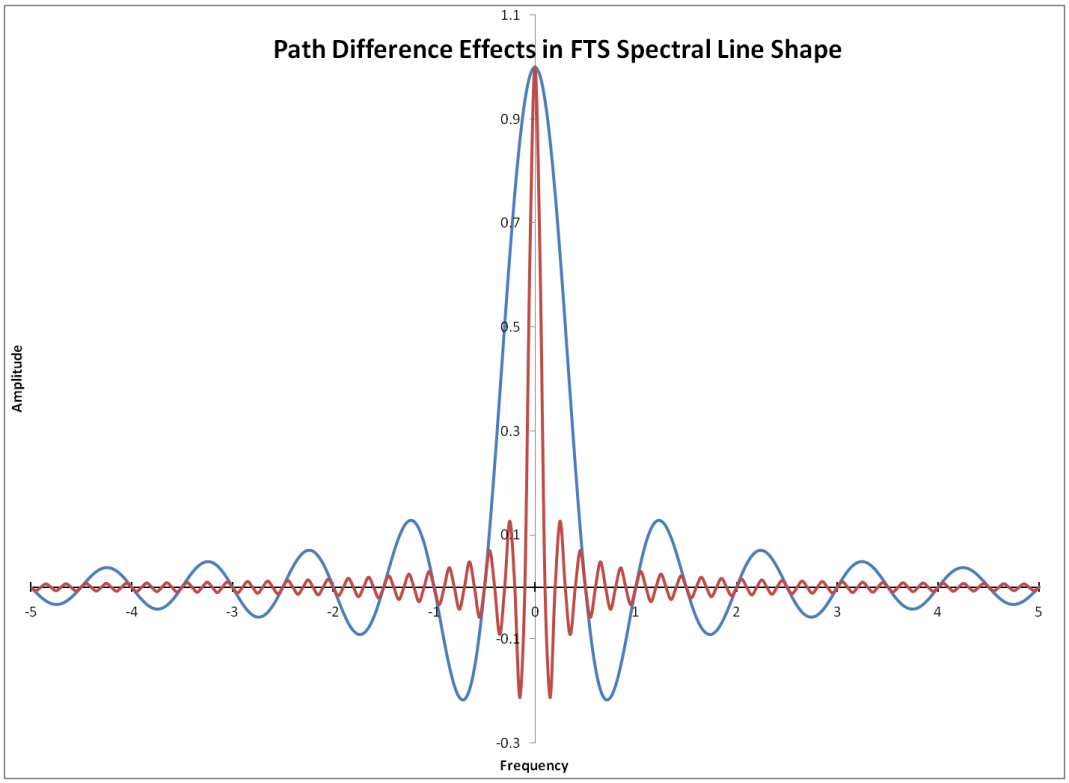

FIGURE 2.16 COMPARISON OF 5X LONGER PATH DIFFERENCE TO FTS SPECTRAL LINE SHAPE

As can be seen from the figure, a dramatic narrowing of the main lobe occurs, drastically improving spectral resolution in terms of instrument line shape. In general, one can determine the full-width at half-maximum (FWHM) spectral line shape for purposes of comparing spectral resolution using the approximation that the FWHM of a unit sinc function is given by 0.60336 (unit sinc width evaluated at the one-half maximum amplitude point) divided by the maximum OPD (Beer, 1992). This shortcut takes the nominal FWHM of a unit sinc and then scales it by the maximum OPD. For the commonly used triangular apodization function, the FWHM of the instrument line shape is approximately one, leaving the estimate of the resolution to just be the reciprocal of the maximum OPD. Unless stated otherwise, a triangularly apodized derived resolution will be used when quoting the instrument resolution and consequent maximum OPD values for simulated spectra in this work.

Since the wavenumber scale and instrument line shape of an FTS can be precisely determined and are independent of the given wavenumber, the spectra produced on an FTS are essentially standardized among FTS systems (Carter, Bennett, Fields, & Hernandez, 1993). This is known as the Connes advantage of the FTS and it allows for easy comparison of FTS spectra to standard reference spectra, whereas DES spectra do not maintain a constant instrument line shape across the wavelength scale and require extensive calibration before any comparison between sources can be made.

2.5.3 FTS Interferogram Formation

From the standpoint of resolution and throughput, one can see why the FTS is aptly suited to the task of remote sensing of gaseous targets. The FTS is best positioned to detect the comparatively weak emission from the likely small concentrations of gas generated from fugitive emissions, and further, to resolve the fine spectral details of that gas. Of course, what hasn't been mentioned is the operational requirement of an FTS that the source must maintain the same spectral radiance and physical position during interferogram collection. With a DES, should the source position move, or spectral radiance change, the result is a spatially smeared and spectrally, or radiometrically, averaged result over the integration time of the detector. While these results are undesired to be sure, they can still be useful. Conversely, should these changes occur to the source while developing an interferogram, severe spectral artifacts are induced in the hypercube. This is because an interferogram is a sample of the interference pattern of the light at the product of a given point in time with the length of the mirror travel. The spatial and spectral relationships of the source must be maintained over the entirety of the time series for the transform of the resulting samples to have physical meaning. Clearly, this causes a problem when using an FTS to collect temporally and spatially transient events (as in the propagation of a gas plume) from a moving platform. This effect was characterized in work accomplished by Mitchell and coworkers (Mitchell, Hemmer, Lewis, & Salvaggio, 2001). The group demonstrated how temporally varying radiance negatively impacts the collected interferogram.

The obvious solution is to complete a scan of the desired maximum OPD before platform or target motion induces changes in the source. This solution involves optimizing many operational parameters, which together comprise the bulk of the effort of this work. What follows is a brief description of an interferogram and the transform process so that a better understanding of the ill-effects of motion can be established.

In most FTS systems, a laser is used to trigger sampling events at predetermined lengths of OPD. This is easily accomplished in that an essentially monochromatic laser (generally He-Ne at 15798.0024 cm^{-1}) produces a regular sinusoidal variation with the phase difference introduced by the moving mirror (Beer, 1992). This difference in optical path will cause repeating intervals of constructive and destructive interference as the wave fronts of the two beams move in and out of phase. These intervals can be used to trigger the sampling of the interferogram. Very precise measurements of optical path can be made using this technique so that an OPD can be assigned to each sample point. It turns out that one need not sample at every triggering event. The theory follows Nyquist in that if the maximum frequency signal is already oversampled by a factor of two, then additional sampling won't improve the result. This concept clears a common misconception that suggests additional sample points past the Nyquist limit at a given OPD in an FTS will improve resolution. As was shown in Figure 2.16, it is a larger maximum OPD, and therefore larger interferogram, that results in better resolution – not additional samples over the same OPD.

Meeting the minimum sampling requirement in the LWIR is not difficult and results in a factor of "data compression" not otherwise attainable if the spectral range included shorter wavelength regions. The compression results in that some integer multiple of laser oscillation zero crossings

would trigger the sampling event, rather than creating a sample at every zero crossing. In this case, to sample without aliasing in the LWIR, a spacing of six laser wavelength intervals will allow a Nyquist sampling rate of the LWIR up to 1316.5 cm^{-1}. Of course, to prevent internal aliasing, an optical bandpass filter needs to be installed to roughly cover the range of interest.

2.5.4 Anatomy of an Interferogram

As previously discussed, interferograms consist of a series of radiance measurements at specified OPD intervals that correspond to the interference pattern of two beams of light that differ in phase. From a Fourier Analysis stand point, it is easiest to envision successive segments in the interferogram as corresponding to the level of sinusoids present of a given frequency if one were to completely reconstruct a spectrum by approximating it with overlapping sinusoidal functions. Spectral features found in the spectrum that vary slowly would correspond to combinations of low frequency sinusoids and this contribution would be found early in the interferogram formation. In visualizing the interferogram this way, one can see why the longer the maximum OPD, the finer the instrument resolution. This is because higher and higher frequency sinusoids are included in the spectral reconstruction, which result in greater fidelity to the original signal. Conversely, getting the first several OPD points correct is critical in achieving a radiometrically accurate measurement as these points essentially carry the signal bias. The very first point is known as the “center burst”, or “white light fringe” (WLF), or zero path distance (ZPD) where there is no difference in optical path and all transmitted frequencies of light pass without destructive interference. The series of triangularly apodized, single-sided interferograms shown in Figure 2.17 illustrate the enhancement to spectral detail through finer resolution as the entire interferogram is formed. The series simply uses the PNNL absorptivity data for ammonia and ranges from 0.01 to 1 cm maximum OPD. The panels can be directly compared to the corresponding transformed spectra found in Figure 2.18.

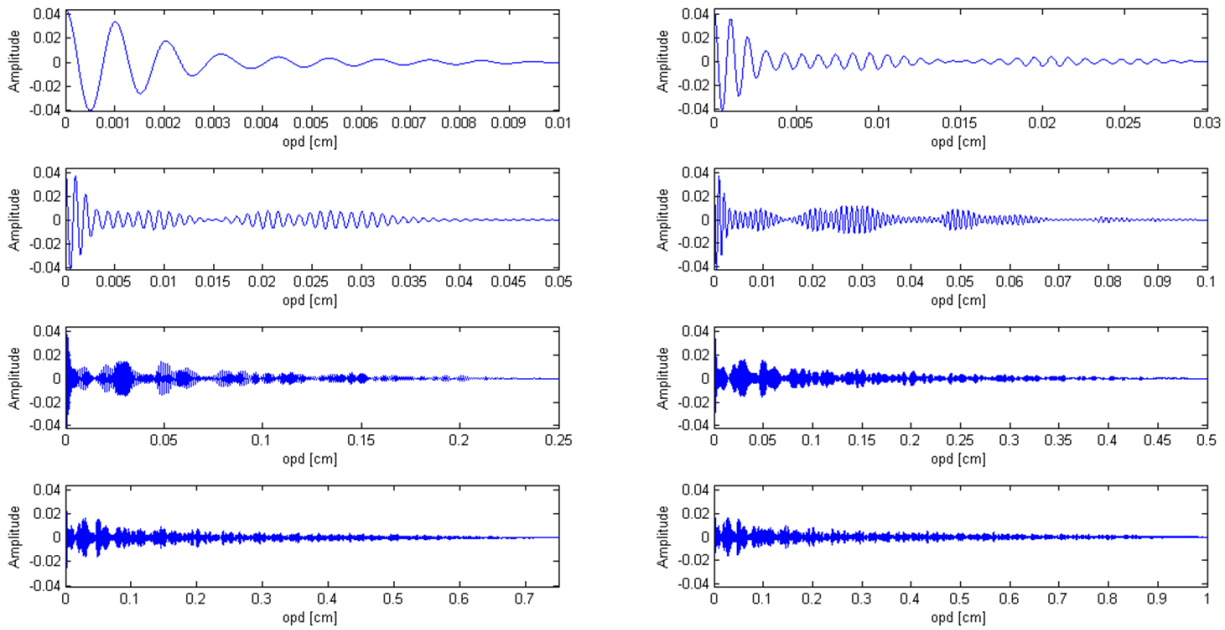

FIGURE 2.17 INTERFEROGRAMETRIC SERIES OF NH_3 BY MAXIMUM OPD

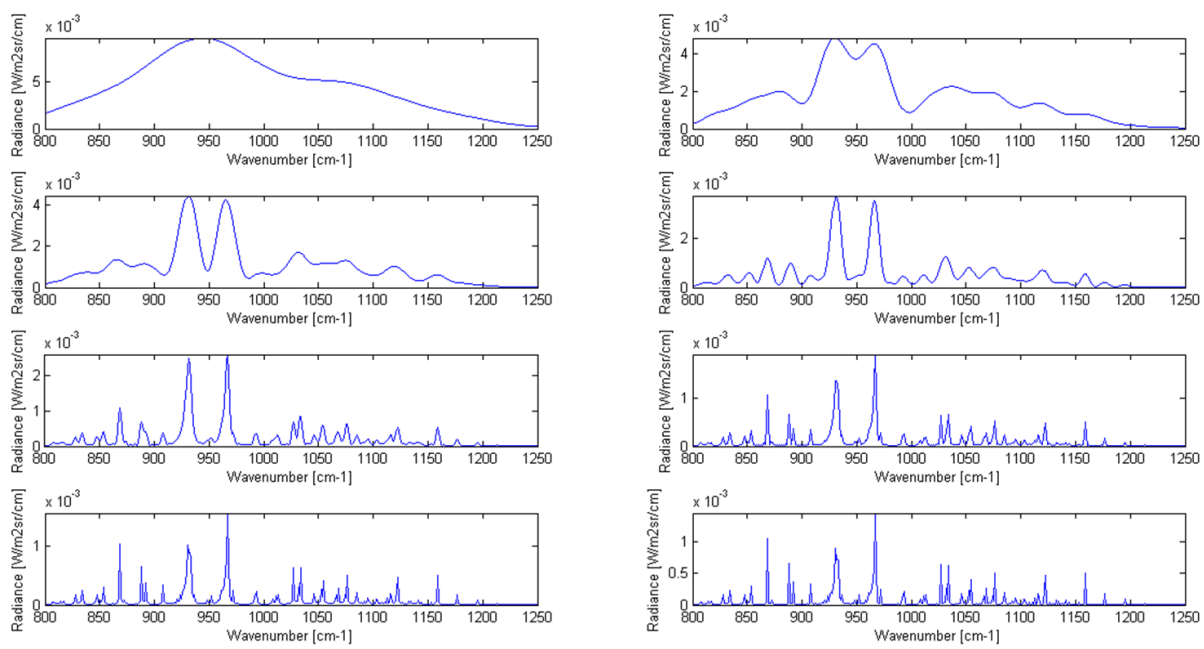

FIGURE 2.18 SPECTRAL SERIES OF NH_3 BY MAXIMUM OPD

The series progression clearly shows how the first several OPD samples establish the foundation to the spectrum, while each successive increase in maximum OPD allows for the inclusion of finer spectral detail until all features are fully resolved. The concept of different elements of spectral features appearing in different and distinct regions of the interferogram is critical in understanding how FTS instruments are affected by error and how to recover evidence of a spectral feature in the

face of such error. Should a portion of an interferogram be lost or corrupted, knowledge of the location of this region would give a general idea as to how the spectral feature of interest would be impacted.

To be clear, the example shown is for interferograms generated from the spectral signature of the gas alone. It does not consider any variances or radiance bias due to temperature effects and contains no background or surface material spectral information. It is purely the regions in the interferogram that are generated from corresponding spectral features in the wavespace. That said, prior work has been done in injecting the pure interferogram of a sample gas into previously collected real interferograms generated from scenes collected in the field. Since the Fourier transform is a linear process, should the gas spectral collection conditions mirror that of the site, the interferograms can add linearly (Flanigan & Quinn, 1996). This lends credence to the process of identifying regions in the interferogram where one would expect to find evidence of a spectral feature from a gas regardless of the background and surface material spectral features that may also be present in the interferogram.

2.5.5 Errors Associated with Interferogram Formation in the Airborne Environment

One of the major influences on the fidelity of the reproduced spectrum from an airborne FTS is error introduced to the recorded interferogram from disturbances to the platform in flight. The two major sources of interferogram error are incongruous OPD values for sample points due to uncertainty in the moving mirror and fluctuations in the sensor line-of-sight (LOS) due to vibration and turbulence known as jitter (Beer, 1992).

Jitter is caused by low frequency vibrations that accompany the motion of a platform and disturb the sensor LOS during interferogram collection. This results in artificial fluctuations to the recorded radiance values due to temporarily viewing another portion of the ground off-target, inducing spectral artifacts into the transformed spectrum. Jitter in an aircraft mounted FTS, for example, may be caused by engine vibration and turbulence. This motion induces a low frequency modulation on the collection of the interferogram. Jitter can be periodic, as in the case of engine vibration and associated beat frequencies from multi-engine aircraft, or aperiodic as in the case of turbulence. Jitter sources are most common below 300 Hz and can have a wide range of duration (Beer, 1992).

The impact to the transformed spectrum from the jitter can vary widely and is most dependent on where in the interferogram the jitter occurred. Should the jitter occur early on, the reliability of the transformed spectral radiance values are lost as the underlying bias to the spectrum has come from another source. Jitter becomes better tolerated, in terms of recovering a usable spectrum, the farther down it is encountered in the interferogram, affecting only the high frequency content.

While the effect of jitter on a mobile FTS system can be severe, there are techniques available via data acquisition and processing that allow for the reduction of many jitter effects. One method of jitter reduction, as reported by Beer, proposes to use second order (or higher) aliasing effects to build a frequency buffer against jitter (Beer, 1992). However, Bennett argues that this step is only necessary for non-imaging FTS systems, as imaging FTS systems allow for removal of the jitter

induced modulation due to the spatially correlated pattern observed across the focal plane array (FPA) (Bennett, 1997). He suggests that the variations in spectra produced by LOS jitter are highly correlated among the pixels in a given frame and due to this correlation, are easily removed via a noise-whitened Principal Components Analysis (PCA) approach. Should one determine the temporal and spatial nature of the jitter shift via a correlation approach, the interferograms could ultimately be improved by substituting portions of the correct interferogram from neighboring detectors on the array.

While knowing where to sample with a sufficient number of sample points is critical, one must also know the exact spacing between these points – or the sample distance. If there is a periodic position error of less than one nanometer in the sample points, ghosting results, while random error can be tolerated to two nanometers (Beer, 1992). This problem plagues the TurboFT FTS system at long OPDs and will be introduced in the next section.

To this point, it can be seen that an interferogram is a series of samples of the interference pattern produced by the interferometer, taken over time at specific distance intervals of a moving mirror. The mirror velocity is ideally constant, while the sampling takes place at determined optical path difference points along the way. It is clear to see that several oppositional relationships are critical in maximizing the performance of an FTS. If one wanted to complete a scan as quickly as possible, this would require a high mirror velocity coupled with a small maximum OPD. The consequences here are reduced spectral resolution coupled with a reduced SNR. So sensitivity is in contrast to resolution and mirror speed. Aside from these theoretical concerns, the faster the mirror speed the more uncertainty in the mirror position during the sampling interval, introducing error to the sampling distance and the resulting spectral radiance.

Mirror error can be thought of as assigning a radiance measurement of a neighboring sample point to the actual sample point because through a piston or tilt movement, the mirror optic has shifted when the sampling process for a given point was initiated. The magnitude of the error is questionable to begin with as the mirror is always in motion during sampling, whereas in ideal sampling, the mirror would be moved to a finite position, stopped, allowed to integrate, readout, accelerate to the next point, stop and repeat. This method is actually employed in what is known as "Step Scan" FTS, but requires a highly repeatable signal and is used to record highly transient events.

Step-scan essentially exploits the fact that the interferogram is built with discrete samples and consists of moving the mirror to the predetermined OPD sample positions and then stopping for as long as needed to gain the required sensitivity through integration. Unfortunately, this type of method is ideally suited to static or highly repeatable events. This is because, as mentioned above, the interferogram as a whole is really meant to be a snapshot in time of an event, but is assembled over some finite period of time. If the event source was static, time would not be an issue, but in the case of remote sensing, few targets of interest are entirely static – particularly relative to a moving platform. Typical uses of this technology include using a laser pulse as a source of excitation and then collecting a resulting highly transient, repeatable event at each mirror position. The observed system then de-excites, the mirror is moved to the next position in OPD and another pulse is fired

allowing for collection of the same event resulting in a different interference pattern (Griffiths & de Haseth, 2007). This method allows for the assembly of a spectrum from events that are so short-lived, no other spectrometer could capture them with sufficient SNR.

However, in most FTS systems the mirror is continuously moving as it makes a scan along the optical path, allowing some finite interval of time to pass where signal is integrated from a finite range of OPD. So there is some uncertainty to begin with and the result is certainly not a series of delta functions in the temporal domain.

Most in the literature tend to view mirror error as an ever-present evil and dismiss its effects as an additive source of random noise to the resulting spectrum (Cohen D. L., 1999), (Zachor, 1977). It is suggested that the distribution of mirror positional uncertainty is normally distributed around the intended sample point and that is how it is implemented in this work. It will be shown that while mirror error certainly corrupts the transformed spectrum, unless there are multiple critical sampling errors, the spectral features of interest are likely to be recovered. It is again the idea that all of the 'pieces' that make up a given spectral feature of interest are found in different regions of the interferogram and the likelihood that all of them will be corrupted is slim, resulting in some semblance of the feature in the transformed spectrum – a clear advantage over DES systems.

2.5.6 FTS Performance Parameters and their Interdependencies

It turns out that multiple parameters drive the SNR achieved in an FTS and their relationships suffer from many interdependencies. Figure 2.19 illustrates the interrelationships between the key parameters affecting SNR in a moving platform FTS system, where in this case SNR simply implies relative "strength of signal" and should not imply a quantitative value.

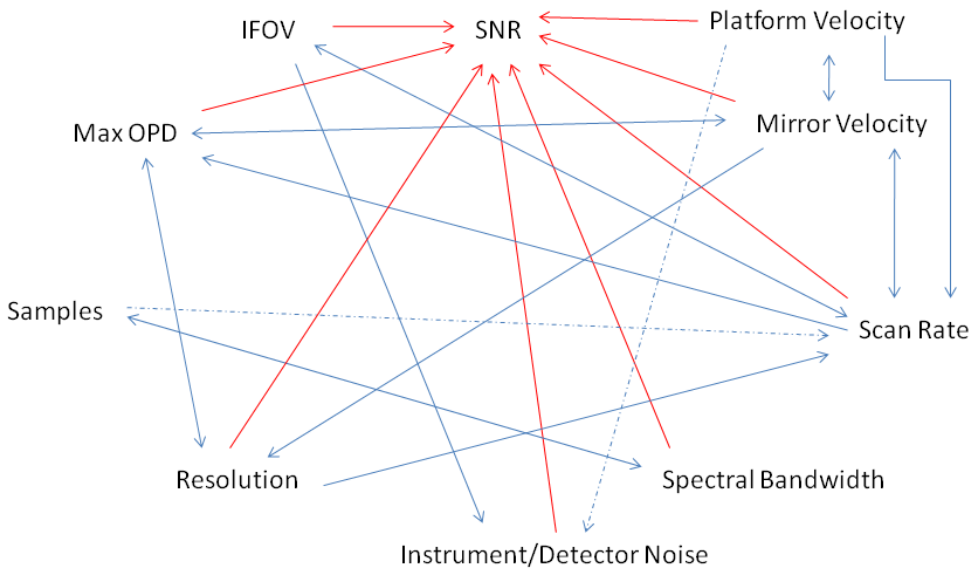

FIGURE 2.19 FTS PARAMETER INTERRELATIONSHIPS

The figure shows the direct parameter influences on SNR with red arrows, while the blue arrows show interrelationships between the parameters of influence. Said another way, the adjustment of

one parameter affecting SNR will require adjustment of the other parameters related to the direct parameter. In many cases the collection conditions will dictate the performance requirements for some parameters which, in turn, will drive all of the others. These other parameters must then be optimized for an SNR that results in detectable and distinguishable target signatures. The dashed lines that connect the 'samples' parameter to 'scan rate' and the 'noise' to 'platform velocity' suggest tenuous links between them that may be of concern under certain collection conditions. As an example, the link between noise and platform velocity suggests that the very movement of the collection platform induces mirror and pointing accuracy issues into the collection, which manifests itself as noise and/or spectral artifacts, which may or may not cause problems with target detection. The connection between samples and scan rate suggests a possible problem that may need to be addressed in that if the volume of samples taken over an array for a given integration time become too large, the readout of the detector delays the next integration from starting. This delay time limits the attainable scan rate of the system, which may cause a greater gap in both the desired spatial and temporal coverage of the scene. This problem further influences platform velocity and IFOV, which in turn affects the mirror velocity, maximum OPD, and resolution.

As can be seen, a single collection condition that needs to be addressed can have myriad effects that cascade through the performance settings of an FTS system. Additionally, a simplified diagram showing only the direct relationships influencing SNR is shown in a cascading fashion in Figure 2.20.

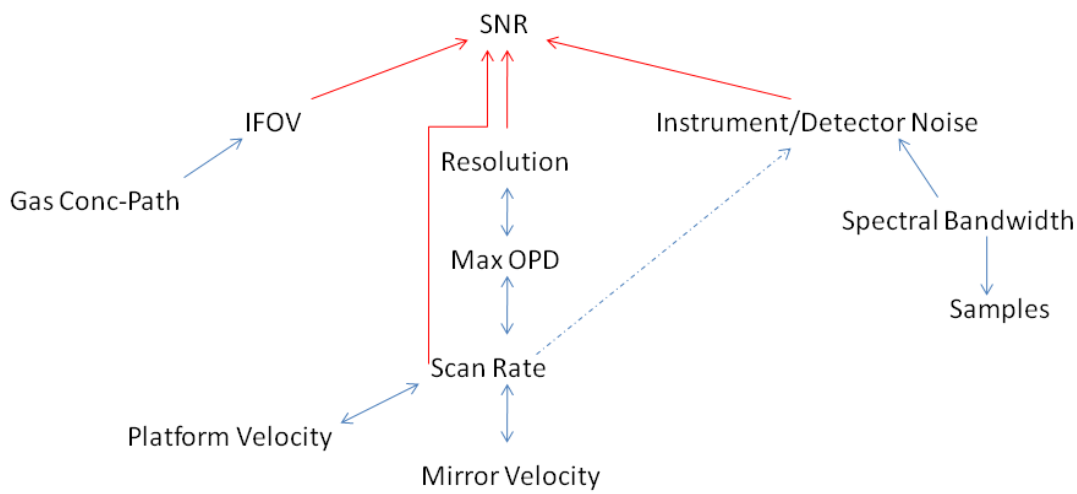

FIGURE 2.20 CASCADING PARAMETER DEPENDENCIES AFFECTING SNR

In reviewing these relationships, some links may not be obvious. Starting with mirror velocity and moving forward, one can see that mirror velocity will directly determine what scan rates are achievable by the system. The arrow is double-headed to show that if the required scan rate is already determined, the mirror velocity must be adjusted to accommodate the rate. Further, if the maximum OPD is known, the maximum scan rate can be found with the knowledge of the mirror velocity (cycle time) and total mirror travel distance. As shown previously, the maximum OPD is the direct driver of the resolution obtained in the transformed spectrum. In general, the platform

velocity, coupled with the IFOV will determine what scan rate is required to accomplish detection once the required frame overlap is determined.

The scan rate also has direct influence to SNR via the total collection time of the interferogram. The SNR is improved as the square root of the total collection time during a scan (Beer, 1992). Beer's expression for the SNR of an FTS is given below.

$$SNR = \frac{\delta\tilde{\upsilon}(signal)t}{\sqrt{[(N_{det/read}+N_{det.int.bknd.})+\Delta\tilde{\upsilon}(N_{instr.emis.}+signal)]t}} \tag{2.8}$$

where $\Delta\tilde{\nu}$ is the total bandwidth, $\delta\tilde{\upsilon}$ is the resolution, t is total time of the scan to complete an interferogram, and N is noise. The first noise term is attributed to detector and readout noise followed by the internal background noise inside the detector chamber, with the final term representing the noise due to instrumental emission. The SNR improvement due to scan time is related to the way random noise is incorporated into the transformed spectrum. Noise present during an interferogram collection is distributed evenly across the transformed spectrum, rather than singularly incorporated into a spectral channel as in a DES system. This is a benefit in that because a transformed interferogram consists of both a real and imaginary part with noise spread to both, removal of the imaginary part during the Fourier Transform process reduces the rms noise contribution. However, this phenomenon can cause problems when trying to detect weak signals among noise. Assuming that thermal and instrument noise are always present, Shot noise is proportional to the number of photons collected. In a DES, this results in major noise contributions when large peaks are recorded in a spectrum with mainly the thermal/instrument noise contributing to the baseline. In an FTS, the shot (and all other) noise is distributed among the entire spectrum resulting in a noisy baseline, which may mask weak spectral lines during detection, particularly when incorporated in large dynamic range spectra.

Without digressing into a discussion of instrumental noise sources in a LWIR system, FTS systems are generally considered thermal/instrument noise limited (Beer, 1992). This point is what lends credence to the assertion that the SNR behaves as the square root of the integration time because the random nature of thermal noise will eventually "cancel", allowing the steady signal to grow over time (Shepherd, 2002). This of course assumes that the noise truly shows no wavelength dependence and over time, would essentially be equal in magnitude across the bandpass. Shepherd suggests that since the thermal/instrument noise is always present and independent of signal, multiple interferograms can be collected and summed so that the sum total of the individual collection times would have the same effect on SNR as if one had used a long collection time to create a single interferogram. Since platform motion prevents slow scan rates, it would seem that the best collection from a moving platform would occur with the rapid collection of multiple interferograms of the target used together in representing a transformed spectrum.

Further relationships pit the SNR against spectral bandwidth and the number of samples. Spectral bandwidth affects SNR in that instrument noise contributions are integrated over some passband and the greater the passband, the greater the contribution of noise. In turn, the spectral bandwidth

desired directly influences the number of sample points required to construct an interferogram while satisfying the Nyquist criterion. Generally speaking, the shorter wavelength (larger wavenumber) limit of the desired spectral bandpass requires additional samples to be taken at shorter intervals of zero-crossings of the laser beam/metrological system. As mentioned above, the larger the number of samples per interferogram, the greater requirement for throughput in the read-out/processing system, which may limit scan rate.

Finally, an anticipated concentration-path length of the target gas can be used to drive the IFOV, which will determine the sensitivity required to detect the gas under a given set of collection conditions. These collection conditions are of course determined with the type of gas and resolvable/identifiable features already in mind.

A summary table of the dependencies of the major performance parameters for an FTS system is shown in Table 2.2.

Selected Parameter Dependencies		
SNR	**Max OPD**	**Scan Rate**
•IFOV	•Resolution	•IFOV
•Gas Conc-Path	•Scan Rate	•Max OPD
•Instrument Noise	•Mirror Velocity	•SNR
•Bandwidth		•Mirror Vel.
•Scan Rate		•Platform Vel.
•Mirror Velocity		•(Samples)
•Platform Velocity		
•Resolution		
•Max OPD		

TABLE 2.2 SELECTED PARAMETER DEPENDENCIES

2.5.7 Contrasting FTS and DES Performance Parameters

While several misconceptions about the FTS system have been addressed elsewhere in this work, what follows is a brief summary of the major points. First, it is now clear to see that assuming the Nyquist sampling criterion is met; additional sample points cannot improve the resolution of a spectrum and they may in fact, reduce the throughput of a given system. Second, the maximum OPD is not related to the spectral bandwidth. Since all frequencies (those passed by the internal optical bandpass) are passed simultaneously, the interference pattern made by all passed frequencies of light is recorded. Choosing a narrow bandwidth does not allow one to save time on the scan, as is true with a DES system, because all frequencies are already present. One may choose to display only a small subset of these frequencies, but to obtain the desired resolution the maximum mirror travel must be the same, regardless of the ultimate bandwidth and line frequencies desired.

Related to the second point, spectral bandwidth does not influence resolution, as is the case in DES systems. For example, if one wants very fine resolution with a DES for a given detector and focal plane, the dispersion of the element must be increased providing a narrower bandwidth across the spatial extent of a given detector element. Thus the length of the total band across the array is now only a subset of what it once was. Conversely, the only driver of resolution in an FTS is mirror travel distance. As long as one can wait for the travel time, the desired resolution, regardless of the ultimate bandwidth displayed in the transformed spectrum, is limited only by the maximum length of the optical path difference. Of course, both systems require sufficient SNR to detect spectral features at finer resolutions. Finally, the spectral resolution produced by an FTS is not dependent on line frequency as it is with a DES. A one wavenumber resolution setting in an FTS results in one wavenumber resolvable spectral features across the spectral bandpass. While dispersion in grating based spectrometers is often linear over a given narrow bandpass, other DES systems are not (Skoog, Holler, & Nieman, 1998).

2.5.8 Double and Single Sided Interferograms

In an attempt to increase the system scan rate, some time saving techniques can be applied to the interferogram collection method. To complete a transform from an interferogram, all one really needs beyond the sample interval spacing is the location of the point of zero-path difference (ZPD), also known as the center-burst, and the final sample point position. One unfortunate consequence of the discrete sampling of what is really a continuous series is that the exact position and magnitude of the ZPD may not actually be known because the closest sample may fall to one side of the maximum. One method used to help capture the ZPD position/magnitude is to collect what is known as a double-sided interferogram. This technique runs the mirror position out to maximum OPD in both directions, crossing through the ZPD point. Two benefits are realized with a double-sided interferogram. One is that the sensitivity is increased because essentially two collections have occurred. The other is that the ZPD is at least straddled, and can be reconstructed with the knowledge of the sample spacing in the region as well as the neighboring slopes. An obvious objection to this method from the standpoint of remote sensing is the requirement that the source stay static, both spatially and radiometrically, over the time to accomplish essentially two full OPD travels.

Another possibility is to start the mirror travel on just the other side of ZPD and then begin a full OPD travel. This saves some time while preserving accuracy, but locating the position just opposite ZPD is mechanically/computationally complex and consumes time that may be needed to complete additional scans. The process of locating the point of ZPD is called phase correction and is applied to both double and single-sided interferograms before transformation. Once ZPD is located, all path differences reported by sample point become absolute resulting in a valid spectrum upon transformation. While single sided interferograms offer faster scan times, they introduce radiometric error due to the high uncertainty associated with the ZPD location. Figure 2.17 shows a series of single sided interferograms.

2.6 Prior Work

This section will briefly examine a few key sensor technologies and systems that have made great advances in the airborne gas plume detection problem. Example airborne FTS-based systems include the FIRST, ACIS, and LIFTIRS. However, it should be pointed out that to date, none of the groups employing these sensors have published high resolution (< 4 cm^{-1}) spectra collected under real conditions from a moving airborne platform.

Techniques and algorithms for the processing of interferograms to extract gaseous spectral features will also be explored. These techniques obviate the need for spectral transformation, saving time and computational cycles during real-time target detection. Further, it allows for detection and/or recovery of gas target signatures in the event of a corrupted interferogram, all of which were explored in this work.

Finally, the only other modeling approach for a scanning airborne FTS system was accomplished at RIT in 1999 as implemented for DIRSIG. The approach to the DIRSIG model served as a starting point, in some respects, for this work. The DIRSIG FTS model represents a good foundational effort and incorporates multiple facets of instrument modeling.

2.6.1 The TELOPS FIRST FTS

The FIRST (Field-portable Imaging Radiometric Spectrometer Technology) sensor, from the Canadian company TELOPS, has been demonstrated to successfully identify gaseous plumes at video rates via an imaging FTS. The FIRST provides adjustable spectral resolution in the LWIR from 0.25 to 150 cm^{-1} (2.5 nm to 1.5 µm at 10 µm) and has a 320x256 detector array (Farley, Chamberland, & Villemaire, 2007). Figure 2.21 shows a sample composite chemical detection image from the FIRST. The background image is a brightness temperature map of the scene with an overlaid image of regions of the scene identified as gas plumes of both SF_6 and ammonia. The gases are detected using a spectral angle mapper run in parallel with a clutter matched filter. The possible gas species shown at the bottom of the image indicate which filter parameters were used as the target basis for the comparison. The filters are run in both emission and absorption and can identify these differing regions within the plume. This often occurs when a plume is emitted from an exhaust as hot vapor, but then cools as it propagates through the scene (Vallieres, Chamberland, & Farley, 2005).

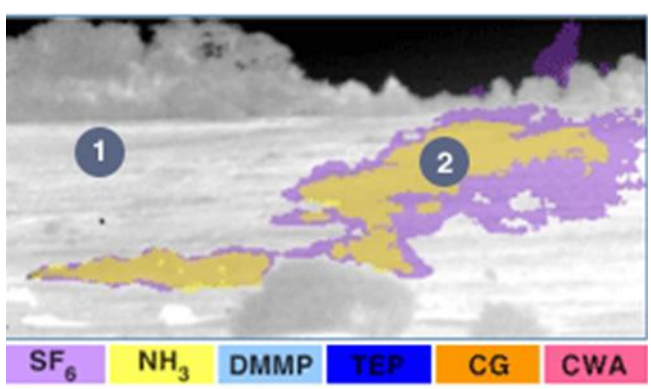

FIGURE 2.21 TELOPS FIRST SENSOR CHEMICAL IDENTIFICATION IMAGE

The collection parameters were not provided for this image, but the system was static and mounted on a tripod. While ammonia is a more representative gas of those of interest in remote sensing as compared to sulfur hexafluoride, it does not pose a particularly difficult target. Figure 2.22 is a LWIR spectrum generated by the model used in this work at 4 cm^{-1} (40 nm at 10 µm) resolution and illustrates the extensive coverage of the ro-vibrational feature through the spectral window. Interestingly, the line position between the major spectral features of NH_3, is where the *only* LWIR spectral feature for SF_6 is located – a substantial Q branch flanked by minor P and R branches from a deformation. Thus, essentially no species discrimination is really required to discern the identity of each in this experiment – simple narrowband monitoring of the spectrum at 10.6 µm for the SF_6, and on either side for the NH_3, would accomplish the detection/identification task.

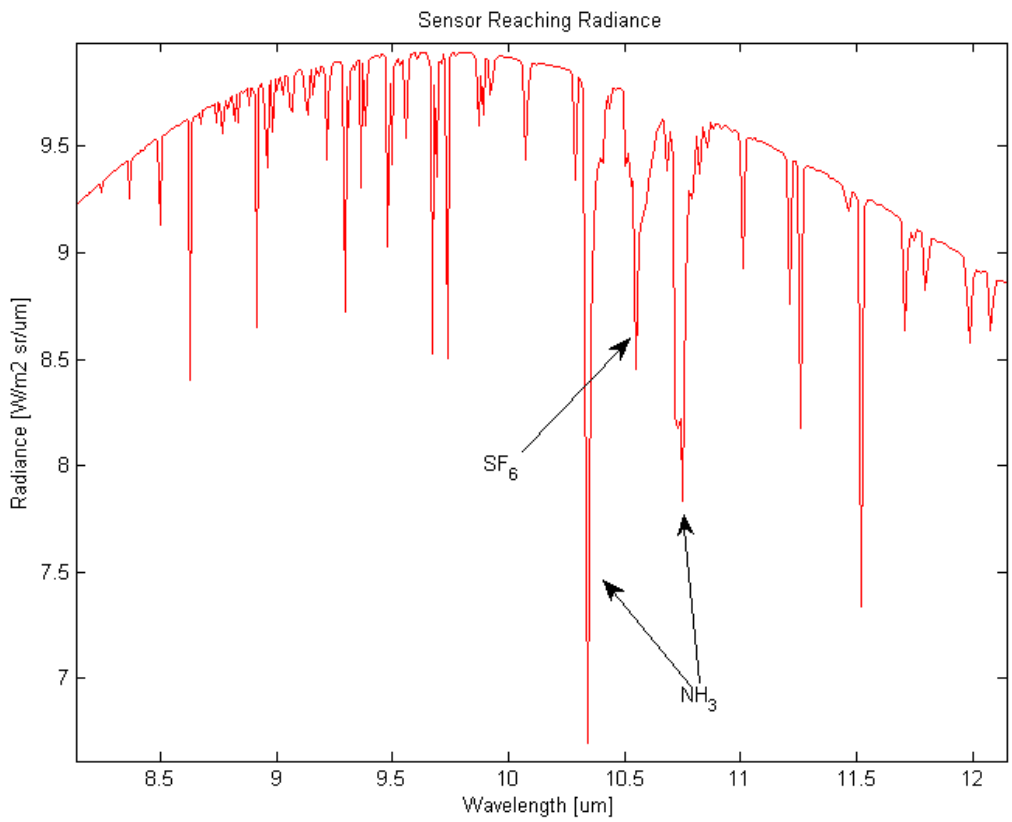

FIGURE 2.22 ANNOTATED LWIR SPECTRUM OF NH_3 AND SF_6 (4 CM^{-1} RES.)

The FIRST was recently adapted to collect data as a staring sensor mounted on an airborne platform (Allard, Chamberland, Farley, & Villemaire, 2008). The design team, knowing the aforementioned potential hazards of collecting an interferogram from a moving platform, took several steps to minimize these ill-effects (Figure 2.23). First, the system was mounted on a stabilized-dampening mount that actively compensated for aircraft yaw and high frequency vibration. The primary mirror used for target monitoring was made steerable via a voice coil mechanism that updated its position at 100 Hz. The mirror had 2 DoF, which allowed not only the ability to continually stare at the target through aircraft translation, but also the ability to cancel aircraft motion along both the pitch and

roll axes. The following diagrams illustrate the FIRST system mounting on the left and the three DoF of an aircraft on the right.

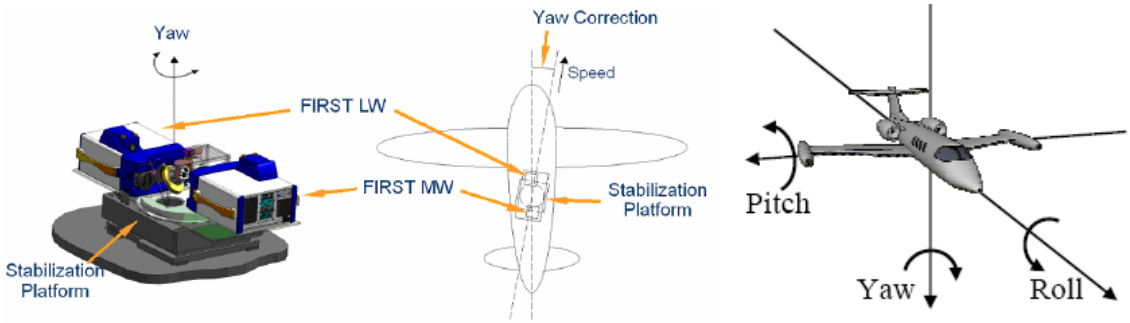

FIGURE 2.23 FIRST MOUNTING (LEFT) AND AIRCRAFT DOF ILLUSTRATION (RIGHT)

In order to determine the acquisition parameters for an airborne collect, several tradeoffs between instrument settings needed to be made so that the system was optimized. The performance tradeoffs, specifically between spectral resolution, temporal resolution and SNR, addressed in the TELOPS paper essentially made it the closest work heretofore to the designated goals of this effort.

Figure 2.24 shows a notional relationship between the spectral resolution, as measured by the FWHM of the instrument line shape as discussed previously, and the rate of production of a datacube.

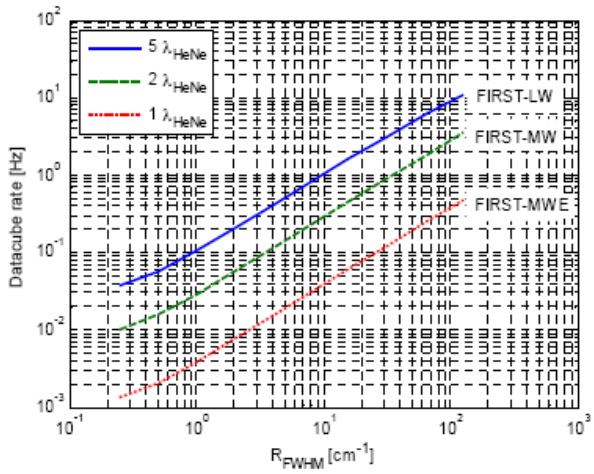

FIGURE 2.24 DATACUBE RATE AS A FUNCTION OF RESOLUTION FOR THE FIRST SENSOR

Figure 2.24 is a log-log plot where the top trace indicates the performance of the LWIR sensor. The difference in the traces indicate that the sample spacing is a different integer multiple of HeNe laser wavelengths as discussed in the FTS section. It is unclear why the integer *five* was chosen for the wavelength multiple as both the work shown in the FTS section above and in the same paper from Allard suggest that a multiple of six wavelengths is sufficient to sample a LWIR interferogram. In any case, the plot illustrates the inverse relationship between resolution and the maximum OPD

required to produce a sufficient interferogram. Clearly, as resolution becomes finer, the maximum OPD required grows, reducing the datacube production rate due to the increased time required to complete a full OPD sweep at a given mirror velocity. The nearly linear relationship on this log-log plot shows that a 4 cm^{-1} resolution would restrict datacube production to approximately one cube every 2.5 seconds. While this represents an extraordinarily fast rate as compared to non-FTS systems such as the AIRIS, it may still not be fast enough as will be shown later.

An additional performance restriction placed on the airborne FIRST system comes from the voice coil actuated primary mirror. The voice coil actuator was not selected solely to point the mirror for target tracking, but rather due to its extremely fast rate of update that allows for aircraft motion cancellation due to erratic transient events like turbulence, buffeting, and wind variations. Unfortunately, a voice coil only allows for total travel distances on the order of millimeters, thus the mirror movement is restricted to +/- one degree. This narrow pointing range severely restricts the staring capability of the system at a given aircraft speed, particularly at low altitude. As the aircraft passes overhead, if one is trying to stare at the same patch of ground, the angular rate of travel for the mirror is faster at a lower altitude and would reach its limit more quickly restricting total collection times. Figure 2.25 taken from Allard (2008), illustrates this concept.

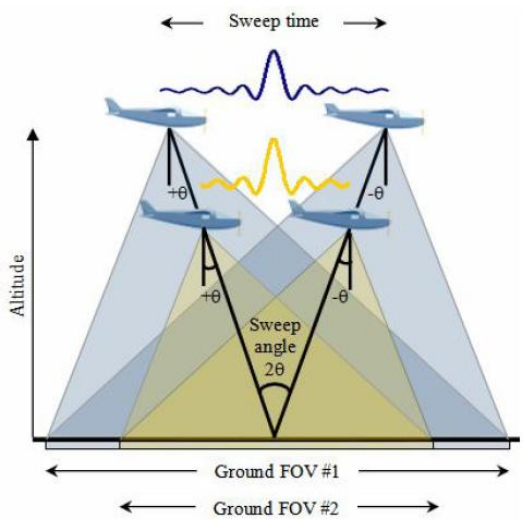

FIGURE 2.25 MIRROR SWEEP TIME AS A FUNCTION OF ALTITUDE FOR THE FIRST SENSOR

Figure 2.25 shows that the physical limits of the primary mirror travel with a fixed field of view (FOV) result in short interferograms and therefore place a physical restriction on attainable resolution at a given altitude and airspeed. The plots below, shown in Figure 2.26 demonstrate the affect on resolution at various altitudes for an airspeed of 50 and 70 m/s (112 and 157 mph, respectively).

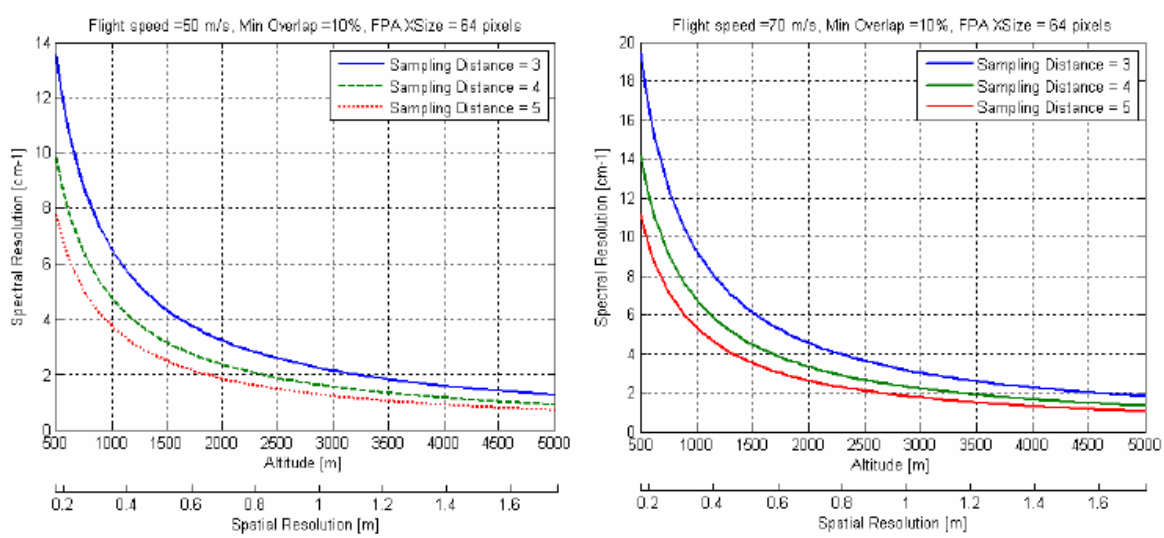

FIGURE 2.26 COMPARISON OF AIRSPEED AS A FUNCTION OF ALTITUDE AND RESOLUTION FOR THE FIRST SENSOR, 50 M/S (LEFT), 70 M/S (RIGHT)

The plots illustrate that a 40% increase in aircraft speed results in a substantial loss of spectral resolution at low altitude, where again, the sampling distance of five multiples of a He/Ne laser wavelength correlates to the anticipated LWIR performance. The collection conditions indicate a 10% overlap of each 64x64 spatial scene, where reduced overlap can marginally improve spectral resolution through slightly increased mirror travel time.

As suggested earlier, if an aircraft is restricted to a one kilometer altitude to minimize atmospheric effects, the FIRST FTS should be able collect spectra at a spectral resolution of 4 cm^{-1} at a rate of 0.4 datacubes per second of the size 64x64x397. Unfortunately, all of the figures presented in this section are theoretical performance estimates, as there are no results published from the system having been flown to date. However, it is posited that there may be some fundamental flaws in the adaptation of the FIRST sensor to an airborne platform and these will be addressed in the approach chapter that follows this section.

2.6.2 ACIS: AIRIS and TurboFT

Another solution to the location and identification of gaseous targets is the Advanced Chemical Imaging System (ACIS) tested by the US Army (D'Amico, Emge, & Marinelli, 2004). The system contains multiple sensors mounted in a gyro-stabilized optical platform on a helicopter. The platform allows for stabilized pointing and tracking while the system collects data. The high spatial resolution component of the system is the AIRIS (Adaptive Infrared Imaging Spectroradiometer), with a Fabry-Perot etalon, used to build up to a 64x64 (with an undisclosed number of spectral bands at an undisclosed resolution) spatial image over a reported 15 seconds per spectral band. In the prototype system, Gittins reported that the maximum observed resolution of AIRIS was approximately 0.08 µm (8 cm^{-1}) and that the spectral range was from 10.0 to 11.5 µm (Gittins & Marinelli, 1998). The prototype system was reported to provide a 48x48 spatial image over 14 bands

in approximately two minutes at nine seconds per band. These numbers seem to support the performance values given for the operational version of AIRIS.

Another component of the ACIS, the TurboFT FTS, is used to provide interferograms at the rate of 100 Hz for an 8x2 spatial sample of the area at 8 cm^{-1} resolution. The concept employed here is to use the TurboFT for wide-area spatial reconnaissance where rapid scanning of a sector can be accomplished at poor spatial resolution from a moving platform. Later, a spectral image can be created over an area of interest with the AIRIS sensor from a static (hovering) platform.

The TurboFT is the closest system to accomplishing the goals set forth in this work through the use of an FTS. The design is revolutionary when compared to other operational remote sensing FTS systems, but is plagued by ghosting at high resolution. The design is still based on interferometry, but uses a spinning refractive optic as the method of introducing the phase difference into the recombined optical path. A rough sketch of the optical path is shown in Figure 2.27 from Jensen, et. al. (Jensen, Ifarraguerri, & Loerop, 2002).

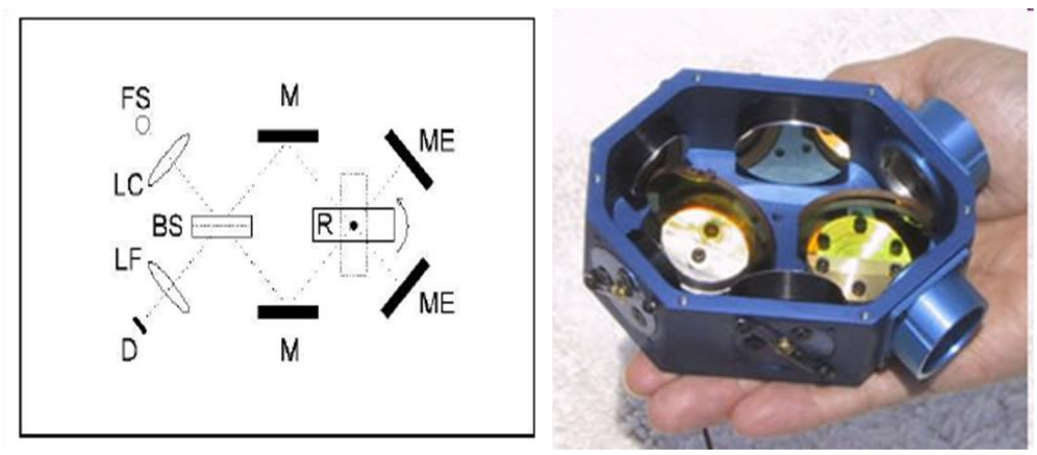

FIGURE 2.27 OPTICAL PATH DIAGRAM OF TURBOFT FTS (LEFT) AND ACTUAL TURBOFTS (RIGHT)

The diagram shows light entering through the aperture in the top left, passing through a beamsplitter where half of the magnitude of the original beam reflects off of each mirror in the middle, passes through the refractive optic at a different entry angle, reflecting off of the end mirror, recombining at the beam splitter and exiting to the detector in the bottom left. The authors claim the system can be switched from midwave to LWIR by changing the refractive optic out with one with a different index of refraction. There are many major advantages of this system as compared to Michelson-based FTS designs.

A major advantage of this system is the speed in which interferograms are collected. The speed of collection is due to the fact that each revolution of the mirror provides four completed interferograms as the maximum OPD is swept out four times per revolution. Or said another way, there are four ZPD points during each rotation. Figure 2.28, also from Jensen, illustrates this point.

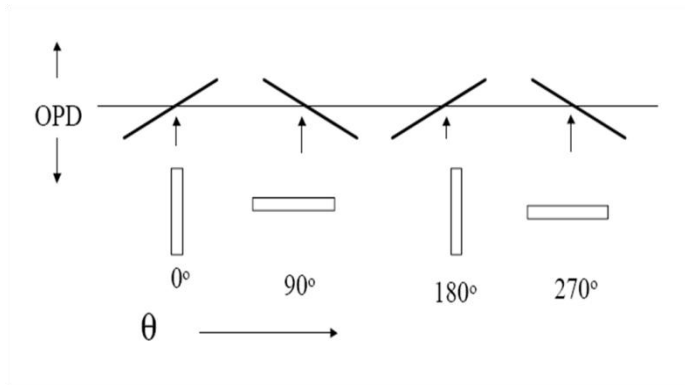

FIGURE 2.28 ROTOR POSITION (Θ) AS A FUNCTION OF OPD

Another major advantage is the simplicity in the design. Sample triggering is accomplished via shaft encoding instead of a laser-gated system. Thus, a reader monitors the shaft position and triggers an interferogram sample point tied directly to the optic position, which provides a known OPD. The ability to keep the rotor at constant velocity is aided by the large rotational inertia of the spinning optic, ensuring a reliable measurement. These advantages combine to yield a system that can be made to be very compact, robust, and require little power.

However, there is a non-linear sampling problem that exists at high resolution. At large OPD the non-linear relationship between OPD and angular velocity becomes significant, resulting in unequal spacing of samples in the interferogram. As discussed in the FTS section, high resolution is obtained via large OPD and the larger the OPD, the more dramatic the non-linearity becomes in the TurboFT. Figure 2.29 from Jensen demonstrates the ghosting effect introduced from a non-linear error in sampling the interferogram.

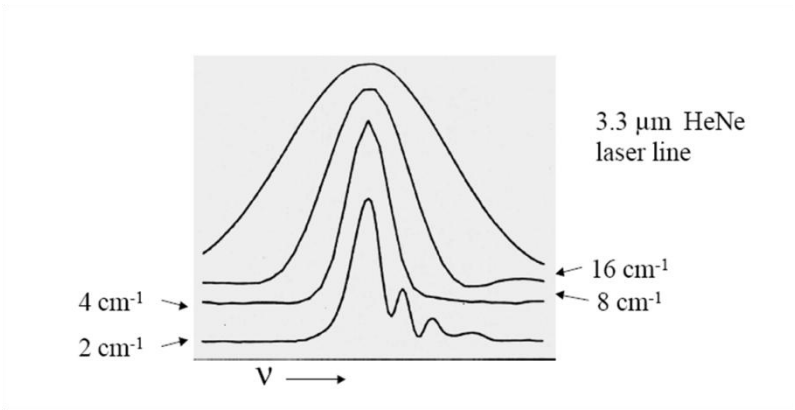

FIGURE 2.29 GHOSTING RESULTING FROM NON-UNIFORM SAMPLING OF THE INTERFEROGRAM AT HIGH RESOLUTION

Figure 2.29 shows the resulting transformed narrowband spectrum of a HeNe laser emission line at 3.3 µm over multiple resolutions. A doubling of the resolution requires a doubling of the maximum OPD traveled for the light path and the example shows that reasonable results are achieved through 4 cm^{-1} resolution. However, as 2cm^{-1} is approached the non-linearity becomes significant and the non-uniform spacing of the interferogram samples results in significant ghosting on the high energy

side of the spectral feature. The author reports that the TurboFT is only expected to operate at 8 cm^{-1} resolution and is not concerned with this design flaw.

Further, the operational rate of the TurboFT is intended to reach 360 Hz/pixel (Jensen, Ifarraguerri, & Loerop, 2002). At this rate, there is concern about the response and sensitivity characteristics of their HgCdTe based detector. Assuming the system can process 5,760 interferograms per second for a 16 pixel detector, this rate shows great promise for nullifying the effects of plume propagation and platform motion in using an FTS to detect gaseous spectral features. This concept will be explored in the approach section that follows.

Clearly, this system is a major step forward over earlier attempts in the airborne identification of gas plumes. Unfortunately, at the time of this writing, no results have been released from an airborne data collect using the Turbo FT. The AIRIS system was flown over Oklahoma City in July of 2003. The system collected a spectral image of the scene first, then gas was released and the image was collected again. The idea behind the experiment was to show a proof-of-concept for the airborne detection of a gas plume in an urban setting. Figure 2.30 shows the urban setting for the test while the image on the right shows the AIRIS detection of SF_6 gas. The white circles indicate the site of the release on the left and the region of detection of the gas for AIRIS on the right.

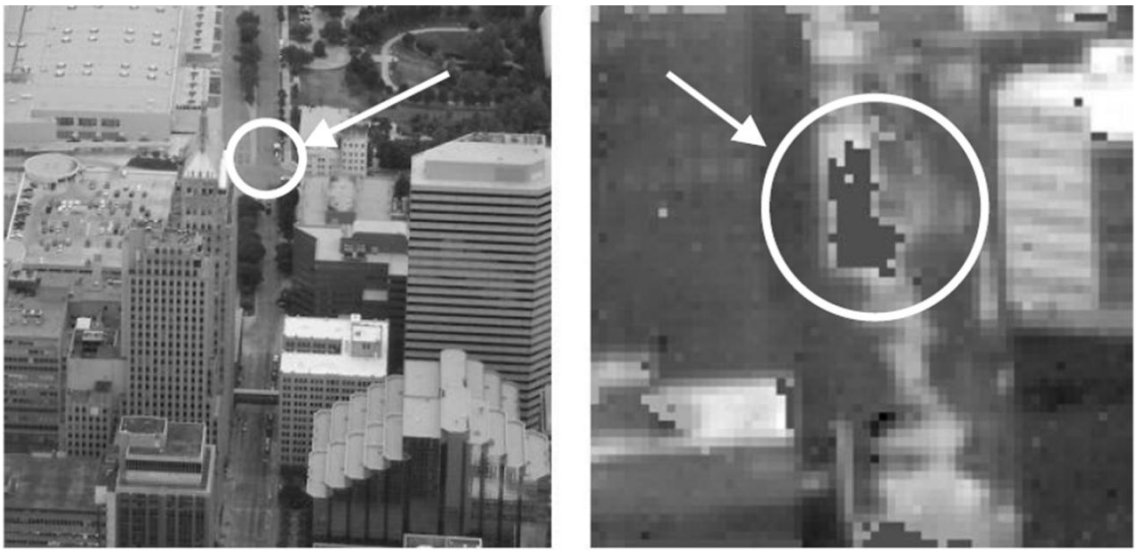

FIGURE 2.30 URBAN SCENE FOR TEST (LEFT) AND AIRIS DETECTION OF SF_6 (RIGHT)

While this test represents a baby-step in the path towards successful operational deployment of such a sensor, the experimental setup provided for a comparatively easy detection scenario. The collection of the scene with and without the gas plume allowed for simple background subtraction to be used to isolate any spectral differences, nullify background material and thermal variations, and eliminate much of the atmospheric influence. The use of a helicopter as the flying testbed allowed the system to remain essentially static during integration to minimize spatial and spectral smearing. Further, the test gas was sulfur hexafluoride, SF_6. This gas is often used in remote sensing experiments for its benign effect on the environment allowing for ease of release and the

unmistakable, isolated, intense spectral feature at 10.55 µm. While there are clearly advantages to using this gas, it is hardly representative of any real gas target of interest one may want to detect under operational conditions. A sample LWIR spectrum of SF_6 generated by the model is shown on the left side of Figure 2.31. A spatial detection example from the paper is provided on the right and represents a background subtracted spatial image from the 10.6 µm spectral plane of the hypercube.

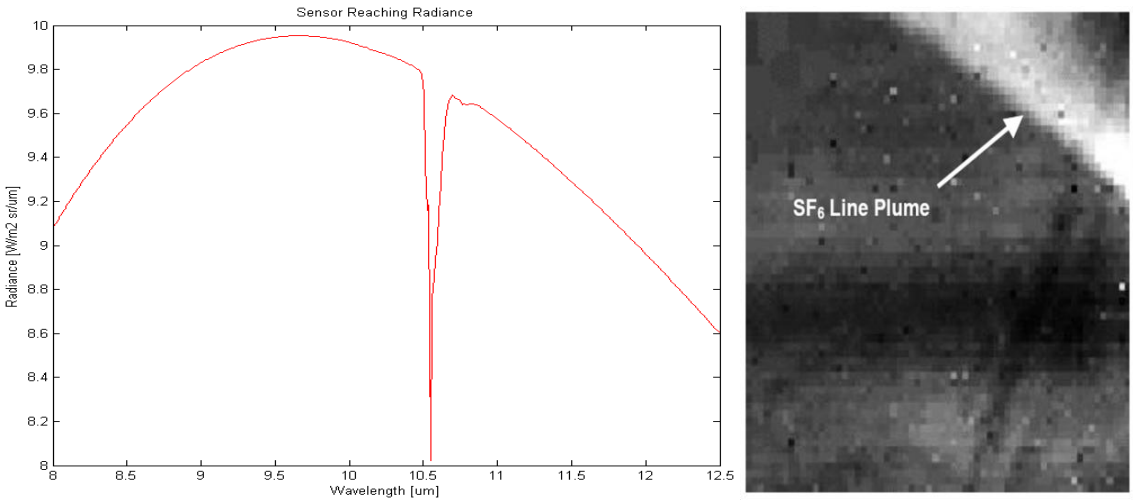

FIGURE 2.31 LWIR SPECTRUM OF SF_6 AT 10 NM RES. (LEFT) AND 10.6 µM DETECTION PLANE (RIGHT)

The comparison of the spectrum above and the detection plane shows that SF_6, at even modest concentration-path length products, would dominate a background subtracted scene. In fact, Gittins reports that the prototype was sensitive to SF_6 down to 0.6 ppm-m at a thermal contrast of 5K!

While the paper claims that the TurboFT was also successful in detecting SF_6 at the urban test site, they state that the data were not available at the time of release of the paper (D'Amico, Emge, & Marinelli, 2004). No further paper has been produced elaborating on the airborne test results from the TurboFT. Still, the TurboFT is a very promising FTS technology and could possibly transform the state of the airborne gas plume detection problem.

2.6.3 Interferogrametric Processing Techniques

In a unique application of signal processing theory to the gas detection problem, a field of work has evolved in developing techniques to isolate and identify the unique interferogrametric signature of various gas species in remotely sensed data. The premise of this field relies on basic assumptions about frequency domain principles. Recalling that an interferogram is comprised of the optical interference pattern of all wavelengths of light entering the interferometer, the magnitude of the interference is related to the phase difference of each wavelength and the total radiance collected for that wavelength. These products are optically summed for each wavelength resulting in a series of multiplexed values at each OPD. If one were to apply a Fourier analysis to these frequencies, it can be said that the observed magnitudes at each OPD are due to the transform of the summation of sinusoidal curves at frequencies representing the spectral feature details at each collected wavelength. Thus, slowly varying spectral features will have an underlying frequency that will

manifest itself as a large spike in a position located near the origin of the interferogram. Conversely, spectral features with abrupt changes should contain high frequencies that would manifest themselves farther downfield in the interferogram.

In applying these principles to the gas detection problem, some assumptions must be made. The first assumption is that even relatively complex background spectral signatures should only vary spectrally at a slow rate and comprise the majority of the collected signal. This would suggest that an interferogram would contain these spectral features in low frequency regions near the center-burst. Further, any quickly varying spectral features, perhaps exceeding the instrument line shape, are assumed to be noise from varying sources and should be located at the ends of the interferogram in the high frequency regions.

If these assumptions are true, then not only should one be able to accomplish basic background removal via high-pass filtering of the interferogram, but one should also be able to remove instrument noise as well with a low-pass filter. What should be left over is the frequency content that is largely representative of the chemical species detected in the scene. In fact, it is suggested in work accomplished by Kroutil, et. al. in 2002 with the US Army Edgewood facility, that gaseous species can indeed be identified by their interferogram alone, resulting in both a reduction in computational complexity and time savings in post processing by obviating the need for a spectral transform (Kroutil, Knapp, & Small, 2002).

Work accomplished by Thomas and Kroutil has shown that interferograms collected from an airborne collect of an industrial stack can be digitally filtered to reveal a region in the interferogram that is sensitive to the spectral features due to the gas plume (Thomas, Lewis, Kroutil, & Combs, 2002). Figure 2.32 shows an isolated region in a double-sided interferogram that reveals sensitivity in line shape due to the presence of methanol.

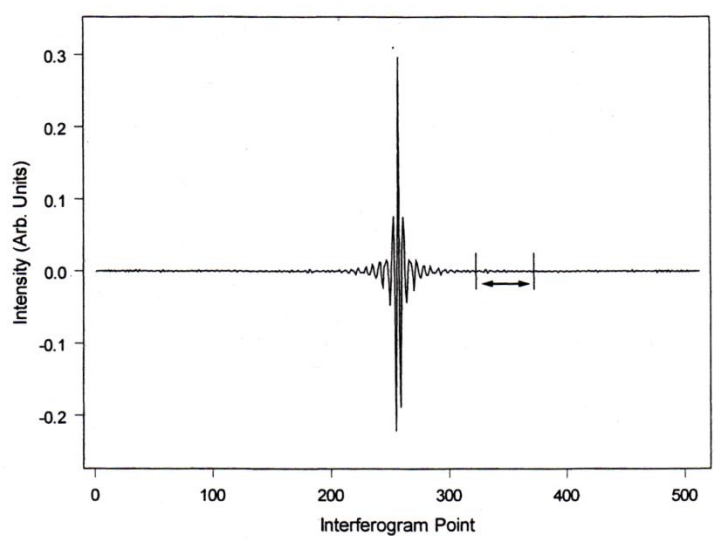

FIGURE 2.32 RAW INTERFEROGRAM SHOWING REGION OF SENSITIVITY TO PRESENCE OF METHANOL GAS

Figure 2.33 shows four different interferograms from four different spatial samples of the scene, where the top two traces contain methanol, and the lower two do not. The interferograms have essentially been bandpass filtered (modulated with a rectangular function) to only show the region that yields a response to the presence of methanol vapor.

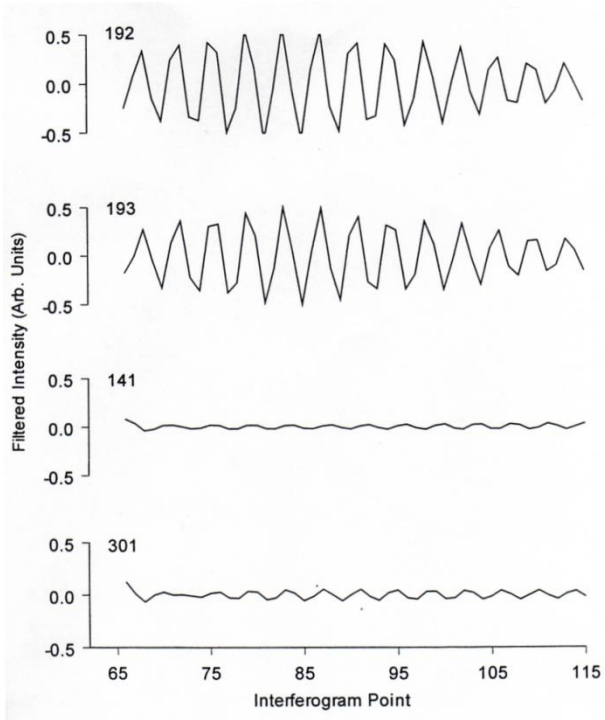

FIGURE 2.33 INTERFEROGRAM COMPARISON OF REGION OF SENSITIVITY TO PRESENCE OF METHANOL GAS

As can be seen in Figure 2.33, the presence of methanol in this case yields a highly discernable and highly repeatable response. From this point, it is reasonable to apply identification aids and many have been proposed such as: a spectral matched filter, linear discriminants, polynomial fits, integration, squared magnitude (power), or principal component analysis (PCA) to quantify the match to an interferogrametric response from the "truth" spectrum of the gas (Kroutil, Knapp, & Small, 2002). Figure 2.34 shows the spectral transforms of the two interferograms, as the solid traces, that contained the target methanol vapor absorption features. Upon transformation to the spectral domain, the spectra taken from neighboring collects without the indication of methanol were used in background subtraction. The dashed spectrum represents a laboratory generated fit of the data that was comprised of methanol vapor at 763 ppm-m, 16 cm^{-1} (0.16 µm at 10 µm) resolution, and a background-gas thermal contrast of 25K. The spectrum is flanked by an equivalent spectrum taken from NIST at 4 cm^{-1} (0.04 µm at 10 µm) resolution for comparison of the underlying ro-vibrational features giving rise to the spectral detection.

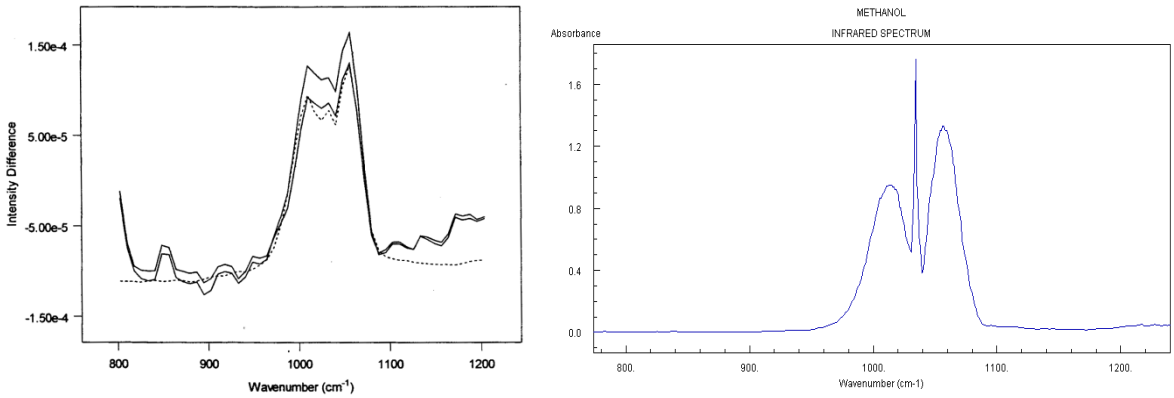

FIGURE 2.34 TRANSFORMED SPECTRA OF METHANOL DETECT WITH LABORATORY FIT (LEFT) AND NIST COMPARISON (RIGHT)

It is clear from the NIST spectrum that the underlying ro-vibrational features belong to the P-Q-R branches and can be visualized in the collected data.

Since Kroutil's initial work, he has since developed software that accomplishes gas detection via interferogrametric analysis and has coded it for implementation in ENVI/IDL (Kroutil, Lewis, Miller, Shen, & Small, 2008). The steps in the software are to first correct any phase errors introduced to the interferogram through the instrument line shape of the spectrometer, thus removing radiometric error in later steps. A segment of the interferogram sensitive to the chemical species is then selected via an iterative statistical analysis that considers the best signal from the species of interest as compared to the background and any other interfering spectral signatures. Once this segment is located, it is processed with a digital filter to remove all but the signature of the chemical species. The finite impulse response matrix (FIRM) filter was chosen for this step as their previous work has shown it to be robust against the large dynamic range of signals in the interferogram. Finally, to determine species identification, a pattern recognition approach is implemented. In this step Kroutil has used PCA, partial least squares, piecewise linear discriminant analysis (PLDA), and a non-linear neural network approach. He recommends PLDA where a non-linear discriminant surface is generated from several linear discriminants that comprise a separation surface of the interferogram. Thus, the region enclosed by the discriminant surfaces contain all interferograms that likely contain the species.

While there is certainly merit to this approach to gas detection, from an operational standpoint, there are many things that can go wrong. When viewing gases at comparatively poor resolution, the possibility for false alarms over a narrow bandpass of the interferogram is possible – particularly since many spectral features begin to look alike at poor resolution and would show up in essentially the same place. Next, at low resolution, there are few regions in the interferogram to choose from to select your bandpass. With all of the proverbial eggs in one basket, if the particular interferogram region was corrupted due to jitter or mirror error, all evidence of the particular spectral feature of the gas would be lost. For further clarity on the subject, reference Figure 2.17 for a comparison of the unique spectral feature locations in the interferogram at varying resolutions.

2.6.4 DIRSIG Airborne FTS Model

The goal of this work is to develop an end-to-end airborne FTS model for use in characterizing system performance under various environmental and instrument conditions. While various FTS instrument models exist, they are not designed with the mobility, flexibility or end-to-end requirement necessary to implement in this work. It turns out the closet modeling effort to this work was accomplished at RIT over a decade ago.

In 1999, Francois Alain completed work on an instrument model of an FTS for DIRSIG as his Master's research (Alain, 1999). While there is no mention of an FTS instrument model/collection platform in the DIRSIG User's manual, it has been verified as a part of the modeling package (Digital Imaging and Remote Sensing Laboratory, 2006).

The instrument model is broken into several modules consisting of the user interface, Michelson-type interferometer, detector, FFT, image cube rendering, and noise module.

When the instrument model is employed, data are collected as a framing array (it was actually implemented as a pushbroom designed to behave like a framing array to save rewriting code, but had apparently been validated). The inputs to the model include sensor altitude, speed, integration time per sample, and the collection time for the entire image. These parameters are then used in the determination of the view angles and sensor reaching radiance through DIRSIG. There is even consideration for platform pitch, roll, and yaw effects as well. The output of the model is a spectral image cube for each frame. The model allows for multiple view angles per image cube, thus simulating the FIRST-type airborne sensor where the interferogram is sampled as the platform translates across the scene, maintaining a boresight on the region of interest. Unfortunately, DIRSIG must be run to generate the sensor reaching radiance for each view angle, placing a substantial computational requirement on the simulation. Alain suggests generating a few primary view angle radiance models and interpolating between them to save time. The number of view angles to be generated can be specified by the user in the sensor flight profile.

Using a given focal length, pixel size, number of pixels, and frame spacing, the DIRSIG sensor submodel models the optical effects to determine the FOV and ground sample points. The interferometer can be adjusted for mirror scan speed and precision in path length, detector integration time, total scan time, interferogram oversampling, and various off-axis effects.

The detector module allows the user to set quantization, spectral response, clipping, apodization, gain, and bias. The BAND section of the DIRSIG configuration file is said to allow specification of the number and size of the detector elements in the FPA as well as to stipulate spatial oversampling. The FFT module takes an interferogram datacube and transforms it into a spectral data cube. It is not clear if this step can be bypassed for one to have direct access to the interferograms. Alain suggested that they would not be accessible in that once the multiple data cubes from various view angles are combined to make the overall output image cube for a spatial position, the data cubes are deleted. Further, Alain argued against providing the data in an interferogram form because of size constraints. Alain showed several examples of interferograms and their sensitivity to various

artifacts used to challenge the model. It is unclear if Alain used a “transform of a transform” to get an interferogram that is identical to the input interferogram for the FFT module in his study, or somehow managed to retrieve the interferograms directly.

Alain suggests that a triangular apodization function, as discussed in this work, was implemented in the code and was done so in a modular form such that other apodization functions may be used. To handle the simulation of noise a Gaussian distribution is added to each interferogram produced by the instrument. A setup file allows a noise mean and standard deviation to be applied to each pixel in the detector array.

Each detector element can be set for ideal sampling over the OPD, or a multiple-time oversample followed with averaging yielding the single sample point. It is unclear from the results of his study how this positively affects the resulting spectrum. He states that the modulation of the Sinc function is minimized via oversampling, but turns to this method rather than apodize the interferogram. In any case, oversampling appears to be an adjustable parameter in the model.

In order to simulate a framing array and allow multiple interferograms to contain a subset of the same ground pixels, the user can adjust the detector ground IFOV as described above. The approach here is to use subpixel sampling which is also user selectable. For example, a super pixel consisting of the IFOV is divided in a 9x9 array of subpixels such that the simulation combines the output of the subpixels. This allows for a marked decrease in jitter (modulation of the true spectrum due to abrupt changes in spectral radiance) due to the unnatural effect of discretely sampling a discretely rendered scene. The blending of a larger spatial area allows a more continuous transition between spectral radiance properties of a given scene, reducing artifacts in the interferogram due to discrete shifts in radiance at each position.

The spectral response given to the instrument for the LWIR was a Gaussian, centered at approximately 1080 cm^{-1} with an approximately 350 cm^{-1} FWHM. No justification or further elaboration was given for this choice.

Finally, the error in OPD, which is related to the uncertainty in the mirror position, was simulated by using a fraction of the sampling interval as either an addition or subtraction from the ideal OPD position. This error fraction is generated from a selection of randomly generated values from a Gaussian distribution with a user selectable standard deviation, which is then smoothed to prevent abrupt changes in the values.

In all, Alain’s FTS model appears quite robust in approach and serves as a good comparison to the model developed for generating data for this work. At the time of this writing, the functionality of the FTS model as implemented in the current version of DIRSIG has not been verified.

3 Approach

The purpose of the airborne FTS model is to simulate the ideal performance of an FTS instrument aboard a moving platform conducting a spectral search for gas plumes that require high spectral resolution to identify. The model allows for multiple environmental, scene-based, and instrument properties to be varied to evaluate detection performance under a given set of conditions. Multiple studies of the impact of individual instrument parameters, or the trade space for performance under two variables, were conducted in preparation for the scene-based effort.

This section details the methods used in the model development, gives rationale for why various trades or assumptions were made, gives an overview of the model validation process, the Monte Carlo simulations, and finally chronicles the development of the scene-based scenarios to evaluate the performance parameter envelope of trades made to simulate an operational employment of the instrument.

3.1 Airborne FTS Model Development

The airborne FTS model is an assembly of various sub modules that support the end-to-end remote sensing model and consist of the source input, error sources, output, and detection metrics. Inputs include radiance contributions from the surface materials, the atmosphere, and the gas plume, if present. A different module in the code corresponds to each of these inputs. The user can specify multiple parameters for each allowing substantial customization. The instrument model accepts these inputs and converts them to an interferogram. During the interferogram conversion process the input can be subjected to uncertainty in the moving mirror position, jitter uncertainty in the primary mirror, interferogram filtering, interferogram detection, and apodization. The interferogram is then transformed back to the spectral domain and is subjected to multiple detection quantification algorithms including Spectral Angle Mapper, Spectral Matched Filter, spectral depth, and various combinations therein. Detection metrics will be covered later.

The model can be implemented as viewing a single pixel and yielding a single spectrum, or can be 'flown' over a scene to yield a complete hypercube along with corresponding cubes of detection metrics. When the model is flown, it can form a single image of the scene, or scan along the flight path, integrating interferogram portions from multiple ground samples. A scene can be designed and built manually or imported as sensor reaching radiance cubes from other sources.

Figure 3.2 illustrates the data flow in the model through the various model components.

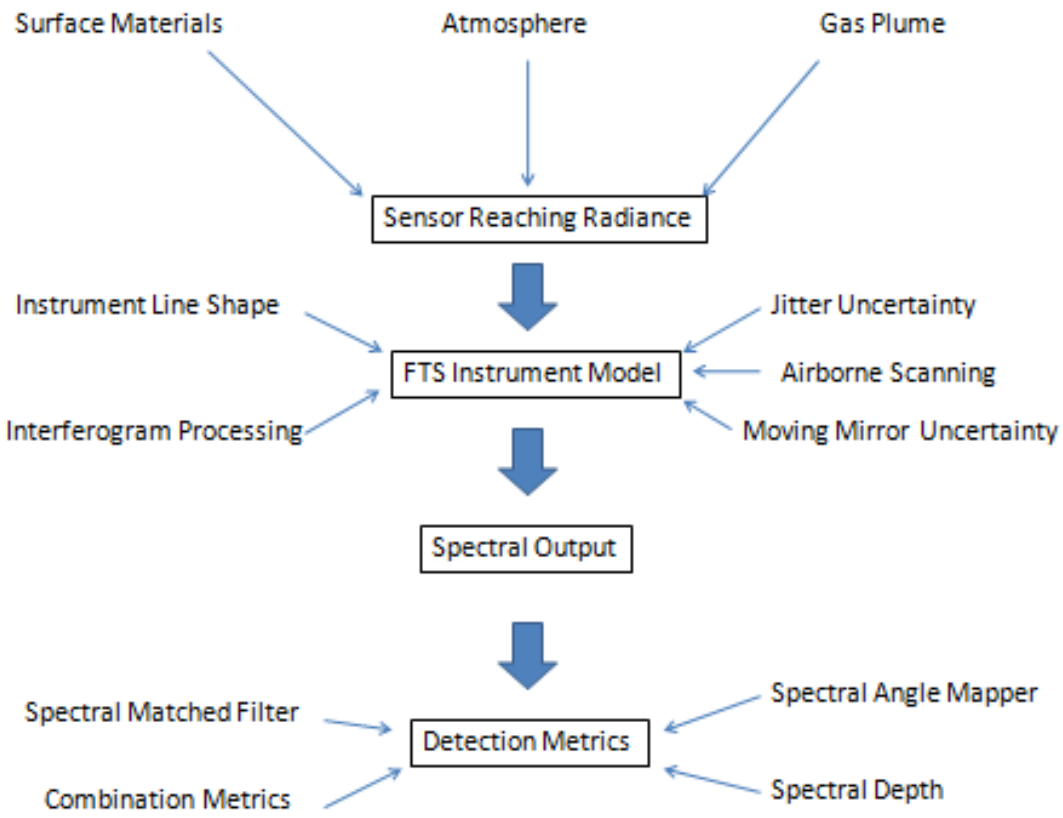

FIGURE 3.1 AIRBORNE FTS MODEL OVERVIEW

Further breaking down the model components, Figure 3.2 illustrates the various inputs from the first stage of data flow.

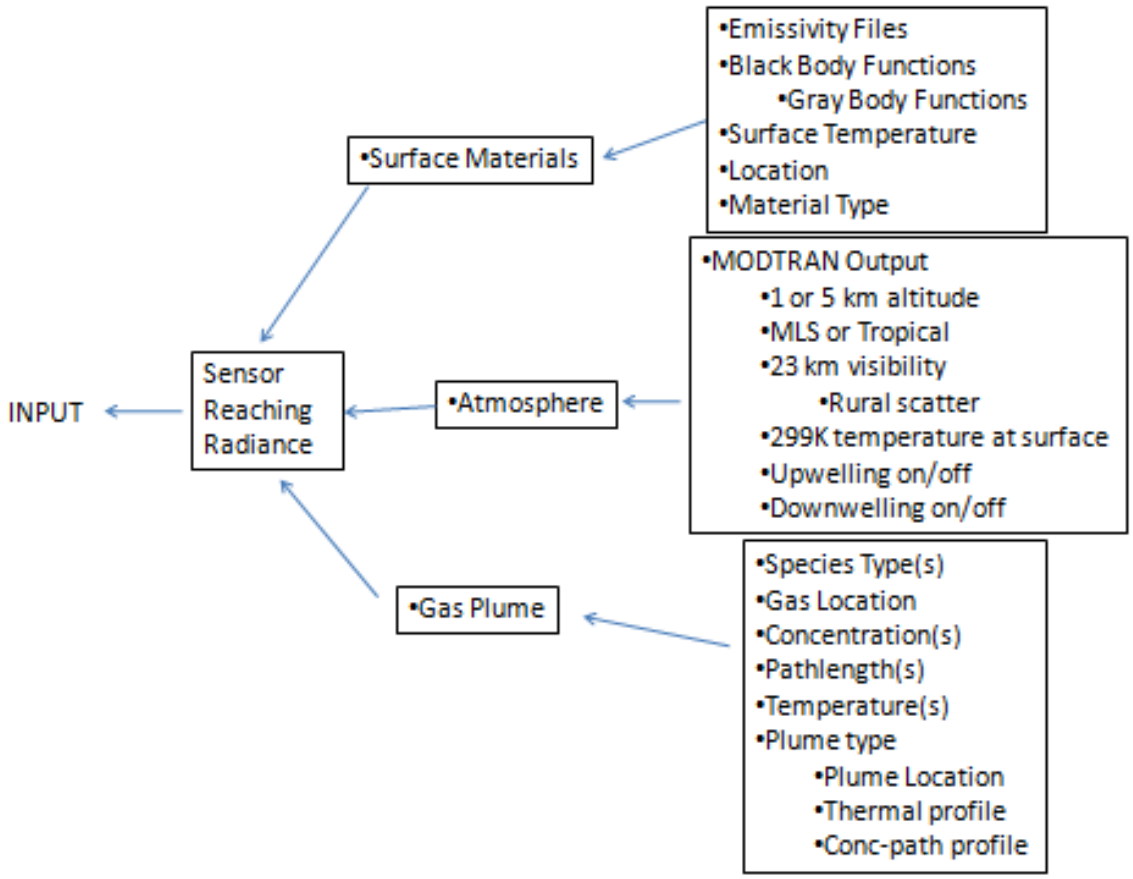

FIGURE 3.2 MODEL INPUT DIAGRAM

Figure 3.3 shows the FTS Instrument Model breakdown.

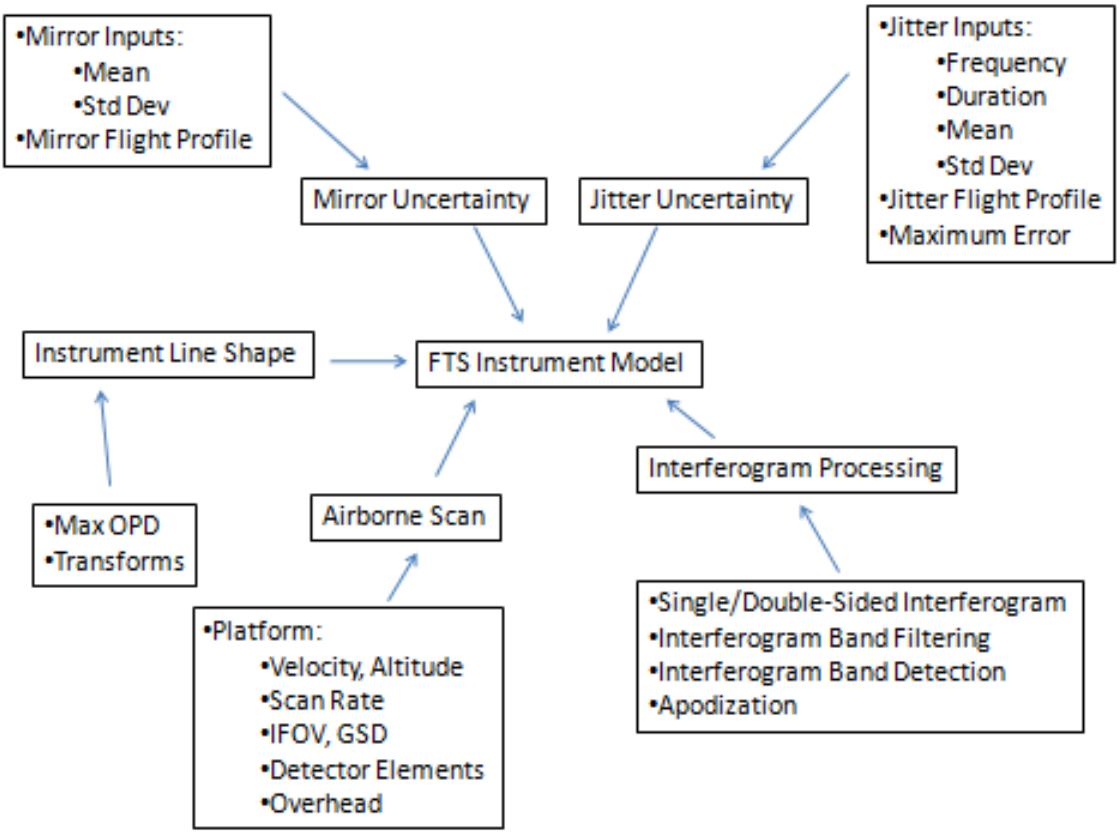

FIGURE 3.3 FTS INSTRUMENT MODULE DIAGRAM

The spectral output module is straightforward and does not need a diagram, while Figure 3.4 shows the breakdown of the detection metrics module.

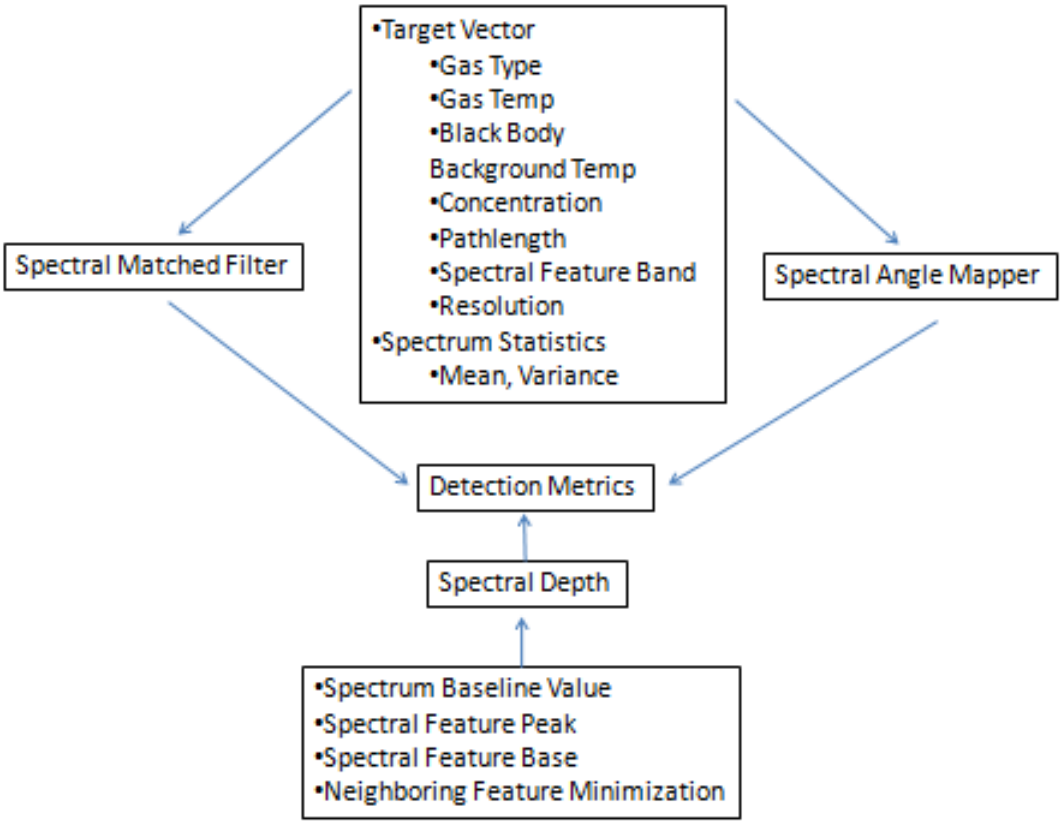

FIGURE 3.4 DETECTION METRIC MODULE DIAGRAM

What follows in the remainder of section 3.1 is a breakdown of each of the modules covering their specific inputs, outputs, and design.

It should be noted that ideal instrument conditions are modeled unless stated otherwise. This means that the model has no optical aberration affects from stray light, off-axis light, chromatic aberration from the beam splitter or compensator, misalignment of optics and/or mirrors, detector response and noise, instrument self-emission, quantization, optical registration on the focal plane, interferogram phase and alignment/sampling, and signal clipping. In addition, the mirror velocity and sampling rate of the detector are always considered sufficient for the assigned scan rate. Rather, the approach to this model was to explore the interferogram formation process as it relates to the detection of gas plumes in complex scenes. Only factors having major impact to the spectral output without direct mitigation were considered. The other aforementioned sources of error can be modeled in future work for even greater instrument fidelity.

3.1.1 FTS Instrument Model

The following sections cover each of the components of the FTS instrument model.

Instrument Line Shape Model

The Airbone FTS instrument model was built around the Fourier transform concept, as one would expect. The transform engine was based on an FTS instrument model given by Skoog, et. al. and

Gross. et. al. (Skoog, Holler, & Nieman, 1998) and incorporates a form of the Discrete Cosine Transform (DCT). This form of the DCT provides the correct instrument line shape function. A cosine is a required function for the model in that the WLF at the ZPD must be at a maximum and the cosine correctly reflects this. The benefit of using a DCT for this work is that the imaginary part of the data that accompanies Fourier transforms is not generated as the cosine is both real and even and operates on real input data. Further, the DCT is defined for any vector of positive numbers, whereas the other forms, such as the Fast Fourier Transform (FFT), require 2^n number of samples. The flexibility in sampling allowed multiple values and sampling scenarios to be attempted based on exact critical sampling numbers, without worrying about zero-padding data and having to divide out the useful portion of the spectrum. Should the instrument noise portion of this work been considered, the sine component function would also be required as the noise is incorporated into the imaginary part of the transform as well as the real part. While the imaginary part is generally discarded upon inverse transform, studying the distribution of the noise would require an examination of the imaginary part. Additionally, some FTS systems use the phase information from the imaginary part to correct sampling error in the interferogram.

Skoog suggests the model is of the form

$$P(\delta) = \int_{-\infty}^{+\infty} B(\tilde{\nu}) \cos(2\pi\tilde{\nu}\delta)\, d\tilde{\nu} \tag{3.1}$$

where the *P* term is the measured power from the interferogram at a given OPD position, δ, *B* is product of the source of spectral radiance and the instrument response, and $\tilde{\nu}$ is the wavenumber. There are two main reasons why this equation cannot be used in the model. First, it requires an infinite spectral range, which is unphysical, and it requires an infinitesimal sampling interval, which again, is unachievable. This equation was modified to be executed in code form and to fit both the finite and discrete sampling environment of the FTS.

The model was written in 32-bit Matlab 2009a, v7.8.0.347 and is designed around the fact that computational advantages can be gained in processing large matrices simultaneously as opposed to serial or loop processing. While the transforms are intended to be continuous in integral form as shown in equation 3.1, they must be implemented as summations in a discrete sampling context. These summations can approximate integral form in the context of Riemann sums and the result of the sampling must be normalized by the total number of samples as well as the interval spacing of the samples. To form an interferogram, *P*, the following summations are accomplished

$$P(k) = \sum_{k=1}^{K} \sum_{n=1}^{N} [SRR(n) \cdot \cos[2\pi \cdot OPD(k) \cdot \tilde{\nu}(n)] \cdot \Delta\tilde{\nu}] \tag{3.2}$$

where *K* is the total number of OPD samples used, *k*, is the OPD variable, *N* is the total number of spectral samples in the spectrum, *n*, is the spectral variable, *SRR* is the sensor reaching radiance vector (spectrum), *OPD*, is the vector of OPD sample numbers, $\tilde{\nu}$, is the vector of wavenumber positions in the spectrum, and $\Delta\tilde{\nu}$, is the interval spacing of samples in the spectrum. Note that only the SRR value is used rather than the product of the SRR and an instrument response. No instrument

response function was implemented for this work as no real detector, nor fore-optics were modeled.

The effect of equation 3.2 is to modulate a series of N cosine values at a frequency dependent on the OPD position by the amplitude of the given spectral SRR. In the code, the first summation is done in one step by presenting vectors for the SRR and spectral band so that the product of these at a given OPD point is a single value. These values are then summed across all OPD points to produce the interferogram vector. In this manner, the interferogram is built up through the summation of the contributions of each of the varying frequency cosine functions at each OPD point. Figure 3.5 shows a test of the instrument line shape function of the model. An impulse function is approximated as the spectral input by nulling out the spectrum values and setting one sample position to a value of one. At full resolution, the result should essentially be a replica of the impulse with minor 'feet' due to the apodization. Figure 3.5 shows the impulse function in red and the instrument response at 0.73 maximum OPD (MOPD) in blue, where the left chart is apodized and the right is not. Note the difference in scale.

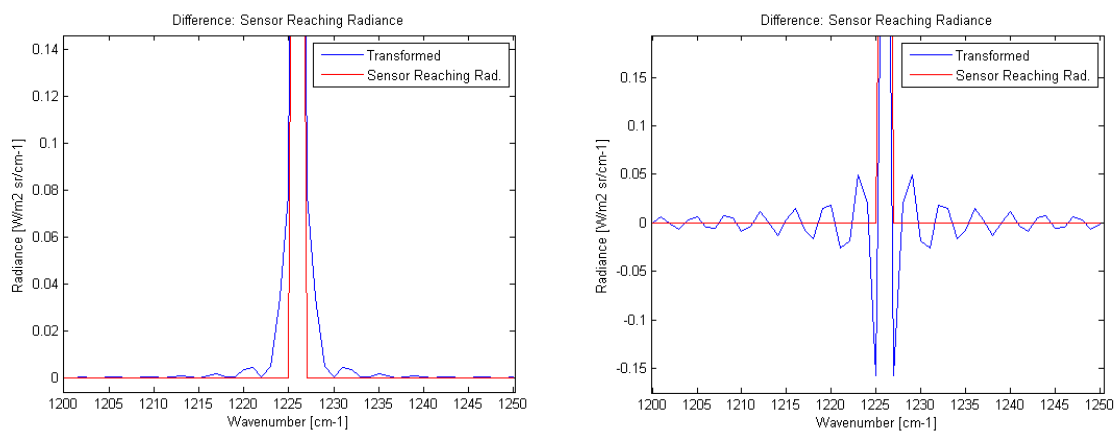

FIGURE 3.5 IMPULSE FUNCTION TO TEST INSTRUMENT RESPONE AT 0.73 CM MOPD, APODIZED (LEFT), UNAPODIZED (RIGHT)

The body of the function on the left is widened corresponding to the reduced resolution and reflects the apodized instrument line shape function. The total height of the response is lowered, reflecting the energy lost to the side lobes. In contrast, the unapodized function shows substantial negative lobes.

After the interferogram has been subjected to various error modules and apodization, it is converted back to wavespace, *S*, via an inverse transform given by the following equation

$$S(n) = \frac{2}{K}\sum_{n=1}^{N}\sum_{k=1}^{K}[P(k)\cdot\cos[2\pi\cdot OPD(k)\cdot\tilde{\nu}(n)]\cdot\Delta OPD] \qquad (3.3)$$

where all variables retain their previous definitions and ΔOPD is the MOPD value. Note that for radiometric accuracy, the inverse transform requires normalization by the total number of samples of OPD. The effect of the ΔOPD term is essentially to scale the spectrum based on the instrument resolution. If this term were absent, all spectral features would retain the same peak height as the

FWHM expanded with decreasing maximum OPD. This would imply not only a warmer gas, but also higher concentration-pathlength as well. This term essentially normalizes the distribution of spectral energy as the FWHM expands, maintaining overall radiometric accuracy.

Interferogram Processing

As mentioned in the background, the apodization chosen for this work was a straight forward triangular function. This function is applied as a right triangle for single-sided interferograms and mirrored for double-sided. The function is normalized to a height of two so as to preserve the total power of the spectrum. Both apodization functions used in this work are shown in Figure 3.6 for the same impulse function shown in Figure 3.5.

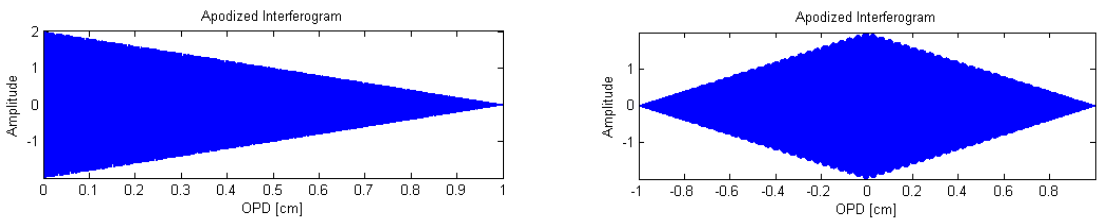

FIGURE 3.6 APODIZATION FUNCTIONS FOR SINGLE (LEFT) AND DOUBLE-SIDED (RIGHT) INTERFEROGRAMS

Since the impulse response function would be infinitesimally narrow in wavespace, it would have infinite support in the spatial domain and Figure 3.6 clearly shows that the approximation of the impulse function elicited the proper response. The interferogram shows the cosine function that corresponds to the wavenumber frequency of the impulse, which fully outlines the envelope of the apodization function.

The major differences between single and double sided interferograms is in their collection and processing times as well as sensitivity and is covered in section 2.5.8. As far as the model is concerned, any type of interferogram from any range of positive and negative OPD can be formed to varying degrees of success. This flexibility also allows one to study the impact of asymmetric interferograms, but was outside the scope of this work. Additionally, any range of sample points can be set as opposed to values limited by 2^n (where *n* is any positive integer). Single sided interferograms are suggested for use in gas detection and identification where speed counts, but quantification and temperature estimates should be derived from double sided interferograms.

Band filtering as applied to interferograms in this work essentially consists of notch filtering the interferogram over regions that are sensitive to a spectral feature of interest. The basis of this filtering is discussed in section 2.6.3. While the benefits of background and minor atmosphere removal were compelling, it was decided against pursuing interferogram filtering for detection as many of the errors that the interferogram was subjected to could easily corrupt the region. Circumventing these errors would require multiple regions of monitoring in the interferogram. However, in the spectral domain, remaining interferogrammetric components corresponding to other facets of the spectral feature still transform to the same wavespace and render some

evidence of the spectral feature in the same band. Thus, in the spectral domain, only one band needed to be monitored for detection, despite interferogram errors.

Uncertainty in the Moving Mirror

Uncertainty in the moving mirror was accomplished by applying a normal distribution to the location of each OPD sample point and this approach has been suggested in multiple efforts by Alexander Zachor in 1977 and Douglas Cohen in 1997 and 1999, to name a few. Mirror error was either classified as coarse or fine. A coarse mirror error was an error in sampling equal to one sampling interval. In physical terms, this would be a displacement in the position of the moving mirror due to a sudden disturbance once the sampling trigger was initiated. A fine error would be classified more as errors relating to the physical uncertainty of the mirror position due to uncertainties associated with the velocity, metrology and moving mechanism of the mirror, and would occur far more frequently than coarse errors.

Mirror errors in the model can generally be exacerbated by increasing the standard deviation of the distribution. The degree of mirror error realized depends on the number of samples used to form the interferogram. In this manner, an interferogram that was critically sampled (having no samples beyond the Nyquist level) would show significant mirror error effects in the spectral domain with only a few samples in error. Whereas, an oversampled interferogram would show only minor ghosting or slight bias affects for the same number and magnitude of sample errors. However, if the magnitude of the mirror error were scaled by the sampling interval, the effects would be essentially identical between the two interferograms. Further, in understanding the impact of the error, the location in the interferogram where the error occurs determines the impact to the resulting spectrum. Thus, one could expect an impact to the bias with an error on the low frequency side of the interferogram while the fine spectral detail and system noise would be impacted on the high frequency side.

It follows that if the number of sample points over a fixed maximum OPD (MOPD) were increased, those points would be sampled at smaller intervals than that of the critically sampled interferogram. If an event that triggered a piston mirror error, such as a vibration, were to occur, the physical displacement of that positional error would be the same regardless of the sampling interval. Thus, oversampling the interferogram really doesn't buy much in terms of providing a buffer to coarse mirror error.

The actual implementation of the coarse error module consists of generating a random number from a normal distribution with 0 mean and a given standard deviation. This number is then rounded to the next integer and added or subtracted from the sample point location integer where the OPD measurement will occur. Thus when a mirror error is recorded, the OPD radiance value for the point either before or after the actual sampling position is recorded and transformed to the spectral domain. Coarse mirror errors can cause substantial changes in the bias of the spectrum over defined intervals, the appearance of ghost peaks, and the absence of anticipated spectral features.

The fine mirror error module allows for multiple positions to 'sub-sample' the interferogram on either side of the actual critical sampling point. Allain proposed, but never implemented in his model, a method to add a small addition to the OPD value given before the interferogram point was calculated. This straight forward approach was implemented in this work as a similar implementation of the coarse mirror error excluding the rounding to integer sample values. He also proposed to smooth the distribution or errors, but the more taxing, abrupt changes were implemented in this work.

Fine error effects at this magnitude of error below the critical sampling interval typically manifest as the appearance of noise in the spectrum as previously described. The following figure series illustrates the affect of fine mirror error. The first figure in the series shows the pre-error state of the collected spectrum. This spectrum contains 290K benzene with a 1km MLS atmosphere over 300K asphalt. It is derived from a phase corrected single-sided interferogram of 4,096 points, which is approximately 1.6x oversampled from Nyquist. The bottom portion of the panel shows a ratio of the two spectra. While a perfect replica should yield a value of one across the spectral range, the ratio of spectra come very close and have very minor differences and slight bias from the transform and discrete sampling processes. The vertical scale on the ratio plot implies differences of less than one part in 10,000. Figure 3.7 shows this initial ratio and is provided as a point of reference.

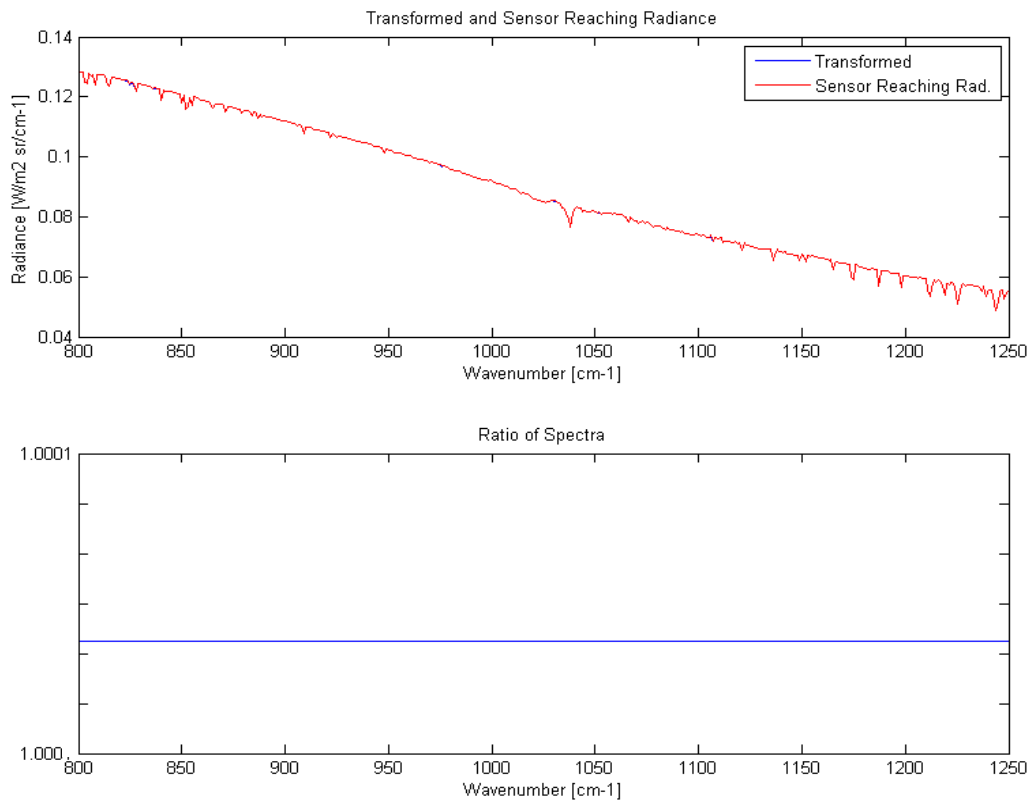

FIGURE 3.7 SRR AND TRANSFORMED SPECTRUM AND THEIR RATIO BEFORE FINE MIRROR ERROR

Figure 3.8 shows the next panel in the series which characterizes the extent of the mirror error. The upper portion of the panel is a chart that shows the magnitude of the mirror displacement in

microns experienced at each location in the interferogram. The vertical extent of the chart measures the displacement in microns of the recorded OPD radiance value for that particular sample point in the interferogram. The lower panel is a histogram of mirror displacement errors in microns from the interferogram. Here the sample intervals have been converted from sample position integers to displacement distance in microns. The first set of symmetric vertical bars indicate one sample interval for the number of points specified (4,096 in this case), and the larger bars indicate a critical sample interval. Note that the error values in this example approach one full sample interval, but do not exceed it.

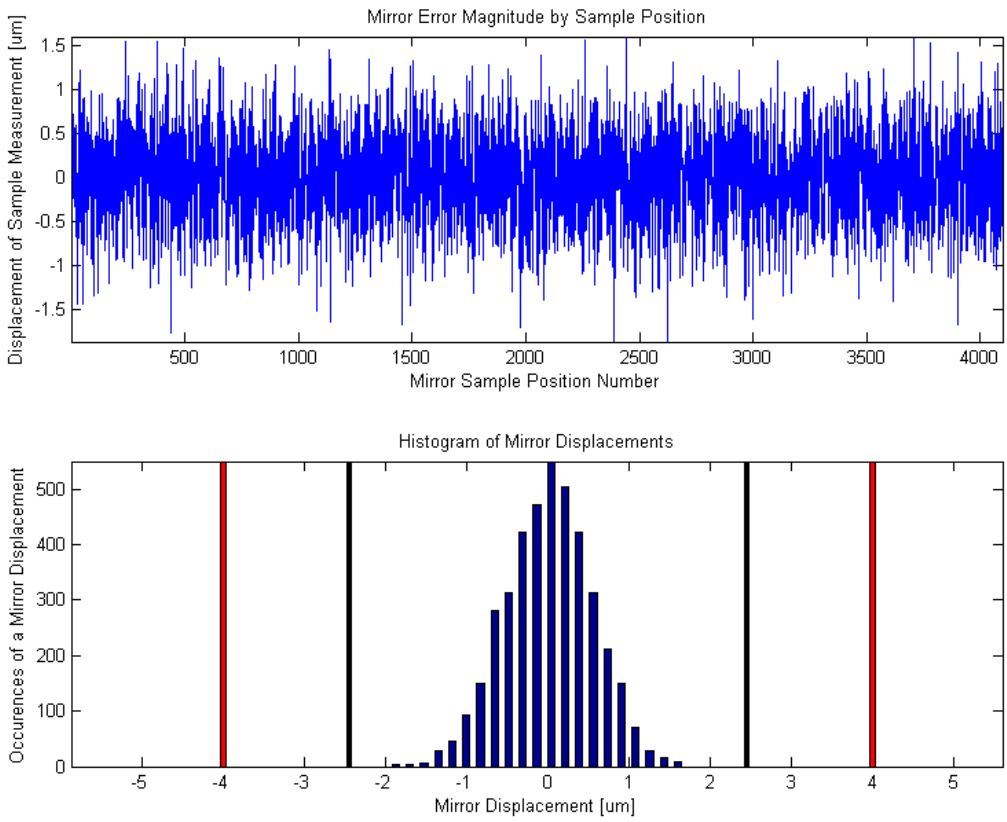

FIGURE 3.8 FINE MIRROR ERROR DISTRIBUTION AND MAGNITUDE (TOP), HISTOGRAM OF MIRROR DISPLACEMENTS (BOTTOM)

The final panel in the series, Figure 3.9, shows the affect of the fine mirror error on the resulting transformed spectrum. As can be seen in the chart, the error essentially manifests itself as noise in the spectrum with minor perturbations to the spectral feature and bias with mean deviation of about 10%. With fine mirror error, all of the samples in the interferogram are in error (sub sampling interval) and in the above case, the mean mirror error at each sample point was 0.39 µm (about 16% of the sample interval). The fact that benzene appears in multiple regions of the interferogram ensures that it will be visible to at least some extent as this and the next example demonstrate. The spectral matched filter scores for this series showed that the value for the target gas subjected to fine mirror error only decreased by 18%.

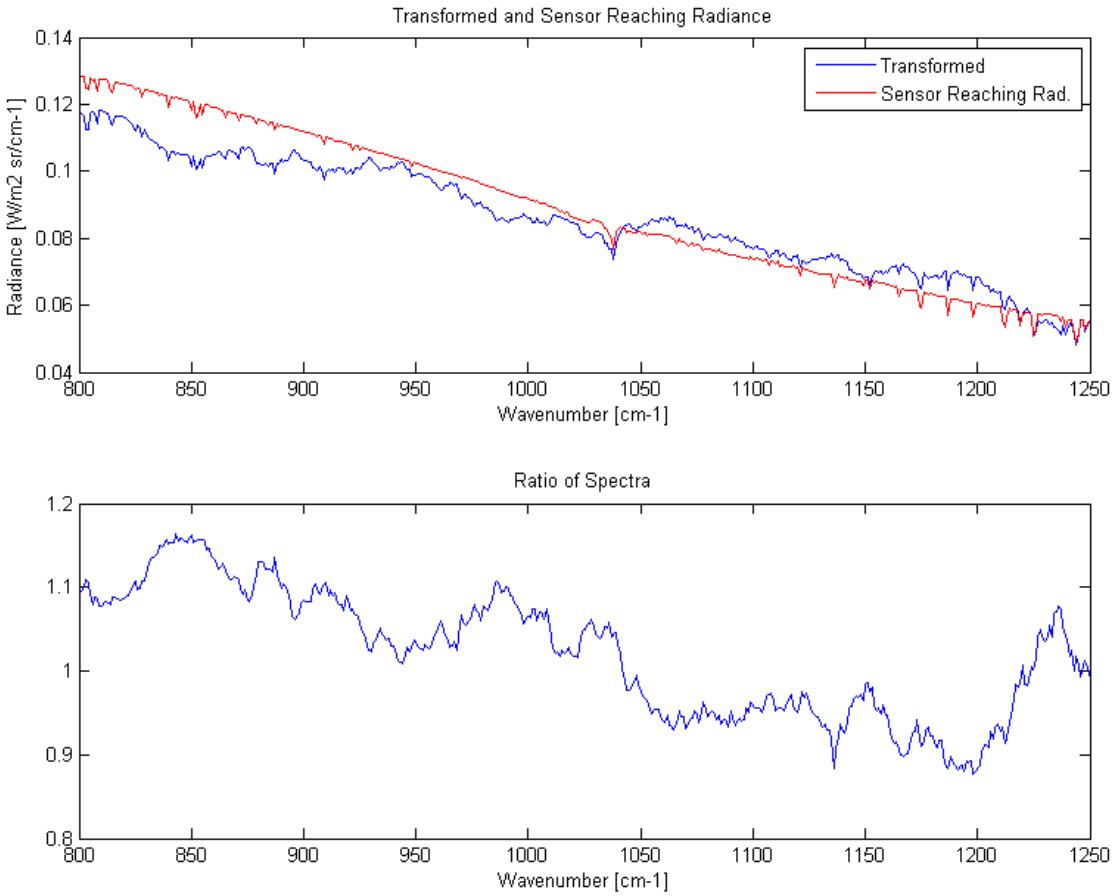

FIGURE 3.9 FINE MIRROR ERROR RESULTING SPECTRUM (TOP) AND RATIO OF SPECTRA (BOTTOM)

Conversely, coarse error can cause significant fluctuations as shown in the resulting spectrum below. In this example, the same starting conditions as in Figure 3.7 apply.

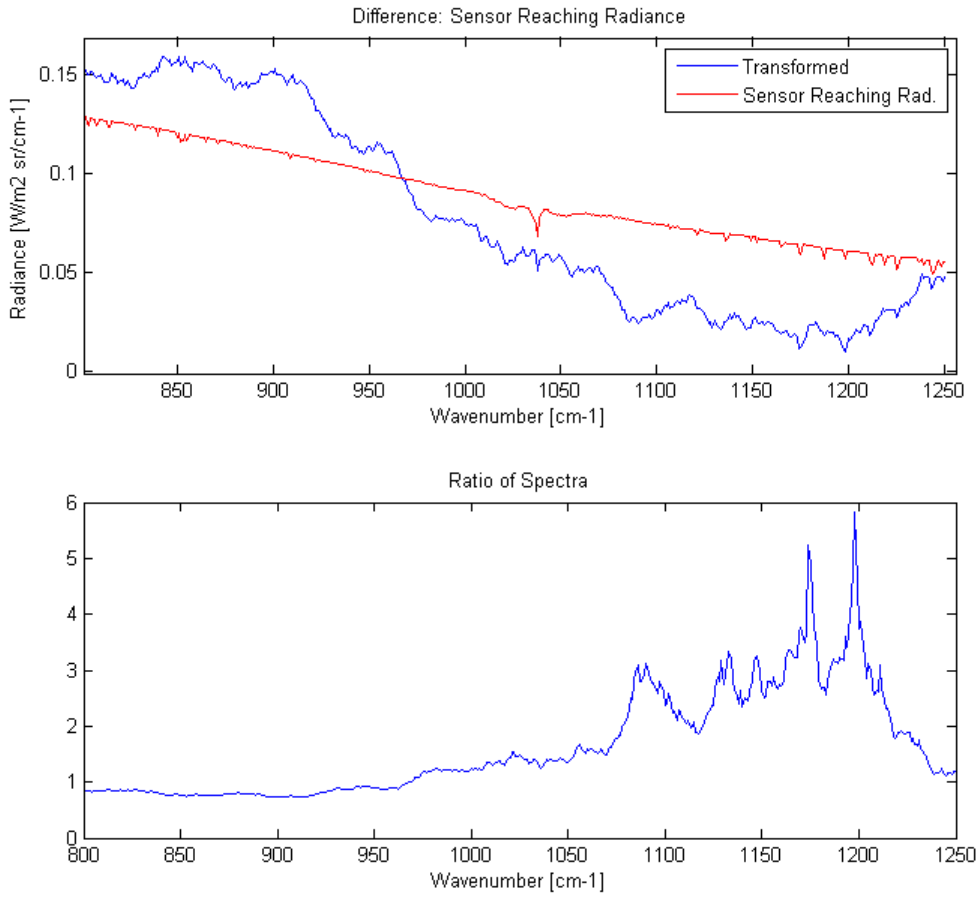

FIGURE 3.10 EFFECT OF COARSE MIRROR ERROR ON TRANSFORMED SPECTRUM (TOP), RATIO OF SRR AND SPECTRUM (BOTTOM)

As mentioned above, despite substantial interruption to the interferogram at both full interval and critical sampling levels, there is still evidence of the spectral feature of interest. The baseline bias of the spectrum is completely lost and there are several peaks that have no real basis in physical origin. In some cases, the radiance values are over 5x different from what they should be.

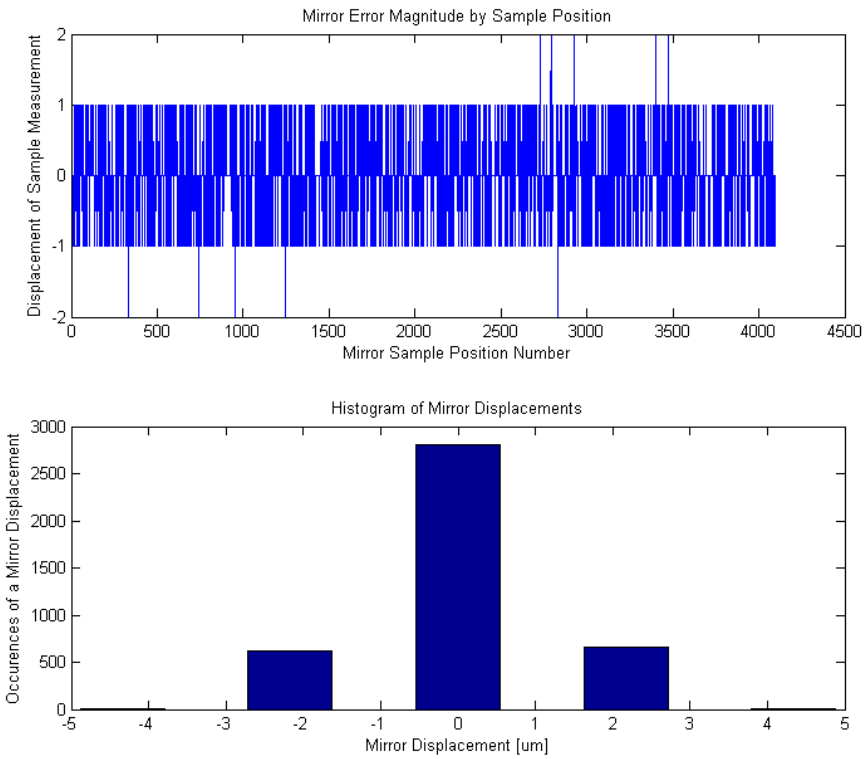

FIGURE 3.11 MIRROR ERROR DISTRIBUTION AND MAGNITUDE FOR COARSE ERROR (TOP) AND HISTOGRAM OF ERRORS (BOTTOM)

The top panel shows the number of full sample interval displacements (2.44 µm) of the mirror at each OPD sample point. The histogram in Figure 3.11 shows that only 10 of the samples (approximately 0.25%) are beyond the critical sampling limit (4 µm) and are largely responsible for the error in radiometric accuracy. However, note that despite the substantial interruption to the spectrum, the spectral feature is still viable. The spectral matched filter detection score decreased by only 22%.

Of course, these are just 'one roll of the dice' examples, to gain a better understanding of the range of impact of mirror uncertainties on the transformed spectra, several hundred trials must be conducted. A Monte Carlo simulation was run on both the mirror and jitter error modules, as well as with both errors applied simultaneously. The results of the Monte Carlo simulations are found in Chapter 4.

Uncertainty in the Primary Mirror due to Jitter

The uncertainty in the primary mirror stems from the turbulent motion and vibration associated with the flight environment. The primary mirror, also called the steering mirror if movable, is responsible for directing the SRR to the aperture of the FTS. While there are both high and low frequency sources of jitter, the primary source of jitter to airborne sensors is low frequency below 300 Hz and this will be the primary focus for this work (Beer, 1992). Many propose that the jitter uncertainty can also be modeled by a random normal distribution, as with moving mirror error, which dictates the pointing accuracy of the mirror in the presence of disturbances (REF).

Jitter has many components that must be carefully considered in the model development for an airborne sensor. First, the extent of the LOS perturbation must be considered. In this work it was allowed to extend up to +/- 5 ground elements away in any direction in a uniformily distributed sense from the targeted ground element in two DoF. That is to say, each X / Y sampled coordinate in the scene was assigned a jitter LOS pointing error in the X and Y direction. The magnitude of the error was a value between -5 and 5 selected randomly from a uniform distribution over that range. When the magnitude of the LOS jitter error is determined in terms of distance from the target, the envelope of distance from the targeted point is normally distributed. In a typical scenario used in this work where the ground element size is 0.2 m on a side and the aircraft is at 1 km altitude, the maximum angle of deflection corresponds to approximately 1 mrad, which represents 20% of the IFOV.

Next, the frequency of the disturbance must be considered. In the case of engine vibration, this disturbance would be considered periodic. For this work, the disturbance modeled was both periodic and aperiodic in essence. In the low frequency case, a 267 Hz periodic disturbance rate was used. The next question is what should the duration of this disturbance be, i.e., how long will the mirror dwell off target? In this work the duration is a random event that is normally distributed around a mean of 300 OPD samples (this number should vary with moving mirror velocity). The samples can extend both forward and behind the given sample point where the jitter event is triggered. In this sense, while the disturbance in considered to occur periodically, the actual effect due to duration, is aperiodic – a more difficult case. Figure 3.12 shows an example of a low frequency jitter spectrum implemented in this work. This particular profile was used for a 20 Hz scan rate and yielded 28% of the interferogram samples off target.

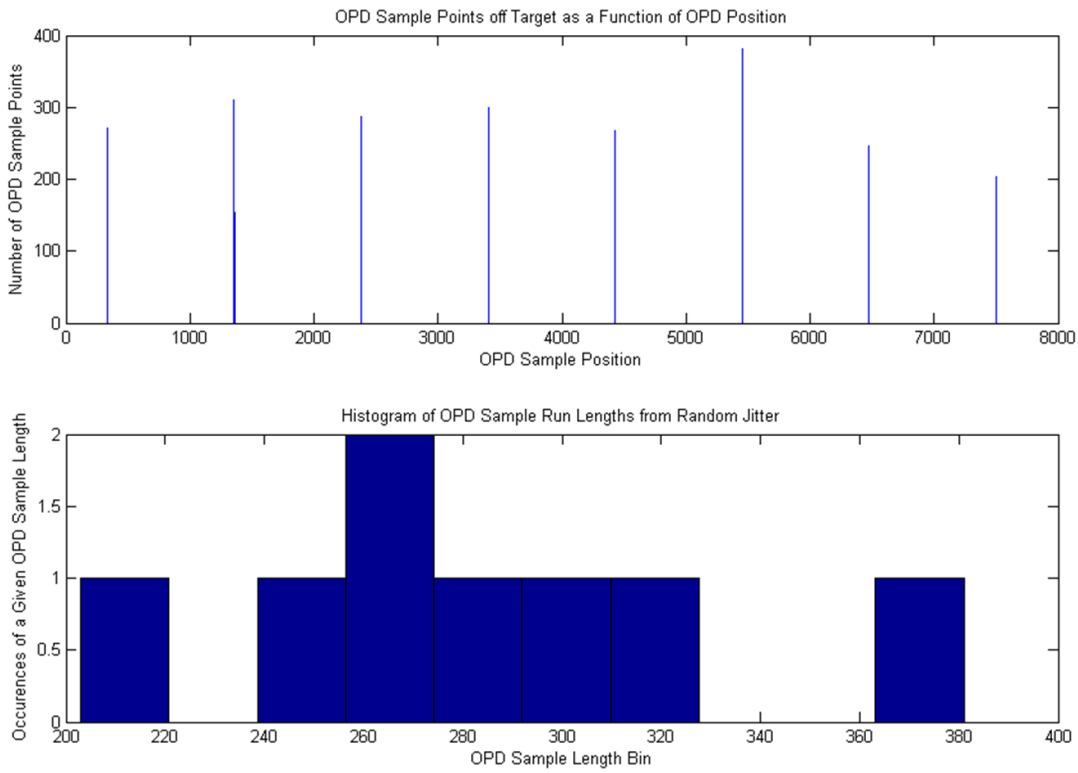

FIGURE 3.12 LOW FREQUENCY JITTER SPECTRUM (TOP) AND HISTOGRAM OF JITTER RUN LENGTHS (BOTTOM)

Figure 3.12 shows first the number of sample points that were off target in the vertical extent, called the jitter run length, while the horizontal shows the position in the interferogram where the disturbance is taking place. A jitter run length is defined as the number of interferogram samples that are collected while the mirror is off-target. The bottom panel is a histogram of the jitter run lengths that affected this particular interferogram of 8,096 samples. As can be seen, the disturbance occurs periodically, while the duration is aperiodic.

To ease the computational burden of implementing jitter, the disturbance of the mirror off target is considered a discrete process, whereas in reality the mirror would spend some time drifting over multiple ground elements with a given acceleration before settling at a maximum displacement, then drifting back to the target position with some likely dampening effect. Again, here the sampling occurs such that the last good sample point in the interferogram is of the target and the very next sample during a jitter event is of the maximally deflected position of the mirror and it stays there until the duration of the event is met, then returns to the target directly. This motion provides abrupt changes to the interferogram and is a more taxing event.

Clearly, one can get a sense that the effect of jitter on the resulting spectrum is dependent on the scan rate of the system. At insufficient scan rates the target might be completely missed whereas with fast scan rates, only a small portion of the interferogram would be affected and in only one location, or missed altogether. Figure 3.13 shows this effect with a critically sampled single sided interferogram of methyl chloride at 200,000 ppm-m with a 30K thermal contrast (target gas grossly

over-emphasized for demonstration purposes) over 300K asphalt. The alternate pixel viewed during the jitter event consisted of no gas, the same 5km MLS atmosphere, and 300K sand. This represents a relatively gentle transition between the two pixels and accentuates the effect of the absence of the gas from the scene.

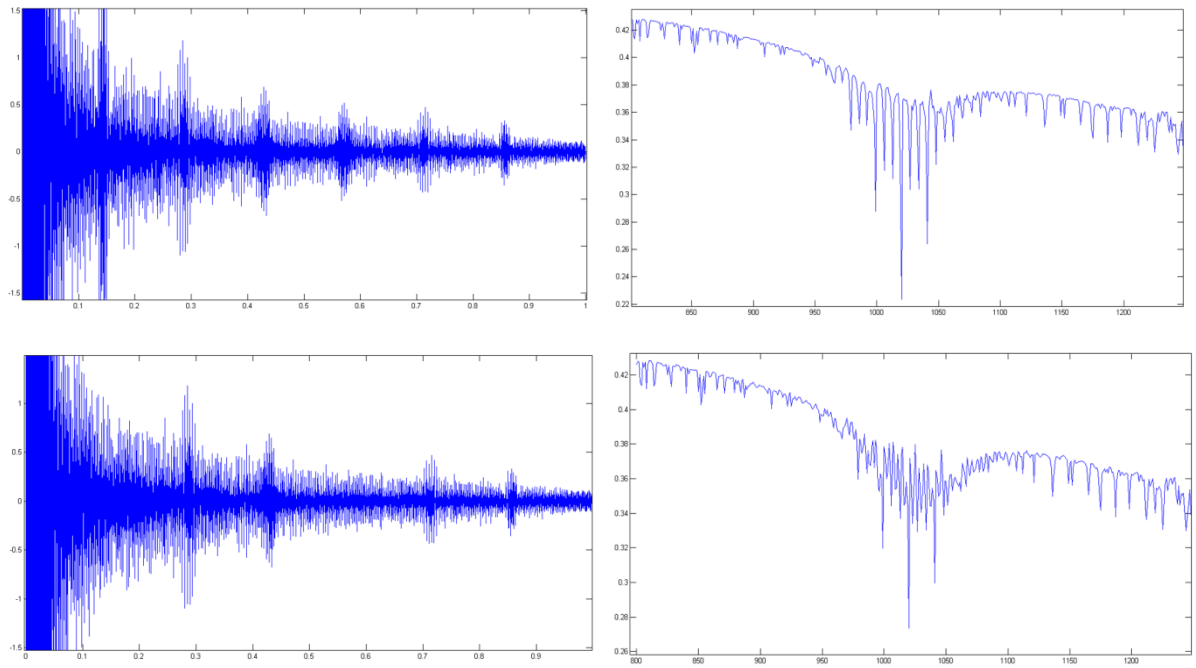

FIGURE 3.13 EFFECT OF LOW FREQUENCY JITTER (BOTTOM) ON CRITICALLY SAMPLED METHYL CHLORIDE (TOP), INTERFEROGRAMS (LEFT), SPECTRA (RIGHT)

As has been shown previously, the periodic nature of the spectral feature manifestation in the interferogram for methyl chloride is quite evident in the normal interferogram at top left. The corresponding spectrum is shown at top right. In the lower panels, the absence of key interferogrammetric portions due to viewing off target pixels shows the effect in the transformed spectrum. While the major identifiable spectral features are still present, several ghost peaks have appeared that have no basis in physical origin. Note that the bias, in the form of the radiance values, and the fine spectral detail are largely intact -- this is because both the high and low frequency content of the interferogram remain intact.

The impact is also dependent on the scene FOV. If the GSD is finely sampled as in a low altitude collect, the spectral variability might be high between ground elements, inducing large magnitude changes in the interferogram due to fluctuations in the recorded radiance values. With large sensor footprints, the spatial extent is likely to appear more spectrally uniform and would have a smaller overall impact.

Airborne Scanning

The main goal of implementing a rudimentary airborne scanning module was to determine what impact the collection of an interferogram over spectrally varying ground elements would have on the resulting transformed spectra. In this way, various combinations of aircraft velocity, scan rates,

and spectral resolutions could be tried to determine an optimal collection scenario that would be able to capture spectral evidence of a target gas plume while overflying the area.

The airborne scanning module uses flight information to determine the number of interferograms collected at a specified scan rate for a scene of a given extent. The most accurate way to visualize how this model works is to think of the FTS model as having a single detector element and that it flies individual flight lines collecting as a push broom over the entire scene in X/Y space, where X is the number of flight lines required to cover the scene and Y is the number of ground elements that comprise the vertical extent of the scene. In more familiar flight collection terms, X would be the cross track direction and Y the along track direction. In this manner, the nadir collection assumption for each ground element holds true as well as the other aforementioned optical assumptions. The way the error and jitter modules are currently implemented in the airborne scanning module assume that the entire scene was built in one pass with a push broom collect as all errors occur at the same point during the collect across the scene. The only problem with this is that the collection geometry doesn't accurately reflect what the model produces. The error and jitter modules are implemented such that they can easily be changed to produce random error over each ground pixel in the scene independently as well. The fine error and jitter module are implemented this way when 'flown' against the modeled scenes.

Using the platform velocity and scan rate, the model determines how much time each scan must take, including built-in overhead, or duty cycle time, if required. It then determines how much distance is covered on the ground during that time. Taking into account the GSD for the given IFOV, altitude, and detector element, the model determines how many ground elements will make up a complete interferogram during the scan. The model then divides the total number of OPD sample points among the required ground elements. This process generally results in one or two extra points left over that will not divide evenly into the number of ground elements. These points are assigned a 0 radiance value as they would have been apodized to near zero regardless due to their position at the end of the interferogram.

Each OPD sample point is assigned a radiance value based on the weight of scene contribution of each of the ground elements as they appear in the IFOV. For example, the first OPD sample point derives its radiance from the entire extent of the first ground element, then as the aircraft moves forward slightly, the next OPD sample point is comprised of mostly the first ground element, but also the first edge of the second. Ground elements are X/Y locations in the SRR cube that contain a single material, atmosphere, and possibly a gas all at specified temperatures and concentrations. The extent of a ground element is spatially homogenous, and the mixture of contributions from two ground elements is just the spatially weighted sum of their respective SRR spectra. Figure 3.14 illustrates this concept.

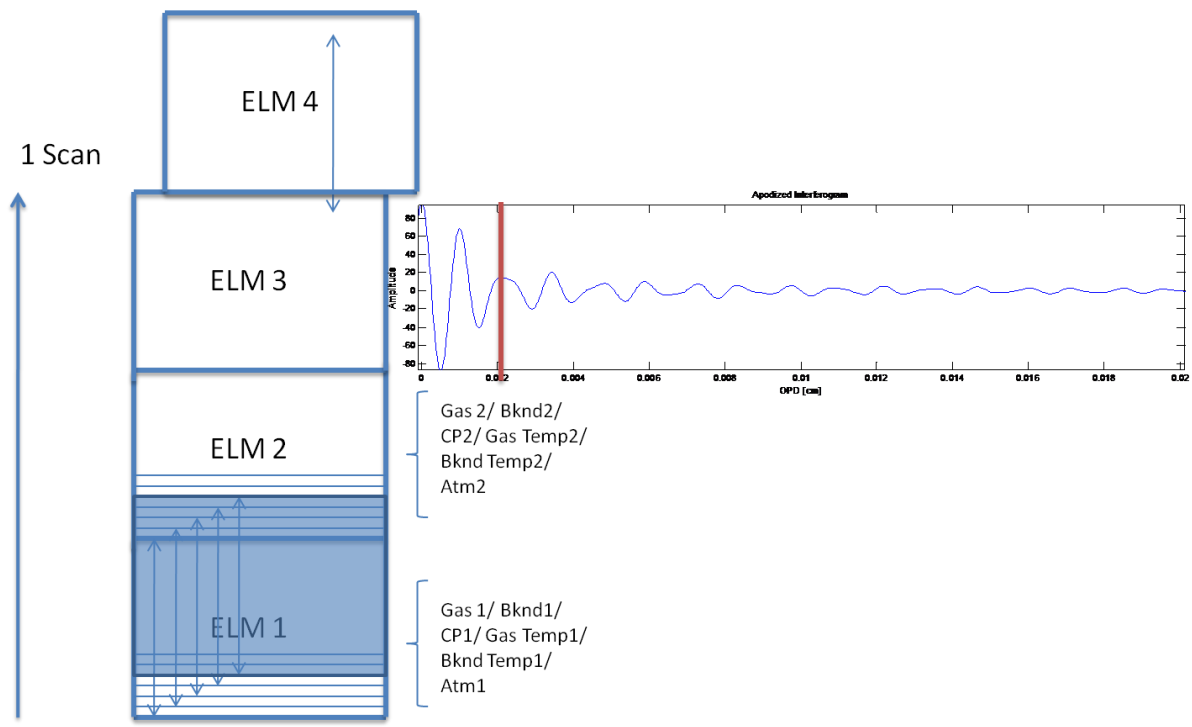

FIGURE 3.14 AIRBORNE MODEL INTERFEROGRAM FORMATION

In the scenario illustrated by Figure 3.14, a total of three ground elements will comprise a complete scan, while the fourth element will have overlap between the two scans. The annotations to the right of the first and second ground elements indicate the properties attributable to each and emphasizes that they are different. The darker blue window represents the IFOV of the sensor. The graduations indicate the distance the IFOV moves due to aircraft motion between OPD samples. The red line on the interferogram to the right indicates where the current IFOV OPD radiance value will be recorded and as the IFOV continues to move up, the interferogram will continue to form to the right. As the scan continues, the IFOV will move to cover the ground up until it is centered over the range extending from the last graduation in element three to the last graduation in element four, with the next scan beginning with the IFOV centered over ground element four. This coverage essentially represents 25% overlap between scans.

As a final point, if the along track extent of the scene in this example consisted of 100 ground elements, the resulting image cube would have only 33 pixels in the vertical dimension. This is due to the fact that each scan represents the spectral content of three (and part of four) of the ground elements. When the original scene is compared to the image cube, it appears that the vertical dimension has been compressed – and it has. This prevents direct comparisons between the image cube and the original scene. Figure 3.15 illustrates this point by comparing a simulated gas plume and a detection plane for a collect of a 2x oversampled, double sided interferogram at 1 cm MOPD, with a scan rate of 60Hz and aircraft velocity of 50 m/s. The detection planes show only the first 50 'flight lines' of the scene and reflect the vertical binning of ground elements in the along track direction. The detection metrics are introduced in a later section.

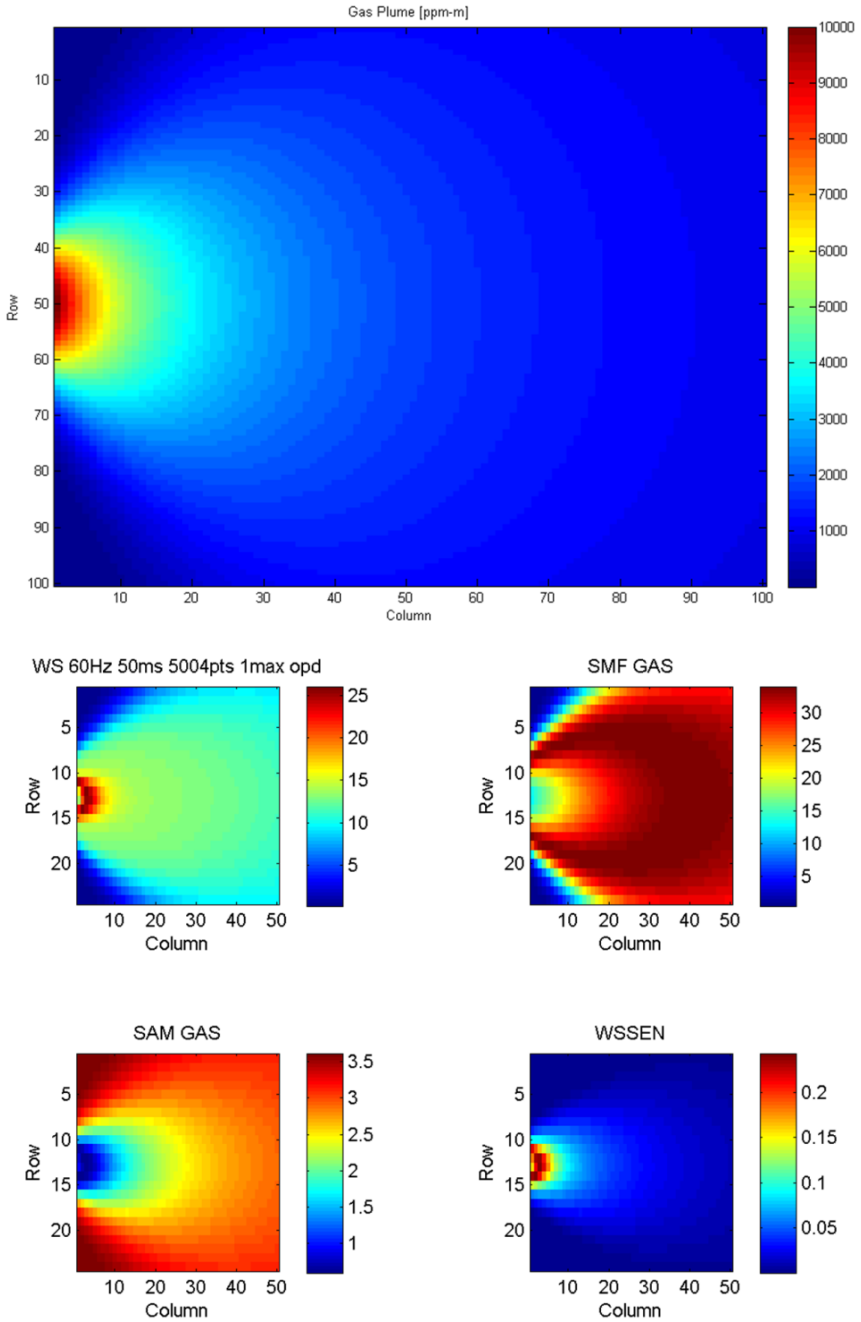

FIGURE 3.15 COMPARISON OF SCENE (TOP) AND DETECTION PLANE (4-PANEL, BOTTOM) FOR NOTIONAL GAS PLUME

Due to the use of canned MODTRAN atmospheres, a rocker point, or point of fixation, cannot be used to attempt to complete each scan over essentially one point on the ground. A point of fixation would require multiple looks through the atmosphere at various slant ranges requiring a unique MODTRAN run for each OPD sample point and that amount of runs is completely impractical (8,006,400 points for a standard scene at full spectral resolution). Another possible approach, suggested by Allain in implementing his airborne model, was to generate a few different atmospheres over the various look angles and interpolate atmospheric parameters between the

points. This approach was also determined to be too computationally taxing for this work. In the end, it was determined that each sample point would be collected as if from a nadir view.

The next question that arose in optimizing the model was what scan rates to select. When reviewing the current state of high resolution FTS work, the FIRST-based airborne model scan rate versus resolution trade space shown in Figure 2.24 gave a good starting place. The chart shows that in the LWIR, the maximum datacube rate is 1 Hz at 10 cm^{-1} resolution and a tenth of that at 1 cm^{-1} resolution, which will serve as an anchor point in the scan rate range. However, in order to cover the rendered scene with a sufficient number of sample points at realistic platform velocities, the scan rate minimum was moved up to 2 Hz, which afforded one complete scan of the scene at full resolution at a platform velocity of 40 m/s (78 knots, or 90 mph).

When using a voice coil as the moving mechanism, the theoretical limits of the technology were considered in arriving at a maximum scan rate. This approximation was made from information from Manning, suggesting that the acceleration of a typical voice coil actuator for a moving FTS mirror is approximately 33 m s^{-2} with a force of approximately 1N (Manning, Bowie, & Griffiths, 2000). Griffiths concurs that typical forces achieved by voice coil actuators in FTS systems range from 1-2N (Griffiths & de Haseth, 2007). The force exerted by a voice coil is given by equation 3.4

$$F = Bli \tag{3.4}$$

where *F* is the force exerted in newtons, *B* is the magnetic field in tesla, *l* is the length of windings, and *i* is the current in amperes. For the example given by Manning above, the magnetic field was 0.2 T, with 5 m of windings, and 1 ampere current. The mass of the mirror was 30 grams, yielding an acceleration of 33 m s^{-2}. The scan rate of a voice coil is derived from the second derivative of the position function, which is a time dependent cosine. The maximum acceleration is therefore given by the product of the amplitude and the square of the frequency. Rearranging the variables, the maximum amplitude/displacement, and therefore resolution, can be plotted against the rate of oscillation (frequency) shown in equation 3.5

$$R = \frac{\omega^2}{2*100*33} \tag{3.5}$$

where *R* is resolution in wavenumber, the 100 term is the conversion to centimeters from meters and the factor of two converts the mirror travel to maximum OPD. Figure 3.16 shows how these values result in the curve shown in relating FTS instrument line shape to scan rate.

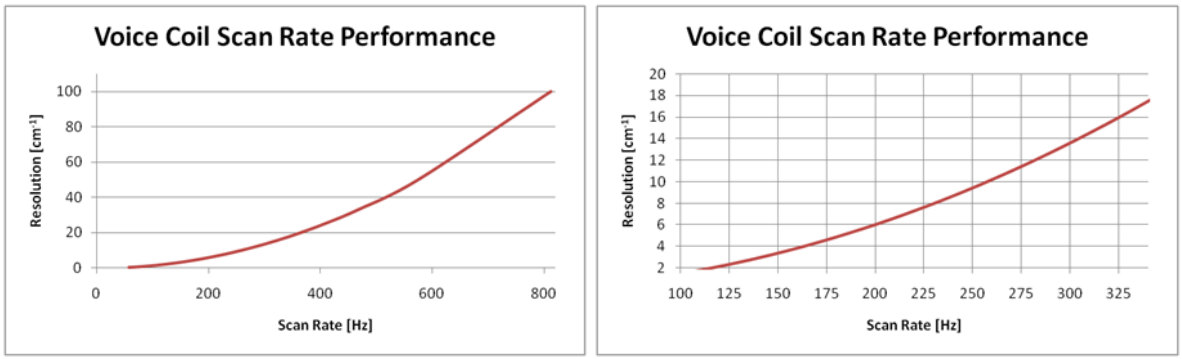

FIGURE 3.16 VOICE COIL SCAN RATE THEORETICAL PERFORMANCE (LEFT) WITH MAGNIFIED REGION (RIGHT)

In comparing the same values as the FIRST performance, the theoretical limits approach 250 Hz at maximum acceleration at an MOPD consistent with 10 cm^{-1} spectral resolution. Taking into account that the mirror must accelerate from a stop, then decelerate, turn around and collect in reverse, the maximum rate should at least be halved, if not more. The limit of the scan rate is driven by the MOPD desired and at a 1 cm MOPD, it would yield a maximum scan rate of approximately 100 Hz. The maximum scan rate chosen for the scene studies in this work was 40 Hz – still far beyond current capability at 1 cm MOPD, but certainly stresses the theoretical limits.

3.1.2 Sensor Reaching Radiance Model

The sensor reaching radiance consists of radiance contributions from various surface materials, the gas plume, the atmosphere, and any interactions therein.

Surface Materials

The model begins with energy leaving the radiative surface, propagating through the gas, while the gas absorbs, or emits on its own, then propagates to the sensor aperture. Both the amount of energy as well as the spectral distribution of that energy leaving the surface is determined by the temperature and emissivity of that material. In this work, emissivity files derived from real data material collects in Trona, California were used. The emissivity files are read in to the model, then are used to modulate a black body function (equation 2.1) at a user specified temperature. Figure 3.17 shows some of the emissivity curves used in this work.

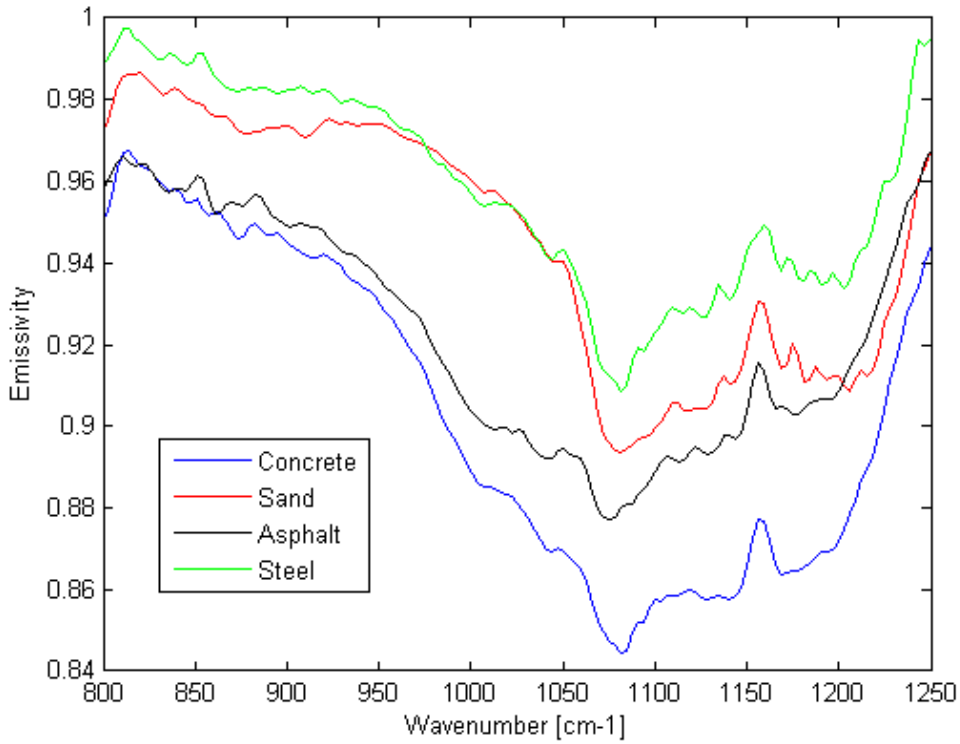

FIGURE 3.17 EMISSIVITIES OF SURFACE MATERIALS USED IN THIS WORK

Note that in general, the lowest emissivity value is around 0.85, while much of the other values are over 0.95. This means that when the black body function is modulated by the emissivity curve, there is only a maximum of 15% variation in the spectral radiance. While these minor spectral variations in the background in these materials have little impact to the success in detection of a gas in the LWIR window, the variation in emissivity plays a much greater role in the resulting gas visibility as discussed in section 2.4.2. To summarize, the reduced emissivity causes either a muted absorption, a loss of thermal contrast, or a conversion to emission, depending on the gas and surface temperatures, resulting in changes to the spectral feature shape being detected. This fact is critical in characterizing the performance issues for the FTS sensor over a scene with varied background materials. Figure 3.18 shows the emissivity modulating 300K black body curves in both wavenumber and wavelength.

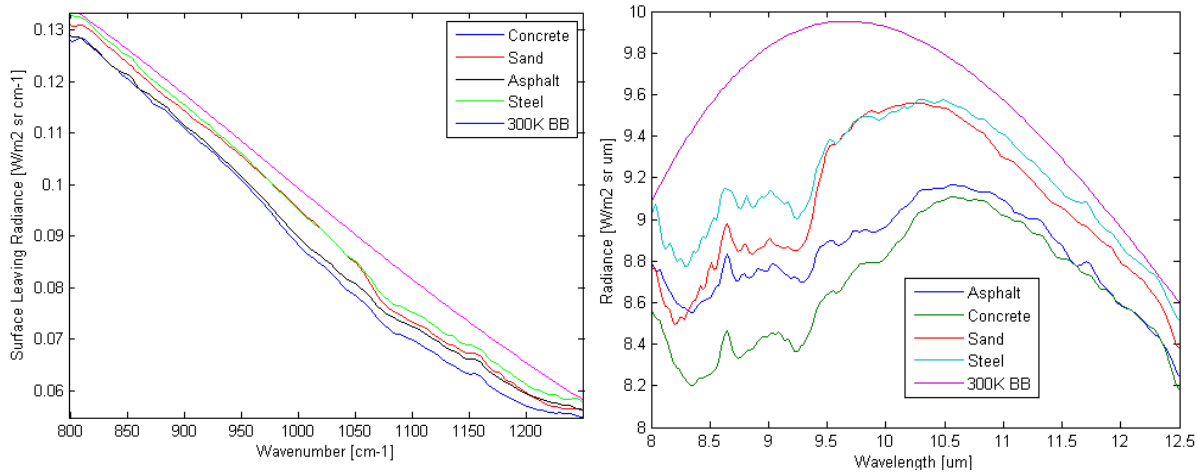

FIGURE 3.18 WAVENUMBER (LEFT) AND WAVELENGTH (RIGHT) 300K SURFACE LEAVING RADIANCE CURVES

Figure 3.18 shows the effect of the various emissivity files on a 300K black body. The variation in wavenumber appears to be minimal, but that is because as a proportion of the total radiance, the magnitude of change is less. The figure in wavelength was provided as well so that the differences are clearer, particularly when viewed in contrast to a 300K black body curve.

The figure shows that these materials generally conform to the overall shape of a blackbody radiator with deviations for various Restrahlen features. Asphalt, sand and concrete look very similar and this may be due to the nature of the field samples.

Asphalt and concrete are not homogenously composed, and therefore can have a relatively variable spectrum that is comprised of multiple spectra generated from its components. The relative strength of each spectral feature is related to the amount of the given component that is present in the sample. Asphalt is generally comprised of rocks of various sizes mixed with tar. Thus, the Restrahlen features associated with the spectra of the various rocks, particularly if large sizes are used, would comprise the majority of the observed spectrum. Generally speaking, concrete is composed of a mix of sand, gravel, and a binder. While the finer the gravel, the more homogenous the appearance of the concrete, large slabs are often mixed with large stones to stretch the use of the ingredients. Over the years, the concrete becomes weathered, removing the binder, revealing the stones used in the formulation. In the spectra shown above for asphalt and concrete, the two samples appeared very similar in composition, thus likely giving very similar spectra.

Finally, steel with rust is shown as a component background material. It appears to have virtually all of the same Restrahlen spectral features as the other materials. This may be due to a fine coating of sediment such as sand on the steel that modulates the underlying black body-like emission of the steel itself. For all intents and purposes, pure metals should nearly approximate a black body model.

One minor issue with using these emissivity files was in the spectral sampling interval of the emissivity data, which is about a third of the sampling of the total combined spectrum that is fed to the model. However, as demonstrated in the Solid Phase Spectroscopy section on page 11, solid

materials have comparatively slowly varying spectral features in contrast to those of gaseous species. Thus, it is a reasonable assumption to upsample the emissivity data to a rate consistent with the rest of the model and not expect any variation between the two data sets at the different sampling intervals.

Another minor issue came in trying to obtain an emissivity file of a manmade material that had substantial variation in emissivity. The only suitable file in the database was of a painted aluminum roof. However, it appears that the resulting emissivity spectrum is essentially a reflection of the atmosphere in part, possibly indicating an error in the emissivity retrieval process. While the data appears unphysical, it was still implemented in a small portion of one of the scene-based scenarios because it had so much variation and was so different from all of the other files used.

Target Gases

Several gases were chosen as potential target candidates from the PNNL library as previously described. The gases were chosen for either their spectral variability, well-established use in the field, or community interest. The spectra span the entire range of the LWIR window chosen for this work and provide multiple examples of multiband spectral features of both narrow and wide character.

All absorptivity data is shown at approximately 0.15 cm^{-1}, or about 1.5 nm at 10 µm resolution. Table 3.1 gives the pertinent information about the gases used.

Gas Name	**Formula**	**Peak Absorptivity** [ppm^{-1} m^{-1}]	**Spectral Feature** [cm^{-1}] Range (peak, if singular)
Benzene	C_6H_6	1.1×10^{-4}	1015-1060 (1038)
Methyl chloride	CH_3Cl	4×10^{-5}	950-1100
Ammonia	NH_3	3×10^{-3}	800-1200 (930,960)
Sulfur hexafluoride	SF_6	5×10^{-2}	940-950 (948)
Sulfur dioxide	SO_2	1.1×10^{-4}	1060-1240
Phosgene	$COCl_2$	4×10^{-3}	800-860 (849)
Carbon tetrachloride	CCl_4	1.9×10^{-3}	800-810 (800)

TABLE 3.1 TARGET GAS OVERVIEW

The PNNL derived absorptivities of each of the gases will be reviewed below in the same order as presented in the table. Figure 3.19 shows the absorptivity spectrum for benzene.

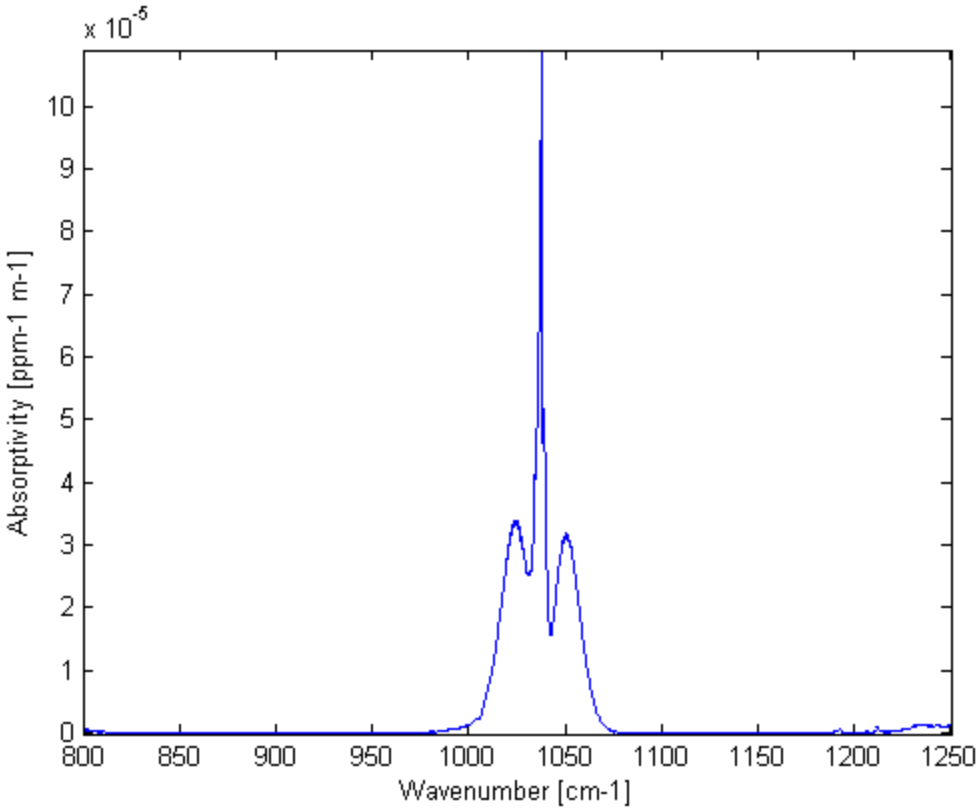

FIGURE 3.19 ABSORPTIVITY OF BENZENE

Benzene was the mainstay gas for the majority of this work for a multitude of reasons. First, benzene's narrow Q branch band head made for a 'needle in the haystack' target when considering the model's maximum one wavenumber instrument resolution. The Q branch shape also mimicked many of the rotational continuum lines of water that are present in the band of interest. Benzene was also of particular interest due to the multiple types of elements that comprise its spectral feature. It has both narrow and wideband elements as well as subelements in the P and R bands that are in between the width of the Gaussian-like base and the Q branch. The multiple bandshapes that comprise the spectral feature ensure that evidence of the gas will occur in multiple regions of the interferogram, which was an important area of study in this work.

Further, while the peak spectral feature occurs at 1038 cm^{-1}, the band position overlaps that of one of the bands of ozone whose peak occurs near 1039 cm^{-1}. This made for an interesting target that was highly dependent on the presence of ozone in the scene, which can be quite variable. At high altitude, or considerable downwelled radiance contribution, the presence of ozone completely enveloped comparatively low concentration-pathlengths of benzene (reference Figure 2.8). Figure 3.20 illustrates the difference between a spectrum generated with benzene and without (upper trace), consistent with typical ozone concentrations from a 5 km altitude. In the trace with benzene present (lower trace), the second edge of the peak at 1038 cm^{-1} is really the only evidence of

benzene presence other than the slightly expanded band shape. Cleary, this is a difficult target from a spectral sense.

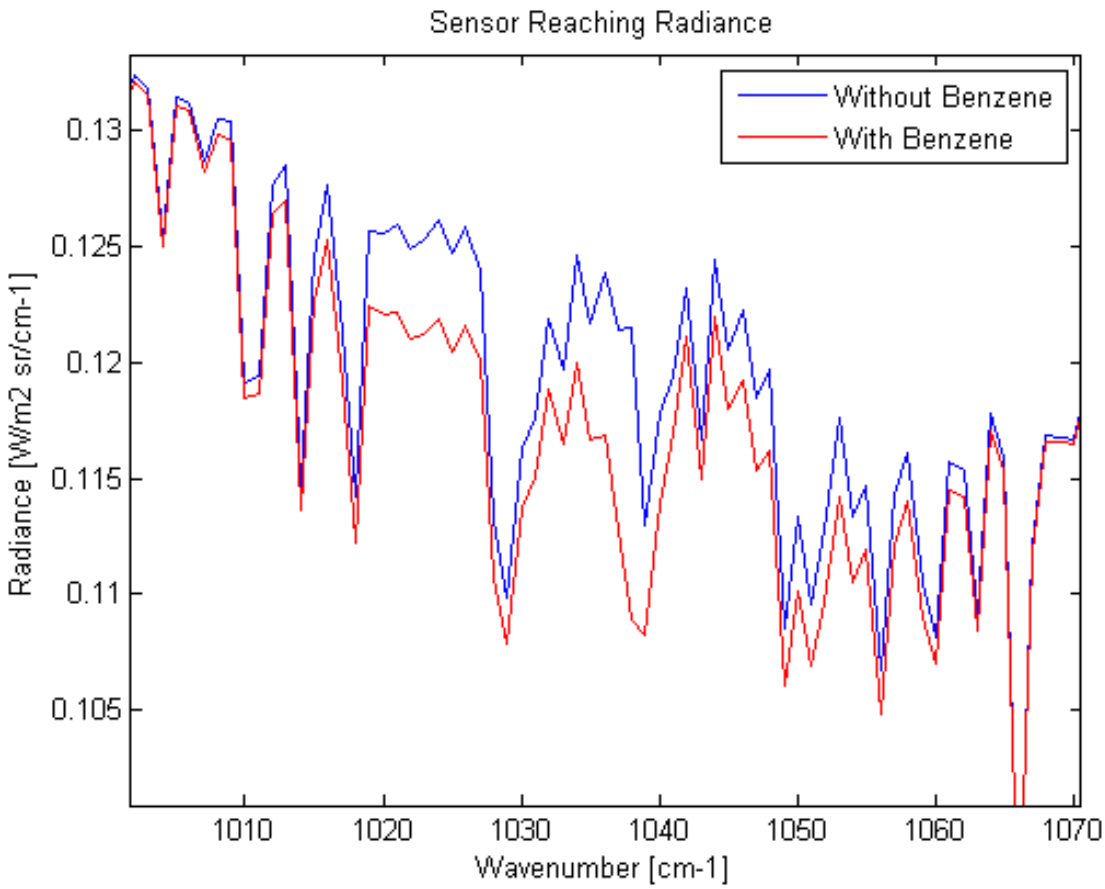

FIGURE 3.20 BENZENE AND OZONE COMPARISON SPECTRA

Figure 3.21 shows the next gas, methyl chloride.

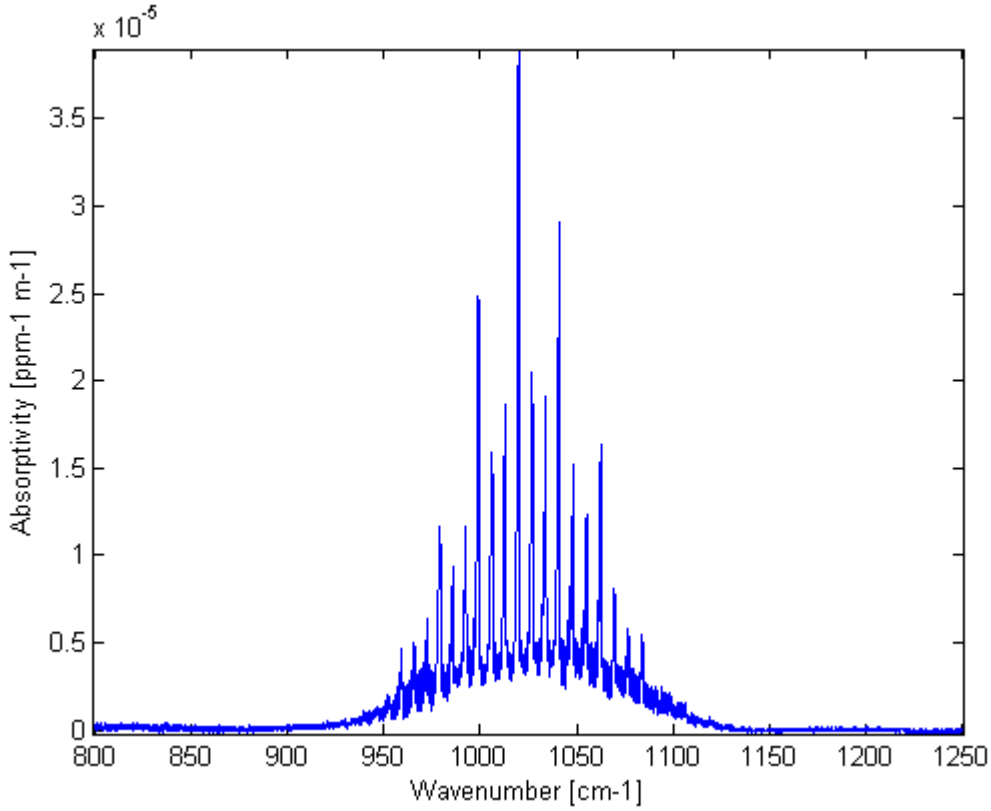

FIGURE 3.21 ABSORPTIVITY OF METHYL CHLROIDE

Methyl chloride made for an interesting target as well because of its repeating pattern of essentially the same spectral feature in this rotational manifold. The resulting interferogram shows regular repeating intervals that correspond to the presence of this gas, which made for a good candidate in understanding the effects of jitter in an interferogram.

Figure 3.22 shows the same format of interferogram and corresponding spectra series by maximum OPD as shown previously. This series is of the absorptivity spectrum of methyl chloride gas illustrating this point, note the differences in scale.

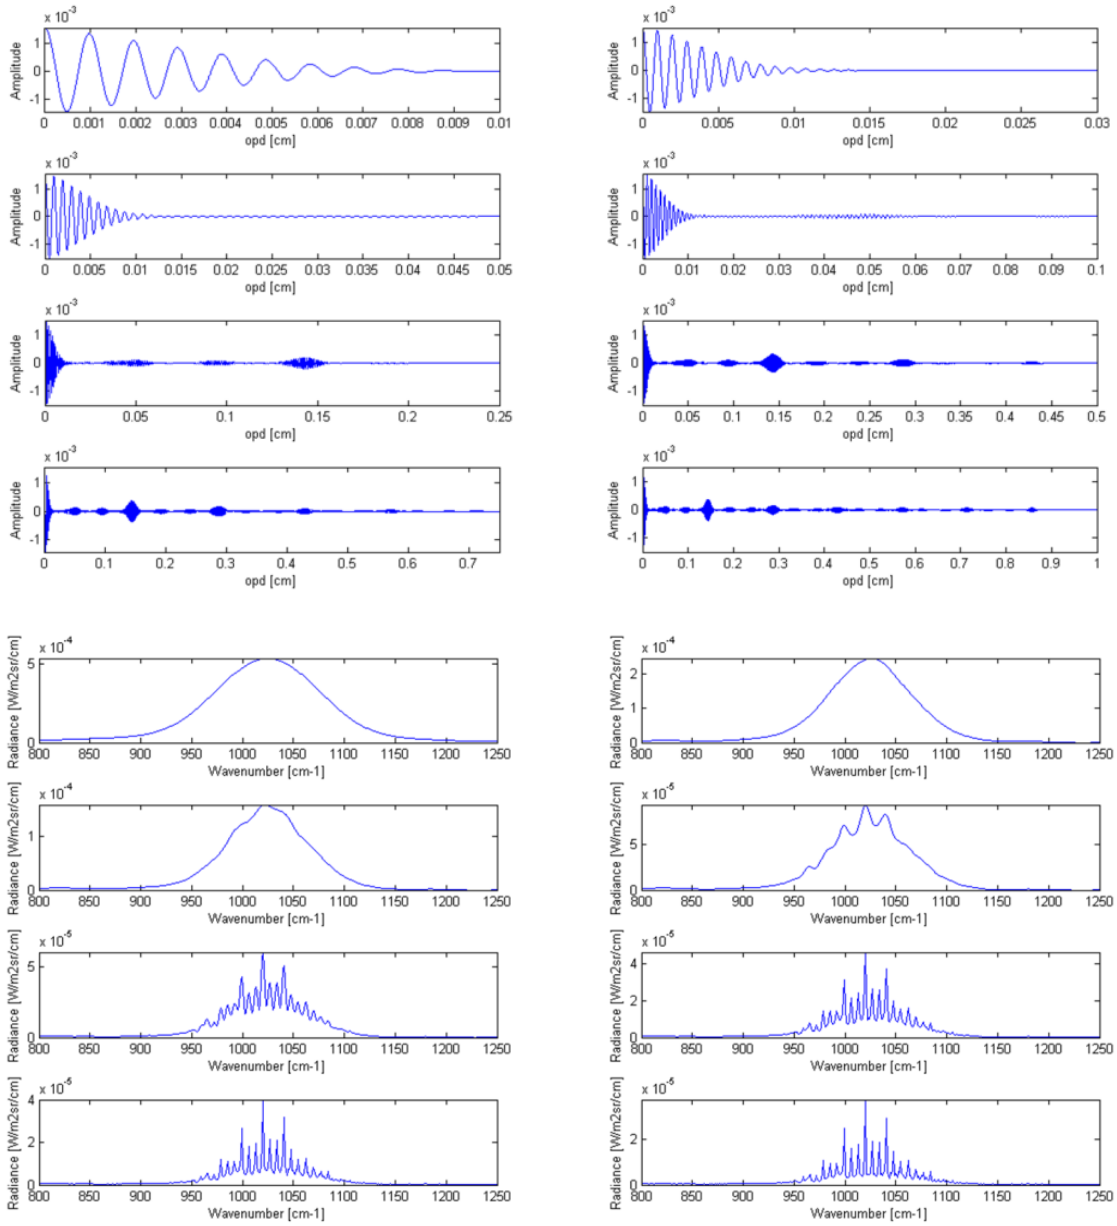

FIGURE 3.22 INTERFEROGRAM (TOP) AND CORRESPONDING SPECTRA (BOTTOM) OF ABSORPTIVITY SPECTRUM OF METHYL CHLORIDE

The significant low frequency contribution comes from the overall Gaussian-like base to the rotational manifold, while the periodic nature of the interferogram features in panel five give rise to the multiple peaks of similar shape.

In addition, methyl chloride was chosen as the confuser gas to be present when detecting benzene. This was due to the similarity and proximity of band heads in the two spectra. While bezene's Q branch occurs at 1038 cm^{-1}, methyl chloride has two neighboring peaks at 1034 and 1041 cm^{-1}, which are easily distinguishable at 1 cm^{-1} resolution, but appear as one band at 5 cm^{-1}. Figure 3.23 shows an absorptivity normalized plot of benzene and methyl chloride, illustrating the proximity and similarity of the peaks.

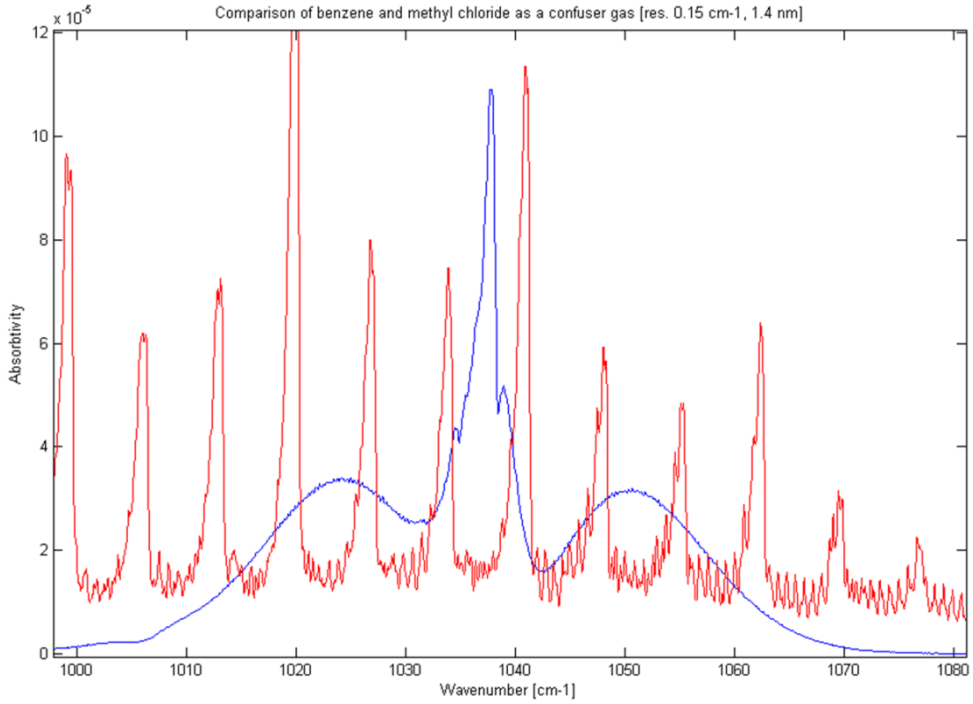

FIGURE 3.23 ABSORPTIVITY NORMALIZED PLOT OF BENZENE AND METHYL CHLORIDE

Figure 3.24 shows the absorptivity spectrum of ammonia.

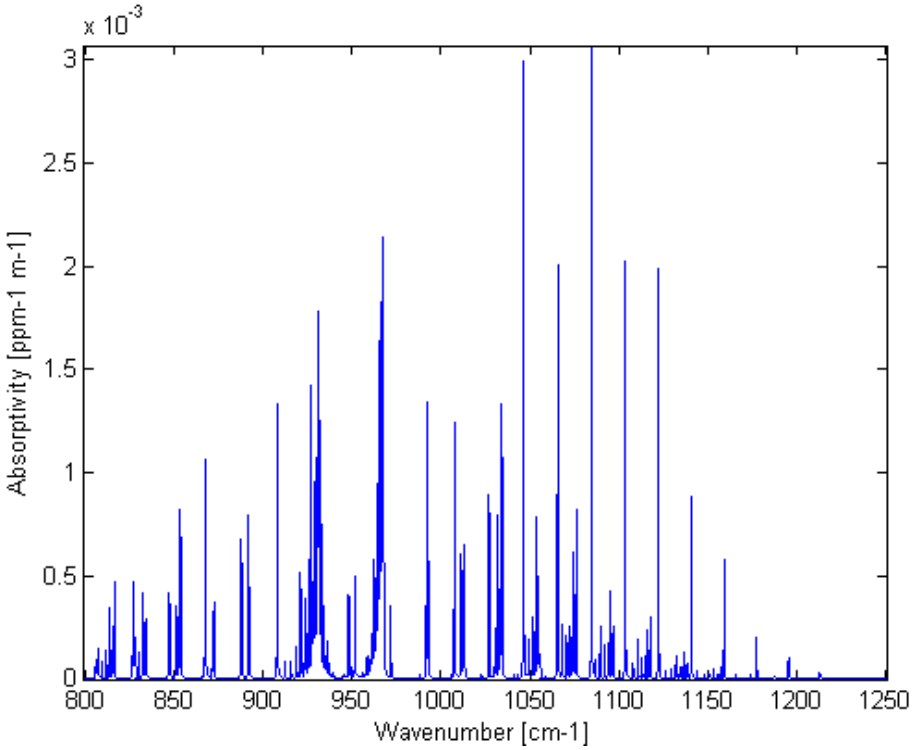

FIGURE 3.24 ABSORPTIVITY OF AMMONIA

Ammonia is a mainstay gas of the gas detection community and is used in many studies (Bennett, Carter, Fields, & Lee, Infrared Hyperspectral Imaging Results from Vapor Plume Experiments, 1995), (Gallagher, Wise, & Sheen, 2003), (Kroutil, Knapp, & Small, 2002). It was chosen for both the significant variability in spectral shape of each of the bands as well as the presence of the bands that span the bandwidth of the model. These two attributes also combine to ensure evidence of the gas will be present not only across the entire interferogram, but also in varying form (reference Figure 2.17). The most common spectral features used in identifying ammonia are the large bands at approximately 930 and 960 cm^{-1}. In fact, one group has used ammonia as a confuser gas of sorts in concert with SF_6 detection at 948 cm^{-1} (Farley, Chamberland, & Villemaire, Chemical Agent Detection and Identification with a Hyperspectral Imaging Infrared Sensor, 2007).

Figure 3.25 shows the absorptivity spectrum of sulfur hexafluoride.

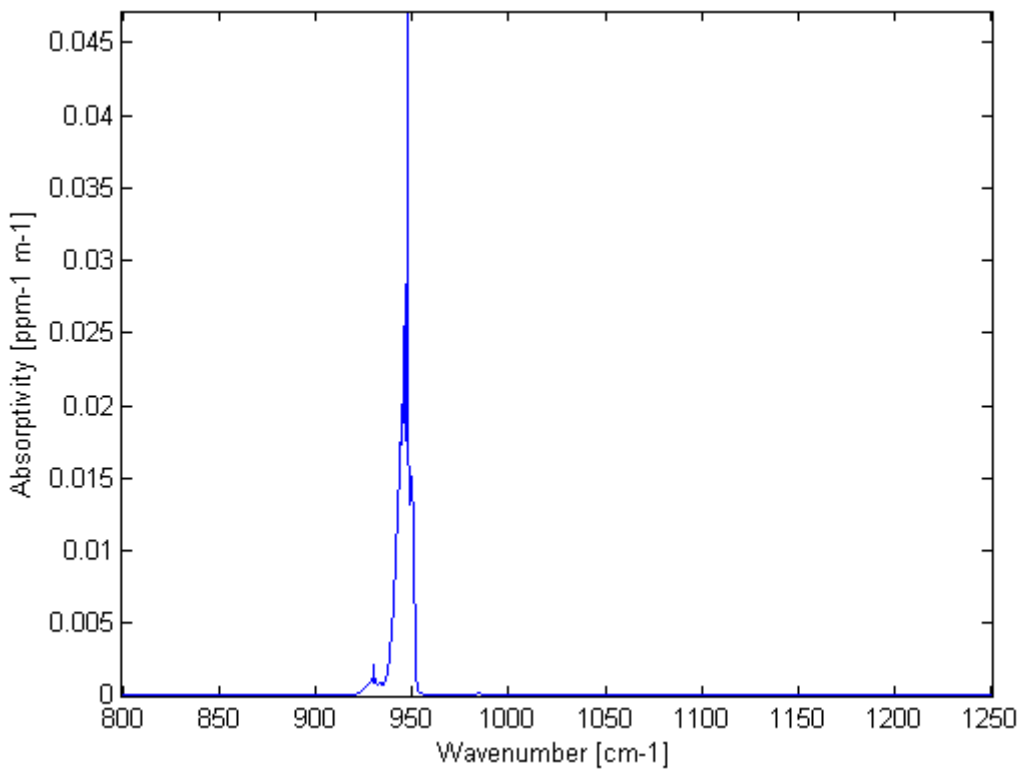

FIGURE 3.25 ABSORPTIVITY OF SULFUR HEXAFLUORIDE

Sulfur hexafluoride is another gas commonly used in field experiments for gas detection. This is in part due to the benign effect on the environment, but also due to its extremely large absorptivity and sharp spectral feature. SF_6 is not of incredible interest to this work as compared to the aforementioned gases, but was included for completeness and for comparison opportunities with prior work. While it may be difficult to ascertain from the figure, SF_6 actually has quite a bit of fine spectral detail that ride on top of the overriding bandshape. A brief study was conducted in looking at the effects of apodization, which primarily affects high frequency detail, on the loss of the fine spectral detail in SF_6 and in being able to recover evidence of this gas with a corrupted

interferogram by removing apodization. Figure 3.26 shows the interferogram and corresponding spectra series for SF_6.

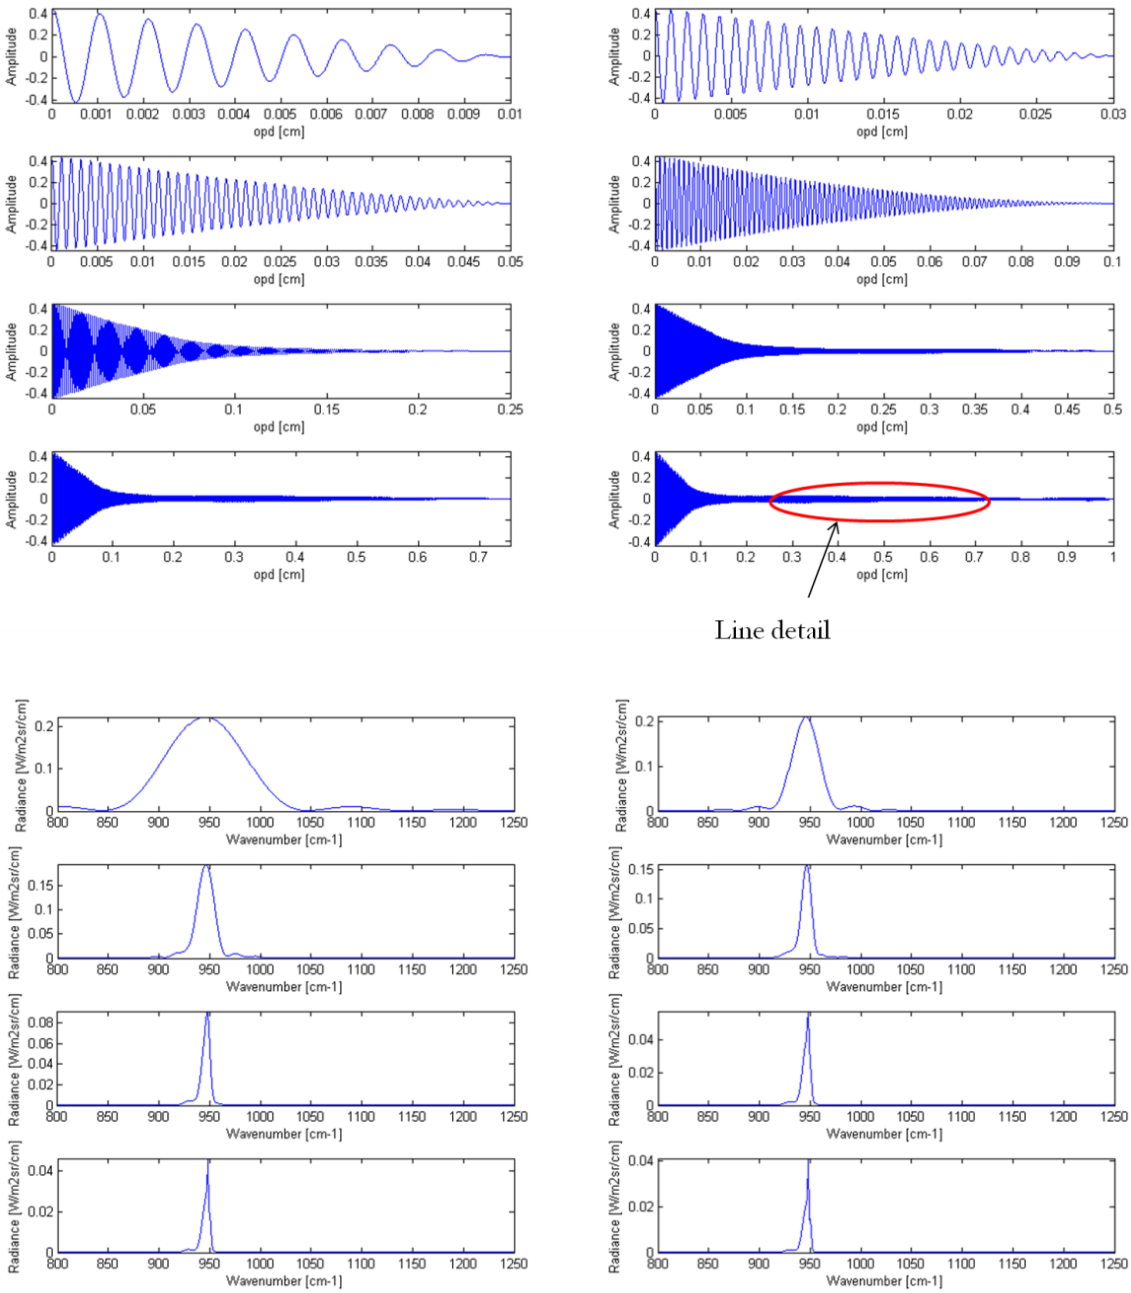

FIGURE 3.26 INERFEROGRAM (TOP) AND CORRESPONDING SPECTRA (BOTTOM) OF ABSORPTIVITY SPECTRUM OF SF_6

The series clearly shows that the majority of the SF_6 spectral feature is comprised of a single band of low frequency content and is the critical part in detecting this gas. However, as the annotation on the figure indicates, there is much fine spectral detail that may represent only a very small contribution to the interferogram, but can aid in the detection of this gas should the initial portion of the interferogram become corrupted.

Figure 3.27 shows the absorptivity spectrum of sulfur dioxide, which is a good example of a P and R branch manifold with very clear evidence of the rotational lines, with the notable absence of the Q branch.

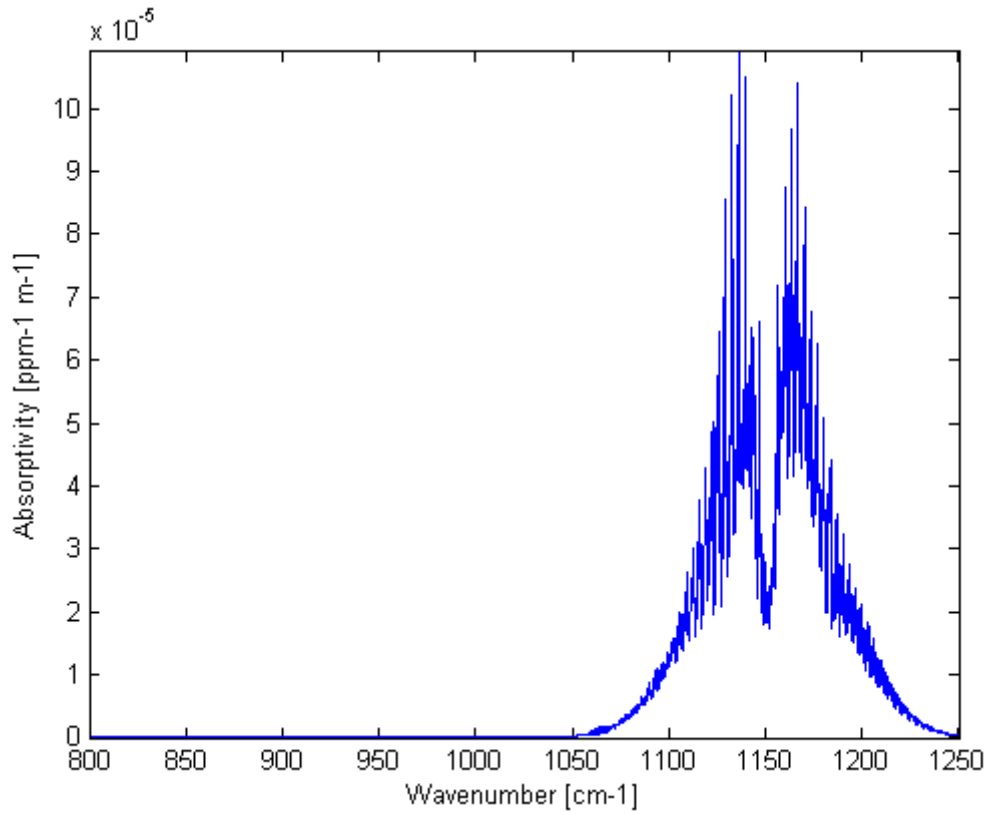

FIGURE 3.27 ABSORPTIVITY OF SULFUR DIOXIDE

SO_2 was included in this work due to the level of interest in the community associated with both natural and man-made sources of the gas and its impact on the environment. In particular, SO_2 content is associated with the venting of volcanoes and a comparison to this data collected by AHI was made.

The interferogram for this species is similar to methyl chloride, but because the rotational lines are so finely spaced, they are somewhat lost at 1 cm^{-1} resolution and merge together to form irregularly shaped spectral features. Thus the regular, periodic interferogram one would expect due to the manifold of rotational lines is not realized. Figure 3.28 shows the interferogram and spectra series for this gas.

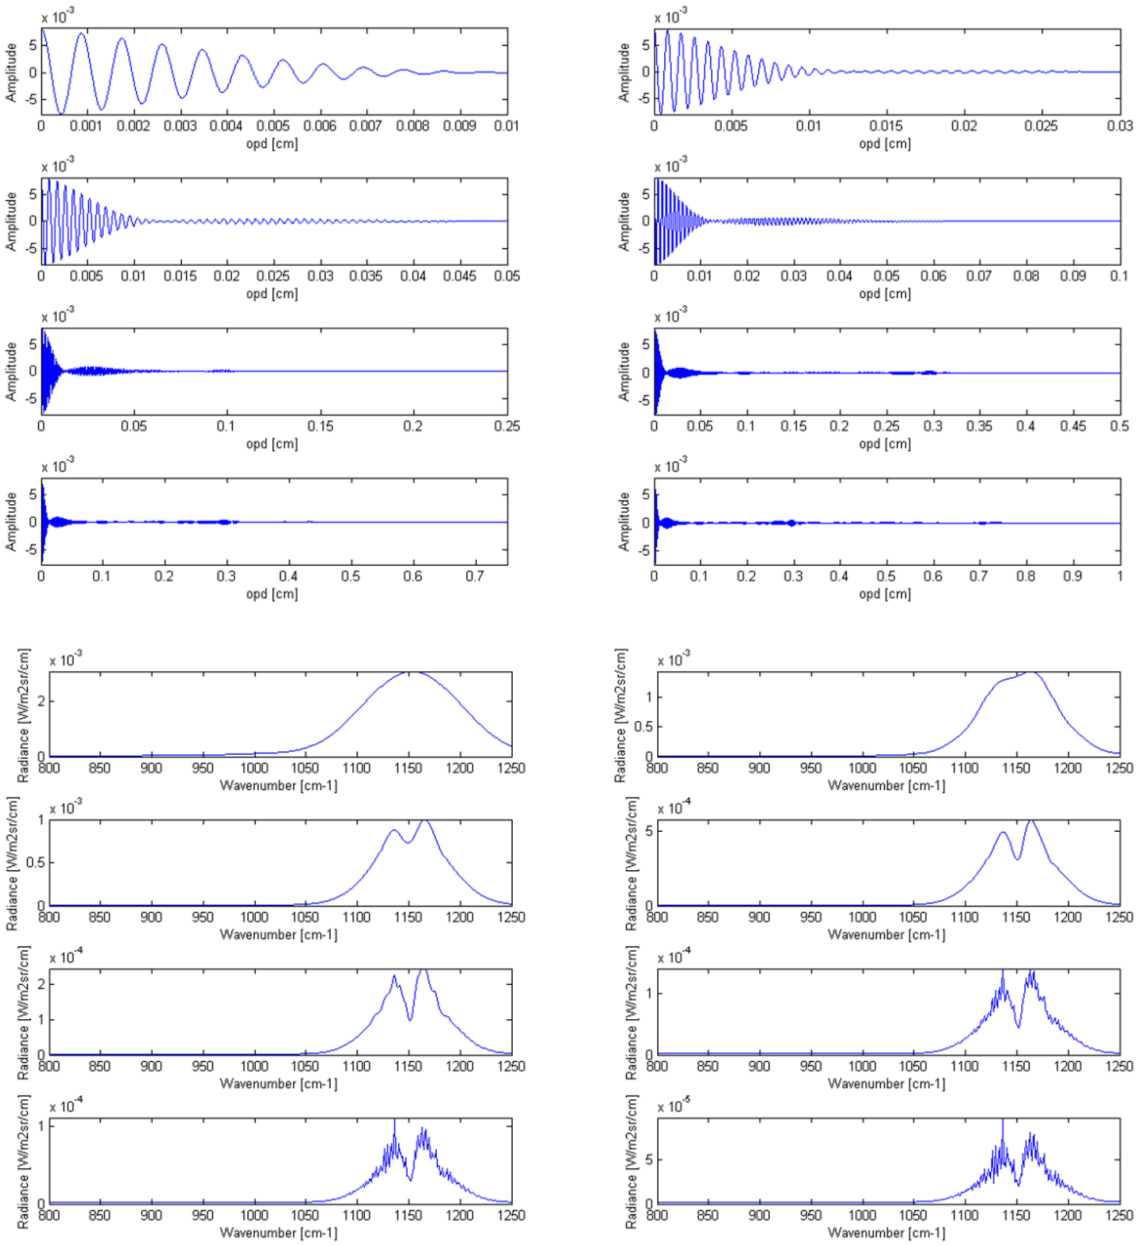

FIGURE 3.28 INTERFEROGRAM (TOP) AND CORRESPONDING SPECTRA (BOTTOM) OF ABSORPTIVITY SPECTRUM OF SO_2

Phosgene was included in this work as an example of a gas employed as a chemical weapon – albeit in WWI. It is similar to benzene in both spectral shape and size, but the P and R branches are far more prominent, almost to the detriment of the Q branch. Figure 3.29 shows the interferogram and spectra series for this gas.

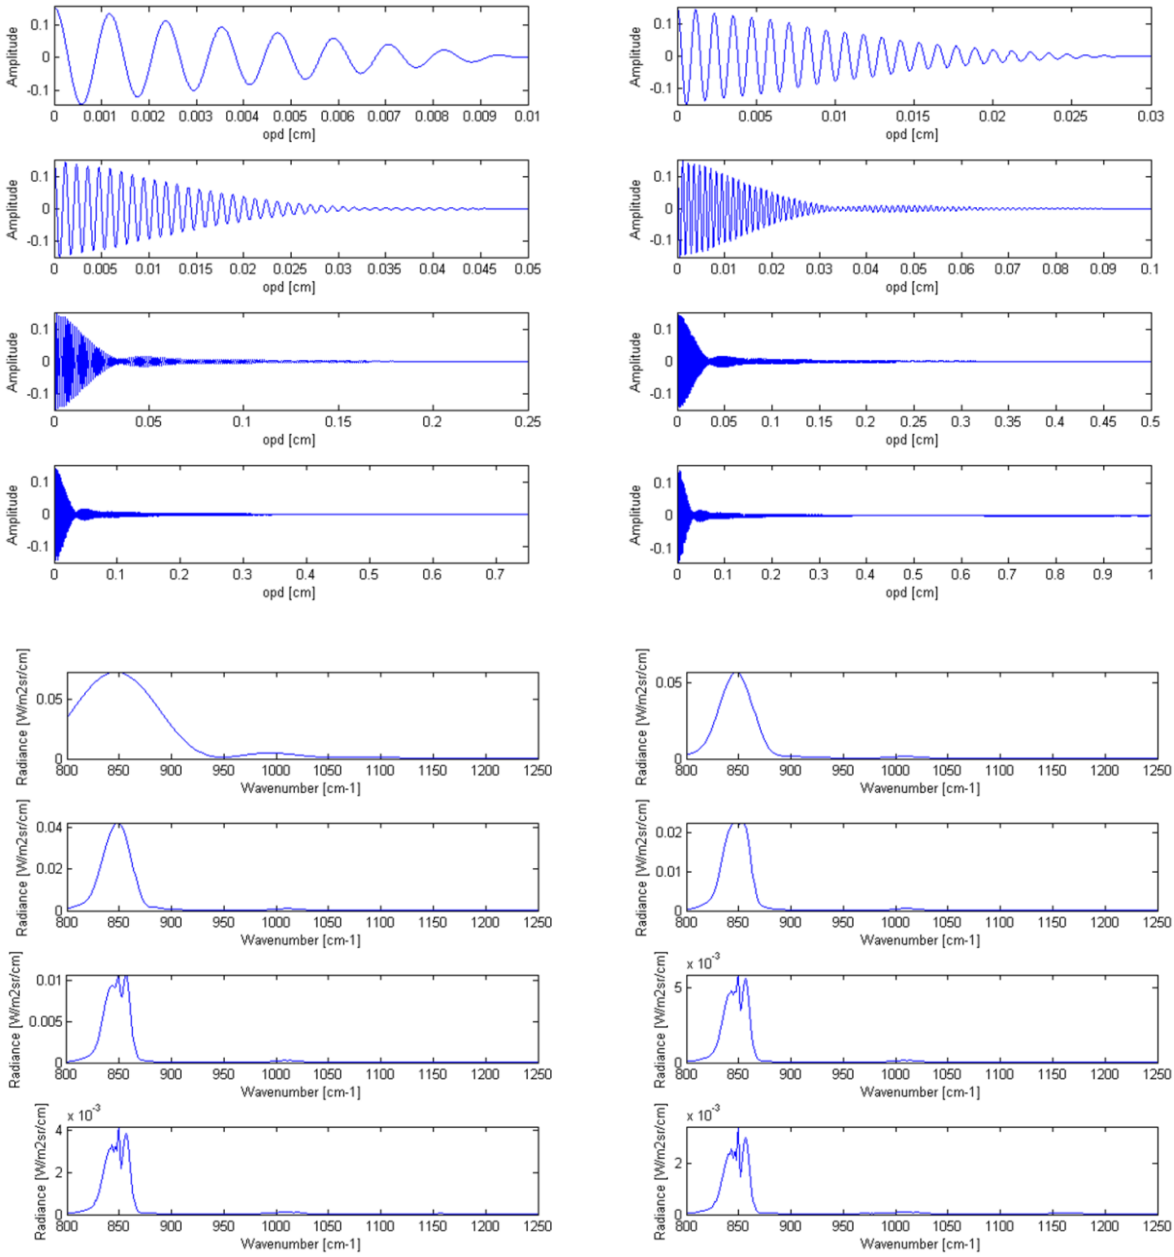

FIGURE 3.29 INTERFEROGRAM (TOP) AND CORRESPONDING SPECTRA (BOTTOM) OF ABSORPTIVITY SPECTRUM OF PHOSGENE

The final gas used in this work was carbon tetrachloride. This gas was chosen for its common place use in urban settings as a degreaser and previously, a dry cleaning agent, as well as its position in the band of interest. The tail end of the spectral feature starts right at the beginning of the spectral window modeled in this work, which may have posed some detection problems, particularly at poor resolution. Figure 3.30 shows the absorptivity spectrum of CCl_4.

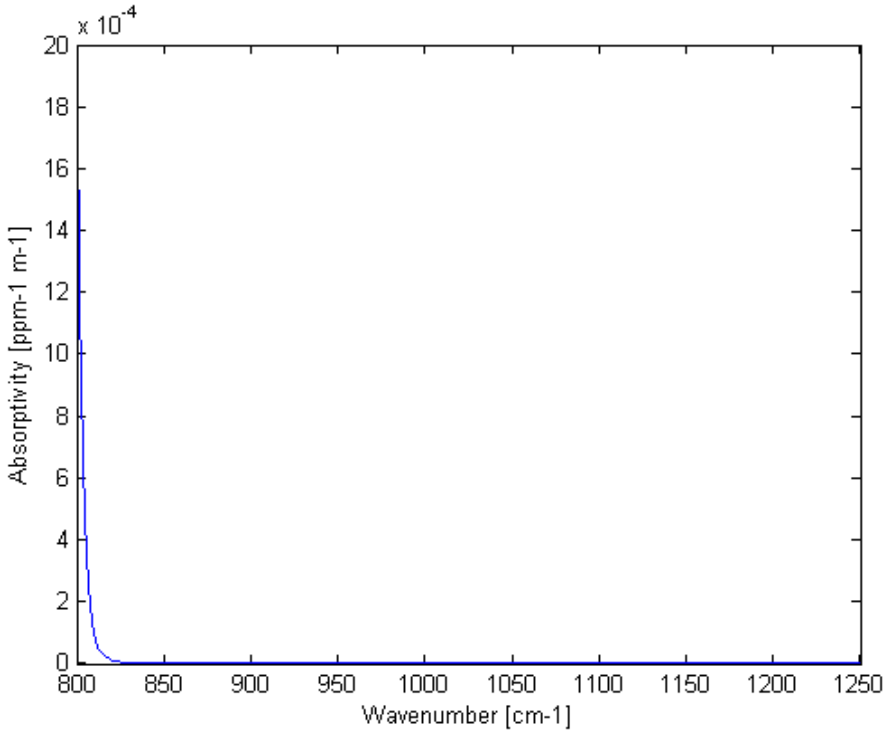

FIGURE 3.30 ABSORPTIVITY SPECTRUM OF CARBON TETRACHLORIDE

The interferogram for CCl_4 is uninteresting and essentially consists of the low frequency content that makes up one half of the singular peak.

A brief point must be made about the absorptivities of the target gases. Table 3.1 contains a list of the absorptivities of each of the gases and in some cases the values for two gases may differ by several orders of magnitude. The absorptivity of the gas is a relative comparison of the 'strength' of a given absorption feature at a specified concentration, pathlength, and temperature of the gas. For example, at a fixed thermal contrast, a 3,000 ppm-m sample of ammonia will yield approximately the same level of absorption as a 6 ppm-m sample of SF_6 at their respective peak wavelengths.

In order to compare the detection performance across the range of gases, it became important to find a way to normalize the absorption 'strength' among them to within the same ballpark. The 'strength' is not only dependent on the absorptivity, concentration and pathlength of the gas, but is also dependent on the thermal contrast between the gas and the background. Since an increase in the magnitude of any of the variables results in an increase to the strength of the absorption and each variable has approximately equal influence, it seemed logical to propose a method of normalization that is the product of all of these factors

$$Strength = \kappa \cdot c \cdot l \cdot \Delta T \qquad (3.6)$$

where all variables are as previously defined. This method is more of a rule of thumb as this relationship is non-linear, particularly as the thermal contrast approaches 0. This method would be

implemented when running a scenario with two different gases. Generally, two of the three variables would be held constant across the gases and the absorptivity value would be derived from the maximum absorptivity value for the gas. Setting the strength of the second gas equal to that derived from a specified concentration, pathlength, or thermal contrast of the first gas and solving would yield the variable value to use for the second gas giving an 'equivalent' absorption strength.

Figure 3.31 shows an absorptivity normalized plot consisting of benzene, ammonia, and sulfur hexafluoride. The normalization allows for the full range of absorptivity to be shown simultaneously for each gas and gives an idea of what strength-normalized gases might look like. This plot allows for direct comparisons of the spectral features and line widths associated with each gas, but cannot be used to make inferences about their relative absorptivities. The gas shown in black centered at 947 cm^{-1} is SF_6, blue is benzene at 1038 cm^{-1}, and the remaining lines in red are due to ammonia. As mentioned previously, many research groups use the major spectral band heads from SF_6 and ammonia near 930, 947, and 960 cm^{-1} as a measure of their sensor/algorithms ability to discriminate between spectral features leading to identification of gas constituents in a scene.

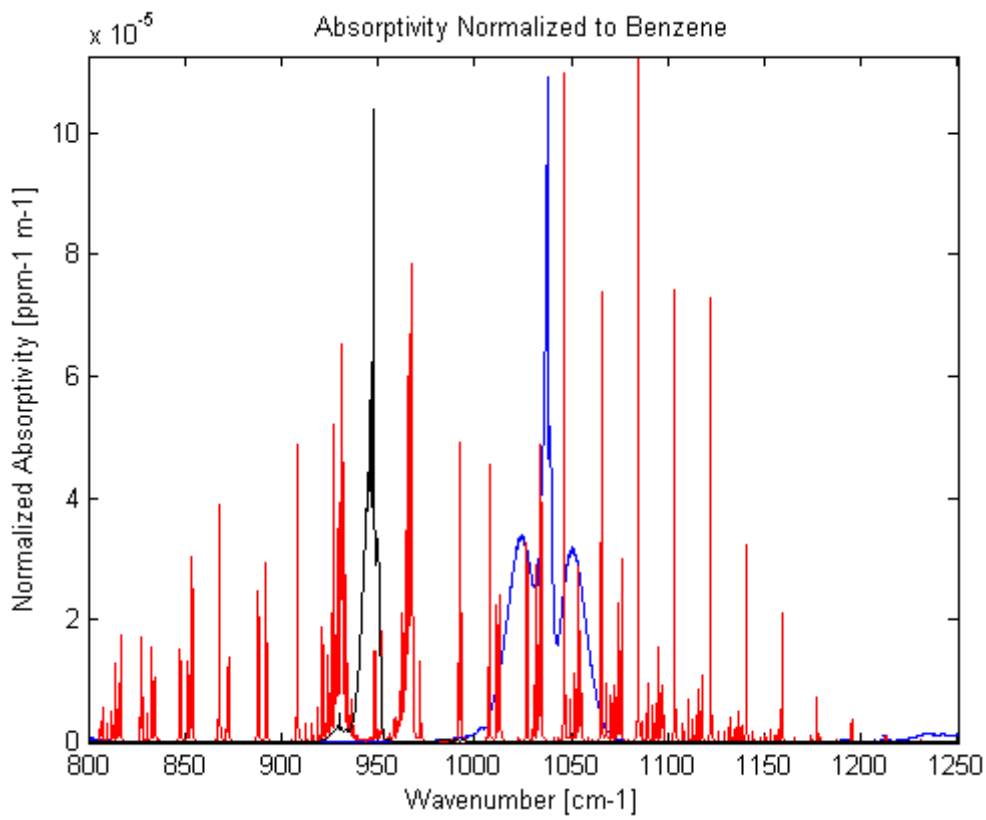

FIGURE 3.31 ABSORPTIVITY NORMALZIED PLOT OF NH_3 (IN RED), BENZENE (IN BLUE, PEAK 1038 CM^{-1}), AND SF_6 (IN BLACK, PEAK 948 CM^{-1})

Another important facet of dealing with gas absorption strength is in specifying the detection conditions of a gas. In order to accurately compare the success of a gas detection for a target vector of specified concentration-pathlength and thermal contrast, one must specify all of the information

in the strength relationship. This is because a given gas at a comparatively low concentration-pathlength with a high thermal contrast will have a very similar detection score to one where the inverse is true for a given target vector. Figure 3.32 is of a detection plane that illustrates this point yielding similar scores across a variety of values for both concentration-pathlength and thermal contrast for a given target vector.

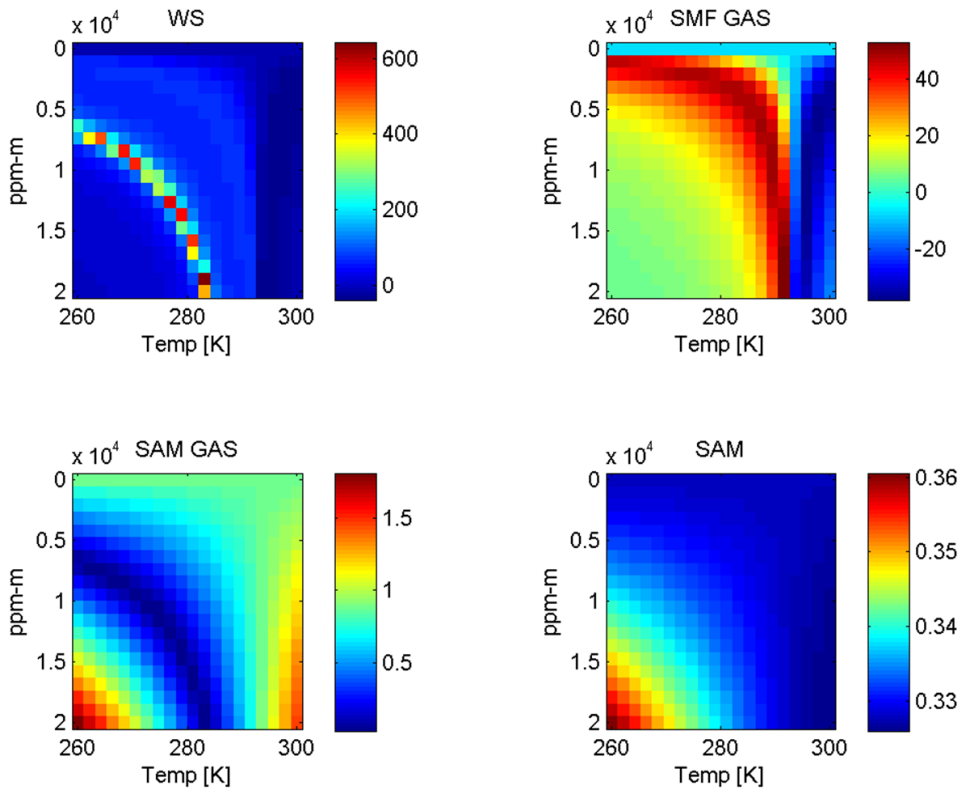

FIGURE 3.32 DETECTION PLANES ILLUSTRATING ABSORPTION STRENGTH DEPENDENCIES

Figure 3.32 Detection Planes Illustrating Absorption Strength DependenciesFigure 3.32 shows the results of several detection metrics (introduced later) across changing concentration-pathlength values and gas temperatures. A clear pattern evolves that shows equivalent 'detectability' at multiple combinations of the strength parameters. The relationship basically shows that high thermal contrast can make up for a low concentration of gas or vice versa and that nearly equivalent spectra result. The model resolution was 1 cm^{-1} and no atmosphere was applied in this scenario, yielding essentially gas leaving radiance. The background was at 300K and the transition to emission can be seen beginning after a gas temperature of approximately 290K. The gas target vector used in the detection was at 4,800 ppm-m and 268K over a 300K background. The comparatively cold temperature was used to get the most detail out of the spectral feature.

Gaussian Modeling of Gas Plumes

With the gas leaving radiance in hand, it is now important to determine how the radiance varies over location in a scene. Gas plumes were initially created with Gaussian models for the spatial distribution of the gas from a point source. The model was further modified to simulate dispersion

of the gas molecules by an essentially reciprocal relationship with distance from the source and a nearly linear increase in path length with increasing distance from the source. Finally the plume temperature was thermalized from the point source and out by a nonlinear increase in temperature. This implies that the plume is in absorption due to a likely scenario where a pipe carrying gas at near atmospheric temperature with some positive overpressure gets a leak and causes the gas to rapidly expand to atmospheric pressure with rapid cooling upon expansion at the point of release. As the plume moves beyond the release point, it rapidly thermalizes to the atmospheric temperature. The Gaussian model used for the spatial distribution of the gas is given by the probability density function (PDF) modified as a summation over all lines in the scene,

$$\sum_{i=1}^{I} \frac{cpl_{max}}{\sigma(i)\sqrt{2\pi}} e^{\frac{-(x-\mu)^2}{2\sigma(i)^2}} \tag{3.7}$$

where μ and σ are the mean and standard deviation, respectively and *cpl* is the concentration pathlength product. The values for *x* ranged from -0.49 to 0.5, the mean was 0 and the standard deviation ranged from 0.1 to 0.3 in .002 increments as the distance increased, but is dependent on the desired plume shape. This allowed for an increasing breadth to the plume with increasing distance. The PDF values were derived for each line of a 100x100 scene from the point of plume origination. As the line value increased away from the plume, the standard deviation of the distribution increased. This model provided the cpl map for the scene. Figure 3.33 shows examples of the types of plumes generated by this model by only varying the standard deviation.

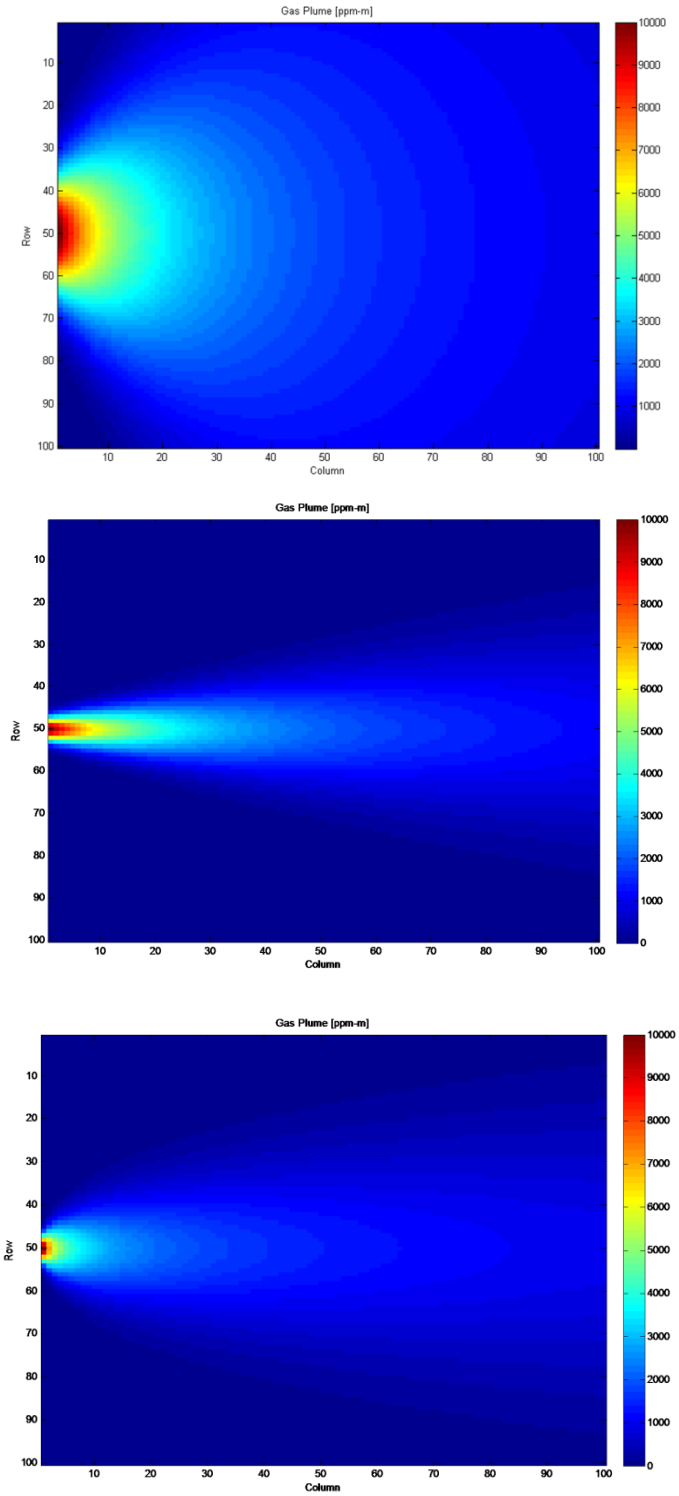

FIGURE 3.33 GAS PLUMES GENERATED BY GAUSSIAN MODEL

The plumes shown were created by varying the standard deviation profile implemented in the model. The images show the spatial distribution of the concentration pathlength product of the gas within the plume.

The plume temperature profile was modeled by the following

$$T(x,y) = T_{min} + M\left(\frac{cpl_{max}}{cpl(x,y)}\right) \tag{3.8}$$

where *T* is the gas temperature, *M* is a coefficient of thermalization, and the *x*, *y* are the coordinates in the scene. *M* essentially controls the extent of the temperature range from source to the end of the plume – the larger the value, the larger the range of temperatures. A value of 8 was used in this work. Once a minimum gas temperature was determined the value of the T_{min} term was derived by subtracting the *M* value from the minimum desired temperature. Since the concentration is highest at the source, this equation yields the coolest temperature there. Figure 3.34 is a close-up view of the distribution of temperature inside the inset shown from the first plume from Figure 3.33 as predicted by the model over the first half of the plume extent.

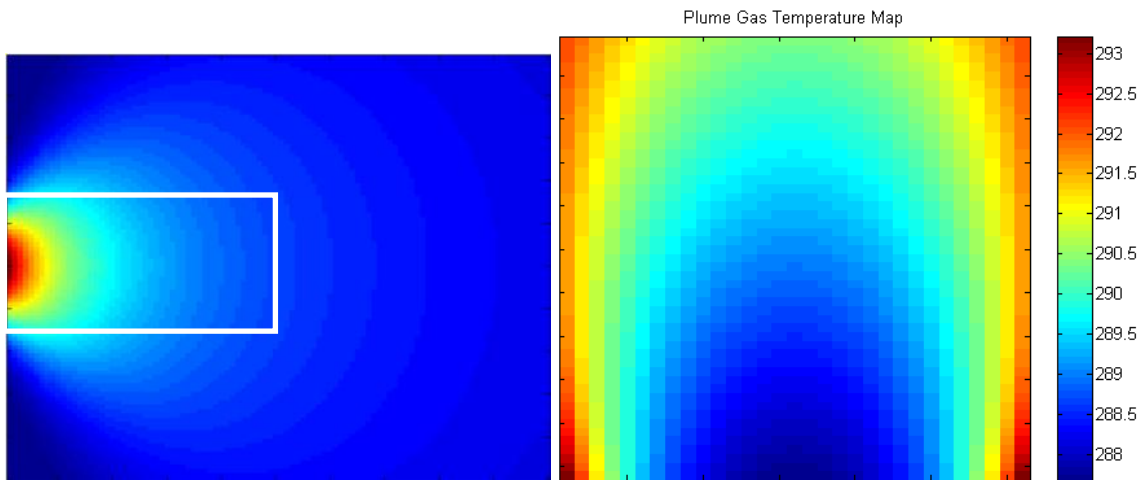

FIGURE 3.34 GAS PLUME TEMPERATURE DISTRIBUTION (RIGHT) OF INSET SHOWN AT LEFT

Atmosphere Model

While it is now clear where the sources of radiance from the background and gas plume originate, that energy must be transported through the atmosphere to the sensor aperture. Without transitioning to another model such as HITRAN, the resolution of the MODTRAN v4r3 model extends to 1 cm^{-1} and is the resolution limiting component of the overall FTS model. However, producing spectral data cubes at 1 cm^{-1} resolution is a lofty goal for a scanning sensor in itself and does not need to be further stressed. Additionally, there were no spectral features in this work that required finer resolution to distinguish gaseous species.

In order to save on execution time, four standard MODTRAN atmospheres were run and are read in to the model when needed. They consist of both 1 and 5 km MSL altitude atmospheres using the MLS and Tropical standard models. The user can select no atmosphere, or a transmission only atmosphere, removing the effect of either or both up and down welled radiance.

In order to generate these atmospheres, MODTRAN was run at nadir, between the ground at 0.079 km and the flight altitude in MSL over Rochester, New York with both the sampling and instrument

FWHM at 1 cm^{-1}. The surface albedo was set to 1, simulating a perfectly reflective surface over the spectral bandwidth and eliminating surface emission. This surface allows for substitution of the customized background radiance profiles to be used in place of the standard MODTRAN files when running the model. While the down-welled radiance profiles are dependent on the surface material albedo, they do not contribute enough to the overall SRR to justify running MODTRAN repeatedly for each surface material at varying temperatures and composition at each ground position. In this manner, a generic down-welled radiance profile was generated at the largest presence it would have and then it was modulated by the reflectance of whatever surface material was present at a given ground location in the model. This method causes very little error in the down-welled contribution (because it is so small to begin with), while the overall accuracy of the MODTRAN output is far more dependent on variable conditions such as temperature and component concentrations that the entire output is a coarse estimate at best, regardless of down-welled contribution.

Figure 3.35 shows a comparison of the down-welled radiance using the default desert albedo file in MODTRAN and the down-welled radiance estimate from the model for sand used in this work. The model uses the down-welled radiance from MODTRAN from a perfect reflector and modulates it with a derived emissivity file for sand from a measured spectrum taken from a desert. The two vertical axes show the down-welled radiance spectra as well as the albedo profile used for each material.

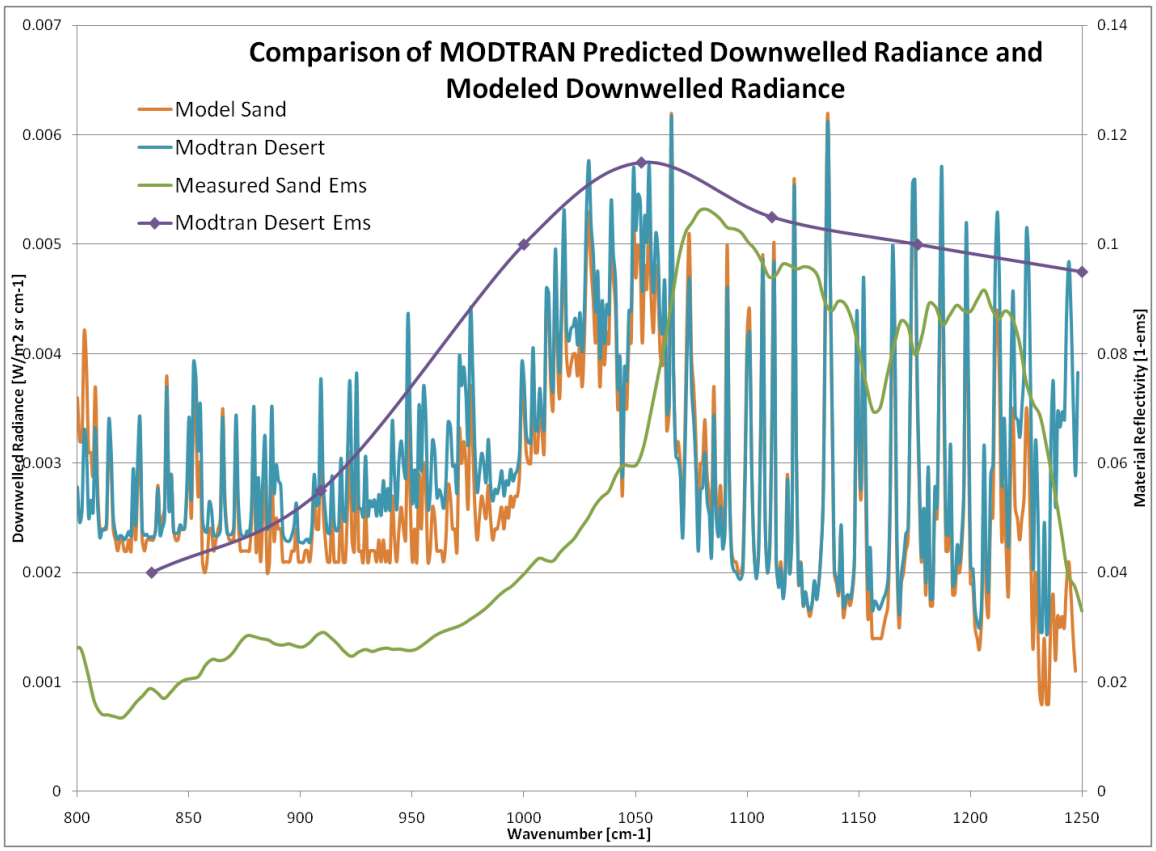

FIGURE 3.35 COMPARISON OF DOWNWELLED RADIANCE FROM MODTRAN AND MODEL

Note that the surface albedo files are very similar with the exception of the last 30 wavenumbers and that the MODTRAN surface albedo file is only represented by six measurements. The down-welled radiance spectra are nearly identical when the deviations from the differences in the surface albedo values are accounted for. Typically, measured total radiance values over the bandwidth of interest average to 0.15 W m^{-2} sr^{-1} cm, meaning that the down-welled contribution is on the order of 2% of the total radiance. Thus, this method of modeling down-welled radiance results in a substantial savings in time and computational complexity with virtually no loss in radiometric accuracy.

Up-welled radiance is essentially altitude dependent and a standard MODTRAN up-welled radiance profile is used for each of the four atmospheres used in this work. Up and down-welled radiance profiles used in this work can be seen in Figure 2.9.

3.1.3 Spectral Output

It is of some importance to point out that the spectral output of the model is inherently in wavenumber. This is due to the fact that spatial sampling occurs on the order of centimeters and in frequency space that becomes wavenumber, or reciprocal centimeters. It turns out that converting wavenumber plots to those in the wavelength domain is not straight forward. Wavelength is inversely proportional to energy as shown in the following development

$$E = h\nu \tag{3.9}$$

$$c = \lambda\nu \tag{3.10}$$

and by substitution

$$E = \frac{hc}{\lambda} \tag{3.11}$$

where h is Planck's constant, c is the speed of light, and ν is frequency in Hertz. Wavenumber on the other hand, is directly proportional to energy as shown here

$$E = hc\tilde{\nu} \tag{3.12}$$

This difference in the relationship to energy means that spectra in wavespace contain equal quanta of energy between wavenumber intervals whereas in wavelength, each successive interval contains less energy. Thus, displaying a spectrum in wavelength would exaggerate the extent of spectral features on the long wavelength side and compress features on the short wavelength side. Figure 3.18 shows an example of the dramatic differences between spectra displayed in equal energy intervals and that in reciprocal energy intervals.

When converting any single wavenumber value to wavelength, one can simply divide that number into 10,000, a conversion from centimeters to microns. To convert a band of values, one must also consider the energy content of the interval between each of the values being converted and scale the result by the energy interval. To this point, a radiometrically accurate spectrum can be converted from one domain to another; however, the nonlinear interval spacing in the wavelength space causes problems in maintaining spectral resolution. For example, if one is trying to match the resolution of a spectrum taken by a grating spectrometer with one from this FTS model, you might choose a few representative peaks that would give a FWHM measurement to set the resolution in the model. However, in doing this, one would find that for wavespace, the resolution is constant across the spectrum and in converting to wavelength, this results in spectral feature mismatches in the DES derived spectrum. It will appear that while spectral features may match in the center of the spectrum, those to the short wavelength side will appear over resolved while those on the long wavelength side will appear under resolved. This in part has to do with some non-linear traits in the dispersion of grating spectrometers, but in a sampling sense, the wavelength spectrum must be converted from energy values that correspond to equal wavelength spacing as opposed to equal energy spacing. To do this, one must interpolate values in a sampled wavenumber spectrum that correspond to equal wavelength spacing before conversion. As such, this introduces some error into the conversion process. Figure 3.36 shows the non-linear relationship between wavenumber interval and wavelength intervals.

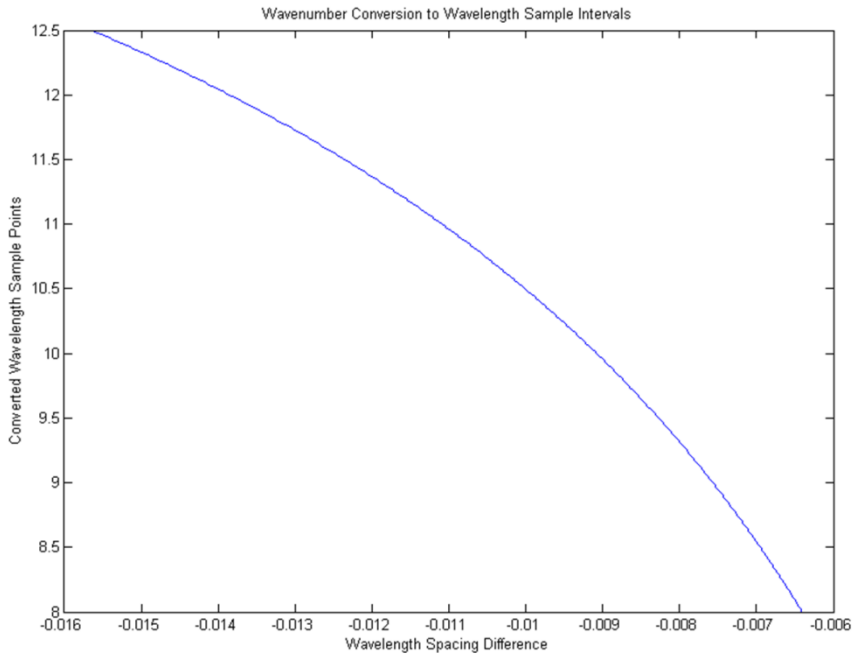

FIGURE 3.36 WAVENUMBER CONVERSION TO WAVELENGTH SAMPLE INTERVALS

The vertical axis shows equal spacing in wavelength values over the band of interest, while the horizontal shows the corresponding wavelength interval difference in equal energy space. As the figure shows, long wavelength intervals contain less energy and require longer intervals with increasing wavelength to maintain equal energy spacing.

As a final note, one can generally obtain satisfactory conversions from spectra in wavespace to wavelength, but going the other way causes changes in apparent resolution and spectral feature shapes that will not match an FTS output.

3.1.4 Detection Metrics

This section details the detection metrics employed for quantification and comparative measurements of change and degradation in target detection performance when target plumes are subjected to various environmental and instrumental effects. Detection metrics rooted in statistical and geometric target detection were the best fit as SNR-based metrics such as Minimum Detectable Quantity (MDQ) do not really fit as no element of noise or detector response is used in the simulation. Further, no comparisons are being made between detection systems for a given gas, which is essentially the purpose of MDQ.

In addition, no attempt at atmospheric removal or modeling was used to improve the detection likelihood, rather the transformed spectra were run against the metrics directly. This was done as a portion of this work was dedicated to measuring the impact of the atmosphere on gas detection in the various collection scenarios.

Much was asked of the detection metrics used in this work. They must accurately distinguish between minor changes in spectral features due to instrument error or changes in atmosphere

and/or background. They must be relatively magnitude independent as the target vector will never be an exact match, and often, depending on the gas temperature and background temperature, the slope of the bias that the spectral feature rides on may be different from that in the target vector. The metrics cannot be confused by encroaching atmospheric features such as ozone and need to work under a variety of spectral resolutions, while providing meaningful measurements in both emission and absorption.

It was found that one detection metric alone was not sufficient in characterizing the detectability of the gas under varying conditions. In fact, accurate, repeatable detection scores were not possible until the combination of three different detection metrics was employed. The three metrics selected were Spectral Angle Mapper, Spectral Matched Filter, and spectral depth.

In developing a target vector to feed the metrics, the goal was to use realistic scenarios. Thus the target vector was produced with a single thermal contrast and cpl over a black body background at the highest resolution of the instrument. This single target was used in the detection of the gas under a variety of conditions, cpl, and instrument resolutions as well as in emission and absorption. As in a real scenario, one would not know these conditions in looking for a gas leak and could only venture a best guess.

Spectral Angle Mapper

Spectral angle mapper (SAM) compares two spectral vectors by calculating the spectral angle that separates the overlap of the two spectra in vector space. The greater the overlap, the more similar the two spectra are and the smaller the spectral angle that separates them. It essentially consists of the dot product between a target vector and the spectrum derived vector. In this work the target vector is selected from the spectral feature ranges given in Table 3.1. The corresponding region from the simulated spectrum is used to compare to the target. The following equation illustrates how SAM is implemented in this work

$$\theta = \cos^{-1}[(S_t)^T(t)] \tag{3.13}$$

where t is the target vector, S is the simulated spectrum and the subscript signifies that the vector corresponds to the same spectral region as the target vector.

It was found that in general, when SAM was above four degrees, major differences in the spectral feature were not being characterized to the fidelity necessary to compare the detectability of a given gas under various conditions. SAM also has a narrow range of acceptable matches for the spectral feature shape, as one would expect, meaning it preformed best over a narrow range of conditions.

The metric defined in this work as 'SAM Gas' is the same SAM metric run over a smaller subset of the spectrum that corresponds to the immediate spectral region of the target. These target gas regions can be found in Table 3.1. To clarify, when the metric SAM is seen, it refers to the spectral angle between the entire transformed spectrum and either the SRR or a target vector, as appropriate, that comprises the entire spectral range. When the metric SAM Gas is seen, it refers to

the spectral angle between a subset of the entire spectrum that corresponds to the unique region that is comprised largely of the spectral feature. The SAM metric was used to characterize changes that would impact the entire spectrum such as atmosphere, resolution, background and FTS performance parameters, whereas SAM Gas characterized impacts on the detection of the target gas specifically.

Spectral Matched Filter

From the statistical approach to target detection, when the target is known, Spectral Matched Filter (SMF) is a good choice. SMF was implemented in this work as the following

$$SMF = \frac{(t-\mu_s)^T(S_{tgt}-\mu_s)}{\sigma_s^2} \qquad (3.14)$$

where μ is the mean and σ^2 is the variance of the spectrum and remaining variables are as previously defined. In this work, the entire spectrum was used for calculating the background statistics as the spectral feature generally only comprised a few percent of the spectrum. The idea of demeaning the target and spectrum is to remove the bias and place them on an equal plane, then the variance due to the spectral variability in the spectrum is used to normalize the terms by the spectral disparity present. Thus, if a spectrum has a wide range of spectral variability, it will reduce the SMF score for a given presence of target, essentially saying that the likelihood of the target being present is reduced by the possibility that the background is also target-like in character.

While SMF is more forgiving in detecting wide ranges of gas under varying conditions, mathematically, it is not really capable of discriminating fine spectral details. In fact, it is easily fooled by a large presence of spectral energy (from whatever source) in the same spectral band as the target. So a trade is made here as well for a greater chance of detection over a wider range of detection conditions as compared to the target vector, but giving up the ability to discriminate between the presence of, for example, ozone versus benzene.

The metric defined in this work as 'SMF Gas' again corresponds to a subset of the total spectrum that corresponds to a unique spectral region due to the presence of the target gas. SMF was not run against the entire spectrum in this work.

Spectral Ratio

To attempt to overcome the problem of the selective nature of SAM with the easily confused response of SMF, it seemed that a combination of the two in a manner that opposes the weaknesses of each would make sense. The Spectral Ratio is just the SMF Gas score over the SAM Gas score. This metric is referred to as the WS metric in all future detection plane images.

A truly good match from the gas plume to the target spectrum gives a high SMF score and low SAM score, which yields a very large number. In the presence of ozone, the SMF score is large, but the SAM score is slightly inflated which offsets the SMF score and correctly downscales what otherwise might have been an equal score to a benzene detection. Unfortunately, this metric was still

confused by increasing contributions from ozone with changing atmosphere and could still not distinguish between contributions from benzene -- yet another metric was needed.

Spectral Depth

The implementation of spectral depth came about for two reasons. First, when running detection scenarios under multiple spectral resolution settings, it was found that the aforementioned detectors, particularly SMF, would favor the low resolution versions of the spectral features as opposed to the higher resolution, where discrimination between species was easier. Further, they were unable to accurately distinguish between confuser gases and benzene, and they erroneously indicated an increasing detection score when only ozone was increasing from atmospheric effects. A final metric was required to ameliorate the overall performance of the gas detection.

Spectral depth was employed in this work similarly to how it is used in the Tetracorder model (Clark, Swayze, & Livo, 2003). A baseline is established at the region containing the spectral feature and the distance to the peak of that feature is calculated. As implemented in this work, an additional point was used in tying the baseline to the contribution of ozone in the presence of benzene. In this manner, as the overall spectral depth increases from the increased presence of ozone, the baseline increases by the same amount, effectively nulling the increase. This allowed for the easy distinguishing of benzene from confuser gases as well as ozone and correctly identified more favorable detection conditions at higher spectral resolution. Figure 3.37 shows an illustration of spectral depth with the offsetting baseline concept in the presence of ozone.

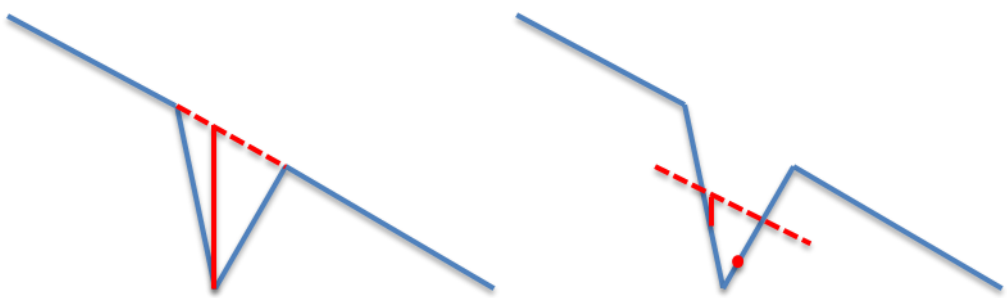

FIGURE 3.37 SPECTRAL DEPTH ILLUSTRATION FOR BENZENE ONLY (LEFT), IN THE PRESENCE OF OZONE WITH ADDITIONAL POINT (RIGHT)

Figure 3.37 at left shows an example of the predicted baseline in the local spectral feature region with a vertical line extending from the mean baseline to the peak of the spectral feature, measuring the spectral depth. The peak at right is shifted slightly along the spectrum, now centered over the ozone spectral feature. The peak shown is largely due to ozone with a small, masked presence of benzene. The additional point that is associated with the ozone spectral feature is shown to influence the mean baseline prediction, shifting it to a lower mean radiance. This reduction in spectral depth helps to reduce FA associated with the increased presence of the ozone spectral feature.

Spectral Energy Normalization

It was found that while the spectral ratio was generally acceptable under most conditions, it could be improved by normalizing the detection score by the spectral depth result. Thus, if a comparatively high score resulted from a low resolution detection of the gas plume as compared to the high resolution detection, the spectral energy contained in the spectral feature – essentially its spectral depth, could be used to normalize the spectral ratio results and correctly identify the more easily distinguishable gas. This step is implemented in the work by multiplying the spectral ratio by the spectral depth. This metric is referred to as the WSSEN metric in all future detection plane images.

Figure 3.38 shows an excellent example of how the detection algorithms can be fooled and how the combination of various detection methods increasingly improves detection performance. The following detection plane was run under a 5 km MLS atmosphere, looking for benzene with a target vector produced at 4800 ppm-m with a 32K thermal contrast over a blackbody background. The scene consists of the gas present in every pixel at 270K over a 300K asphalt background. The larger strength of the target vector makes it appear more concentrated in the detection plane than it really is. The vertical axis represents a trade in maximum OPD which corresponds to instrument resolutions between 1 and 5 cm^{-1}. The horizontal axis represents a trade in concentration of the target gas from 0-20,000 ppm-m. This type of system trade would indicate what resolution is most effective for detecting a given concentration of gas. Said another way, it shows which conditions produce a spectral feature that produces the largest metric score.

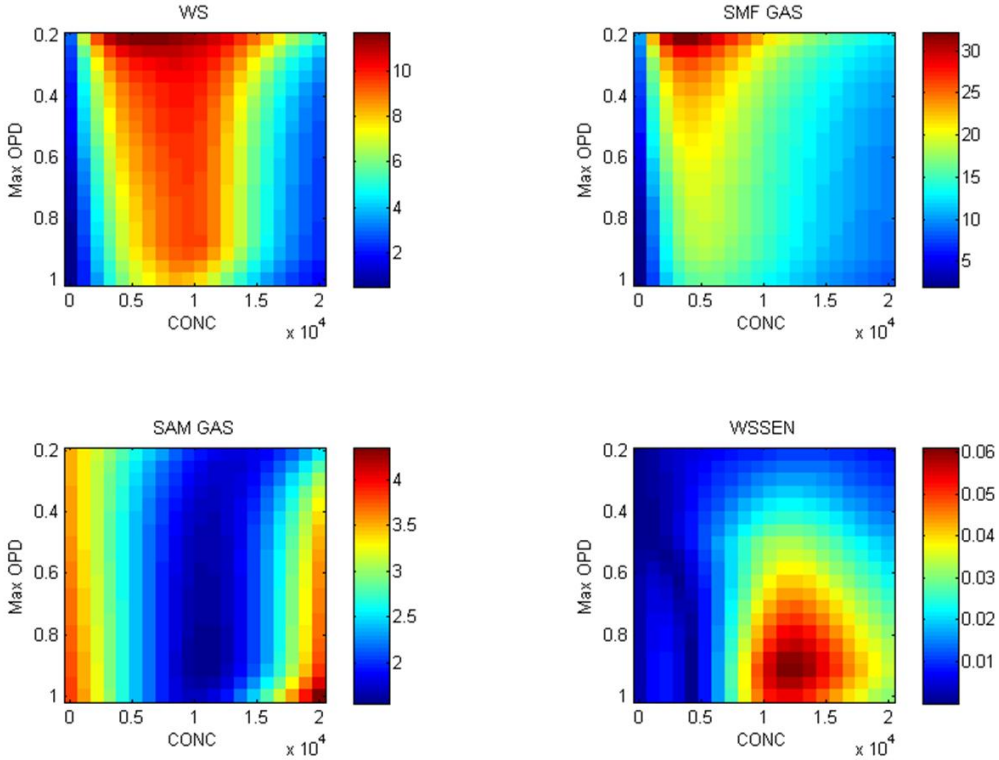

FIGURE 3.38 DETECTION PLANES ILLUSTRATING DETECTION METRIC CONCERNS

The top right panel shows the SMF Gas score map and as mentioned, the SMF Gas score was at a maximum at the lowest resolution due to the combined energy of the significant ozone band at 5 km and the nearby benzene, which have combined at this resolution. Remaining scores in the SMF plane indicate the correct identification of the target vector across the remaining resolutions. The lower left panel shows the SAM Gas metric result from the target vector band in the spectrum. The SAM Gas metric was only moderate at low resolution and correctly improved toward higher resolution, but appeared to require a larger than expected concentration of the gas to yield the closest spectral match.

The Spectral Ratio in the upper left panel shows the ratio of SMF Gas to SAM Gas and while it shows an improved detection result, it is still centered on the low resolution version of the spectrum. Finally, the plane in the lower right panel is the Spectral Energy Normalized - Spectral Ratio. This metric correctly identifies and distinguishes benzene from the presence of ozone, and does not yield a high score until the resolution coupled with increased concentration allow for benzene to emerge past ozone. An analysis of the spectra involved shows that indeed, the presence of benzene can be clearly seen as distinct from ozone where indicated by the metric as compared to the earlier metrics. This result shows that even though one is searching for 4,800 ppm-m benzene, due to the significant masking presence of the neighboring ozone spectral feature, a significant increase in the cpl of benzene is required to be confident in the presence of the target. Analysis of the impact of various performance and environmental trades, including confuser gas impact on target gas detection is shown in Chapter 4.

Detection Metric Response

Figure 3.38 shows that the detection metrics do not necessarily indicate the maximum score response at the target gas CPL that corresponds to the given target vector. This phenomenon occurs for a number of reasons. First, the target vector was produced with no knowledge of the atmosphere or background materials. This was done to reflect an accurate measure of the impact of the environmental change made to a ground element or scene as scene by the FTS model. Thus, changing atmospheres or backgrounds would each produce a unique score result as opposed to a favored atmosphere and background that reflects the given target vector. In this manner all target vectors for this work were constructed from the target gas absorptivity which modulated a blackbody at a given temperature coupled with a blackbody at another temperature representing the background and providing the thermal contrast.

Figure 3.39 shows a series of detection metric planes that give the detection score results by CPL. The CPL ranges from 0-30,000 ppm-m and the target was 290K benzene gas at 20,000 ppm-m over a 300K background. This series shows the changes in detection response as a progression by first matching the conditions of the target vector, then adding an atmosphere with the same blackbody background, then changing the background to asphalt with the same atmosphere.

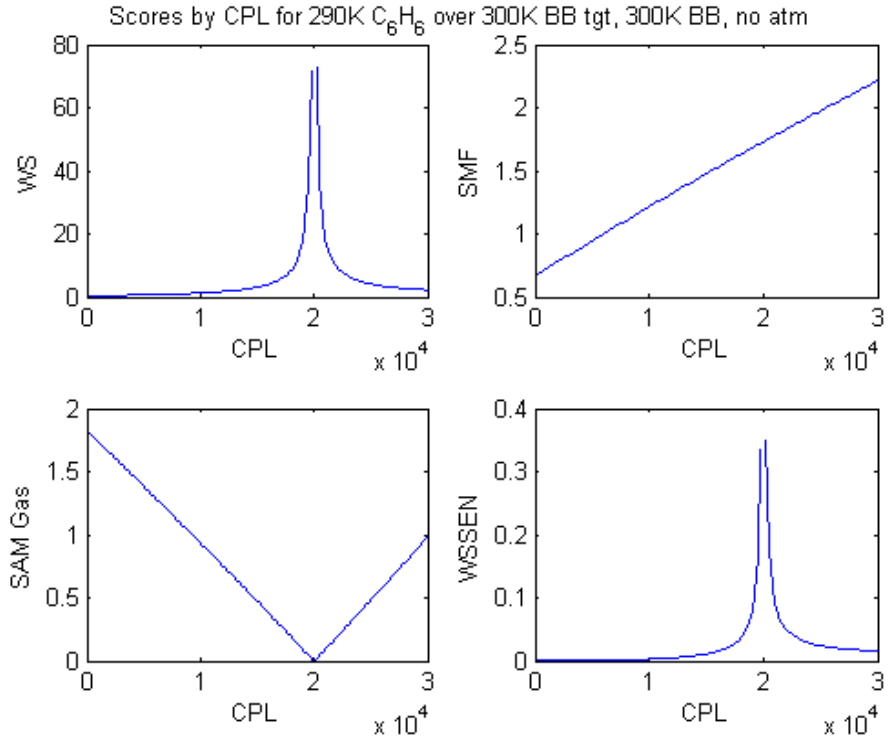

a) Detection planes for 290K target gas over 300K blackbody background, no atmosphere

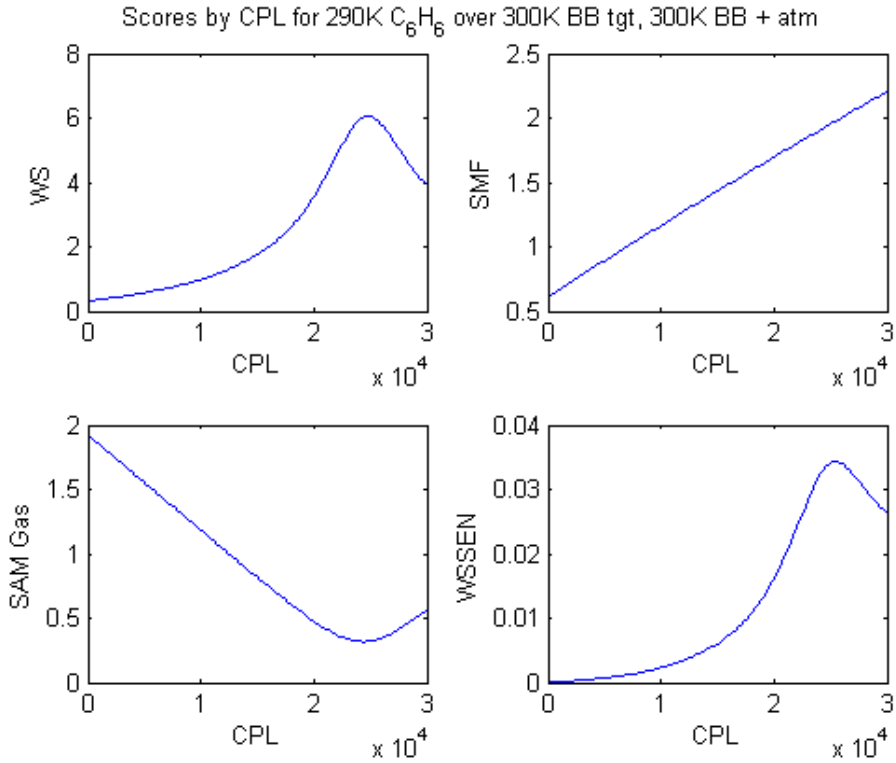

b) Detection planes for 290K target gas over 300K blackbody background, 1km MLS atmosphere

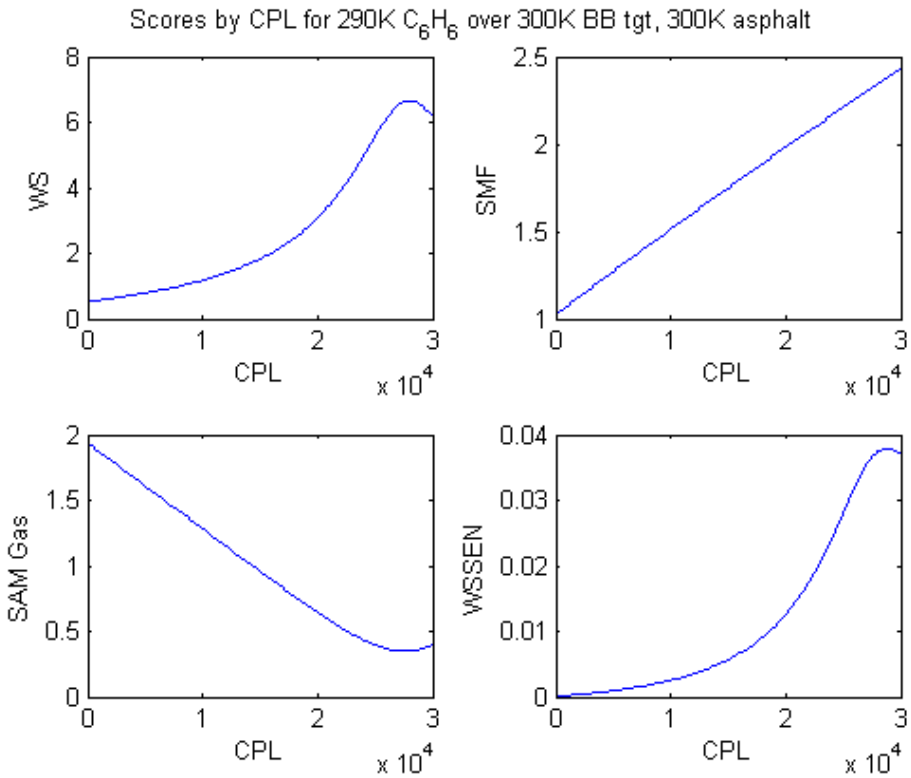

c) Detection planes for 290K target gas over 300K asphalt background, 1km MLS atmosphere

FIGURE 3.39 DETECTION PLANE RESULTS FOR 290K, 20,000 PPM-M TARGET VECTOR BY BACKGROUND AND ATMOSPHERE

Figure 3.39 in panel *a* shows the substantial response of the detection metrics to a perfect match to the target vector. Of course, the nature of the SMF is such that it is largely magnitude dependent, whereas the other metrics are more selective. In panel *b*, it is clear that there is no direct match to the target vector as seen in the previous panel and the maximum score has shifted to a higher CPL than the actual target. This is due to the atmosphere both attenuating and affecting the shape of the spectral feature where a stronger signal is required of the spectral feature before it becomes more "target-like." The much larger shift in preference of CPL for a maximum score result is due to not only the change in background spectral variation, but more importantly, the drop in emissivity of the background material which makes the background appear slightly cooler than it really is. This has an attenuating effect on the strength of the absorption feature and requires a larger CPL for the gas to once again appear more "target-like." This trend can be seen by background material in nearly all of the environmental impact performance trade studies shown in the results section in chapter 4.

The spectral resolution and the temperature of the gas in the collected spectrum also affect the detection metric preference for a given gas CPL. The poor resolution preference has been demonstrated in Figure 3.38 and brought about the need for the additional detection metrics and combination metrics beyond using SAM and SMF alone.

3.1.5 Model Validation and Comparison

This section describes the methods used to validate the model performance and accuracy and compare the output to real sensor systems.

The radiometric accuracy of the model was validated against GATS Atmospheric Science spectral calculators and found to be in exact agreement (GATS, Inc.). The GATS Atmospheric Science company supports NASA, Science Applications International Corp. (SAIC), Ball Aerospace, National Center for Atmospheric Research (NCAR) and more, in satellite-based atmospheric remote sensing. They provide free high resolution spectral modeling software for use by researchers, teachers, and students.

Next, the model was used to try to emulate results produced by a trio of high resolution ground-based FTS systems. The results were from a study by Gruber in comparing the efficacy of using passive sensors in quantifying concentrations in gas plumes, where active sensors are far more reliable (Gruber, 1998). The FTS systems consisted of a Bomem, MIDAC, and ILSCAD units and the MIDAC was used for comparison due to having the highest resolution -- it produced 3 cm^{-1} resolution scans at 1 Hz. The setup of the experiment placed the FTS systems approximately 240 feet from a wall while an SF_6 gas plume passed perpendicular to the sensor LOS, just beyond the midpoint of the range. While the atmosphere was reported to be at 28.5 C, no information was given about the temperature of the gas, the background, the material of the wall, nor the gas flow release rate. The active FTS systems estimated the CPL to be 30 ppm-m and derived a measured atmosphere temperature of 31 C.

In building a scene to match this setup, several assumptions had to be made. First, the material of the wall that produced the best restrahlen feature match was actually sand. The temperature of the wall that gave the proper thermal contrast was 306K (33 C) and the temperature of the gas was set to 302K. This temperature was approximately the same temperature as the measured atmosphere from the MIDAC FTS, thus it was assumed the gas had well thermalized by this point. While the active sensors measured the gas CPL to be approximately 30 ppm-m, a setting of 35 ppm-m gave the best match to the provided spectra. The resolution was set to match the MIDAC at 0.33 maximum OPD. Finally the atmosphere was scaled to approximate a 240 ft path.

The results of the model output as compared to the published spectra are shown below. Figure 3.40 shows the first series comparing transformed SRR from 800-1250 cm^{-1} – the vertical lines on the collected spectrum indicate the bounds that were used in the modeled spectrum. Note the radiometric accuracy of the model when compared to the collected spectrum.

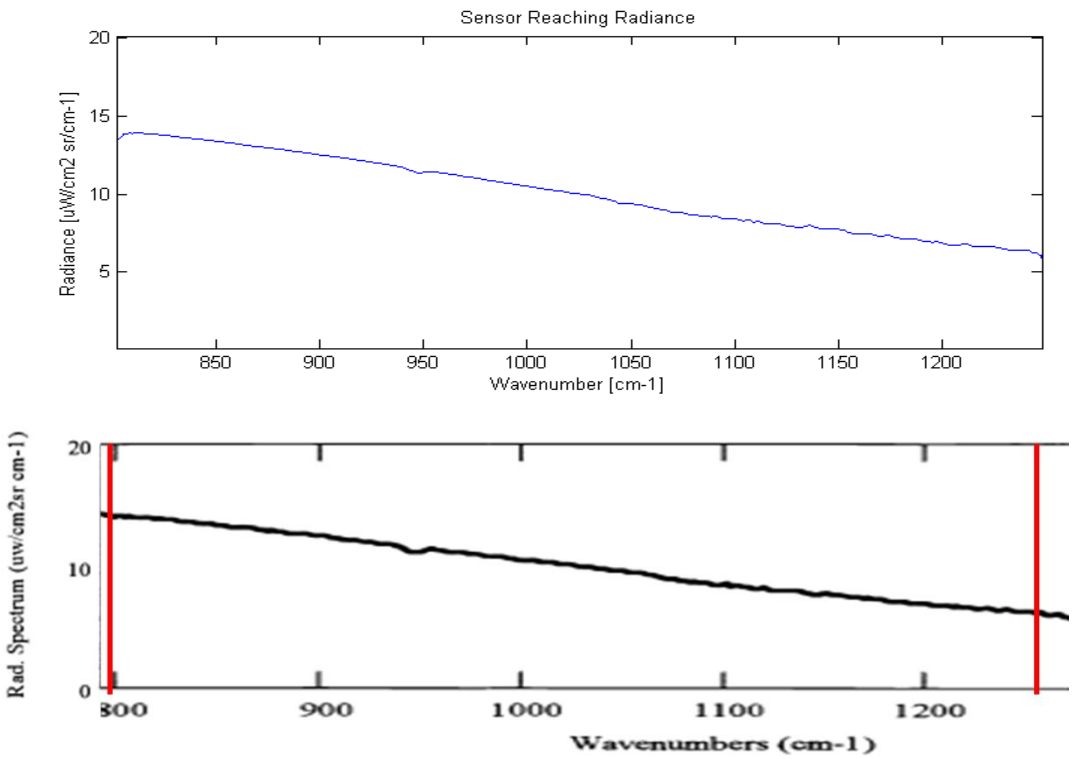

FIGURE 3.40 COMPARISON OF TRANSFORMED SRR BETWEEN MODEL (TOP) AND GROUND-BASED FTS (BOTTOM)

The next series of spectra show a comparison of the brightness temperature of the scene as compared to the collected data. Figure 3.41 shows that the overall background and atmosphere are well modeled within tolerance for estimations made for the various scene temperatures.

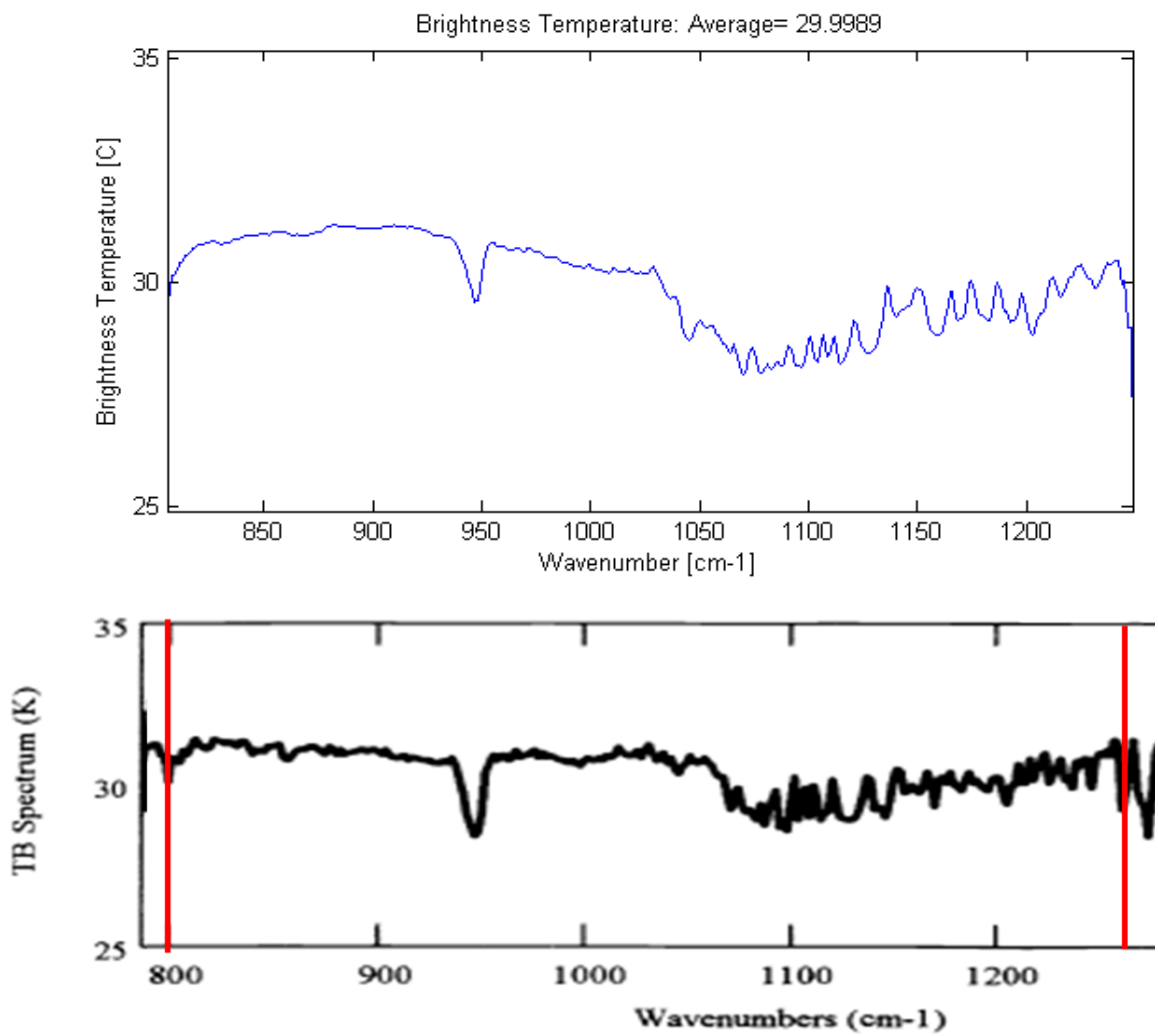

FIGURE 3.41 BRIGHTNESS TEMPERATURE COMPARISON FOR MODELED (TOP) AND GROUND-BASED FTS (BOTTOM)

The results of this comparison gave good indications that the end-to-end version of the FTS model was working correctly and could easily model customized scenes with minimal effort.

The next comparison was to the Airborne Hyperspectral Imager (AHI) sensor collection of SO_2 during an over flight of a volcano in Hawaii (Mares, 2004). AHI is a DES system and not an FTS, but as stated earlier, no comparatively high resolution data are available from an airborne FTS. As mentioned in the Spectral Output section, the comparison between not only FTS to DES derived data, but also when converted to the wavelength domain, causes minor disparities between the spectra.

In modeling the collection parameters, it was reported in the paper that the aircraft was at 4,000-4,500 ft MSL. Of course, this is an over flight of a volcano, so the actual AGL distances are likely to be much shorter as ground altitudes over the volcano ranged between 3400-3700 ft MSL according to Google Earth. Therefore, it was estimated for the model that the AGL altitude was 250 m (820 ft). The spectrum was reported to be collected in 32 bins between 7.5 and 11.5 µm, yielding 0.125 µm resolution. This corresponds to a resolution of roughly 15 cm^{-1} at the center of the band and a

MOPD setting of 0.08 cm. The paper reported an estimate of the background to be a 335-363K 0.96 emissivity gray body, with the vapor plume likely to have a 150 m pathlength and a range of 325-1400 ppm concentration.

The best fit to the data used sand as the background with a temperature range of 346-360K, SO_2 gas with a range of 335-350K at 700-725 ppm with a 150 m path. The gas concentration was chosen as a middle range of the reported values. The atmosphere was generated as a 250 m Tropical model. All values used in the model were within the ranges specified in the paper and considering the estimates that had to be made for altitude, atmospheric temperature, background material, gas temperature and concentration, the fit of the data is quite good.

Figure 3.42 shows multiple spectra from various flight lines that represent different portions of the SO_2 plume. Several spectra from this plot were modeled and found to be in good agreement. Note that the long wavelength side of the modeled spectrum appears under resolved while the short wavelength side is slightly over resolved. This is a consequence of the transformation to the wavelength domain.

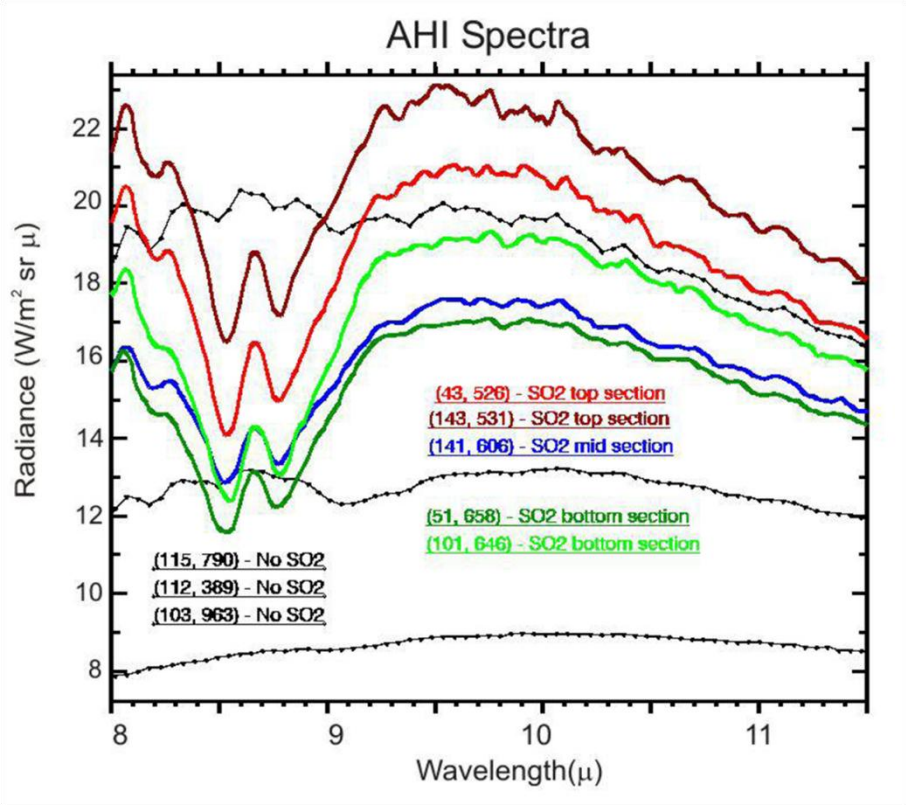

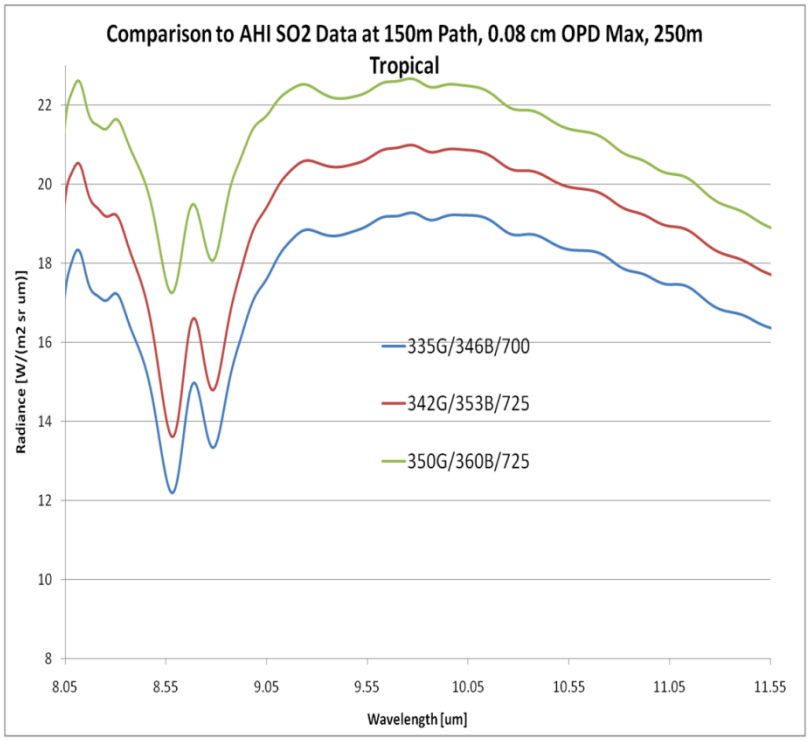

FIGURE 3.42 COMPARISON OF AHI COLLECTED (TOP) AND MODELED (BOTTOM) SPECTRA OF SO_2

The legend in the modeled spectra gives the temperature used for the gas, the temperature of the background, followed by the concentration of the gas used. Note the radiometric accuracy of the model in comparison to the collected spectra.

3.1.6 Scene Building

In order to test the airborne scanning module over a scene more complex than a gas plume over a single or blended background, a material class map was built and populated manually. The goal of the scene was to provide a significant amount of spectral, thermal, and spatial emissivity variability while keeping the scene components realistic in nature. The scene constructed for this work is shown in Figure 3.43 as a temperature scaled image. Note that there is good thermal variability not only within a material class, but between classes as well. The scene consists of water as modeled by a blackbody, sand, concrete, steel pipe, asphalt, and aluminum. The scene simulates a gas pipeline distribution facility and is largely comprised of sand. There is a small pond in the upper left quadrant of the scene and continuous concrete pads supporting various steel pipe structures. There is an asphalt drive at the bottom of the scene and a concrete walkway to a shed with an aluminum roof. The assignment of temperature values to the various material classes were based on a time of day just after the sun had set and loosely took into account the thermal inertia of each material class.

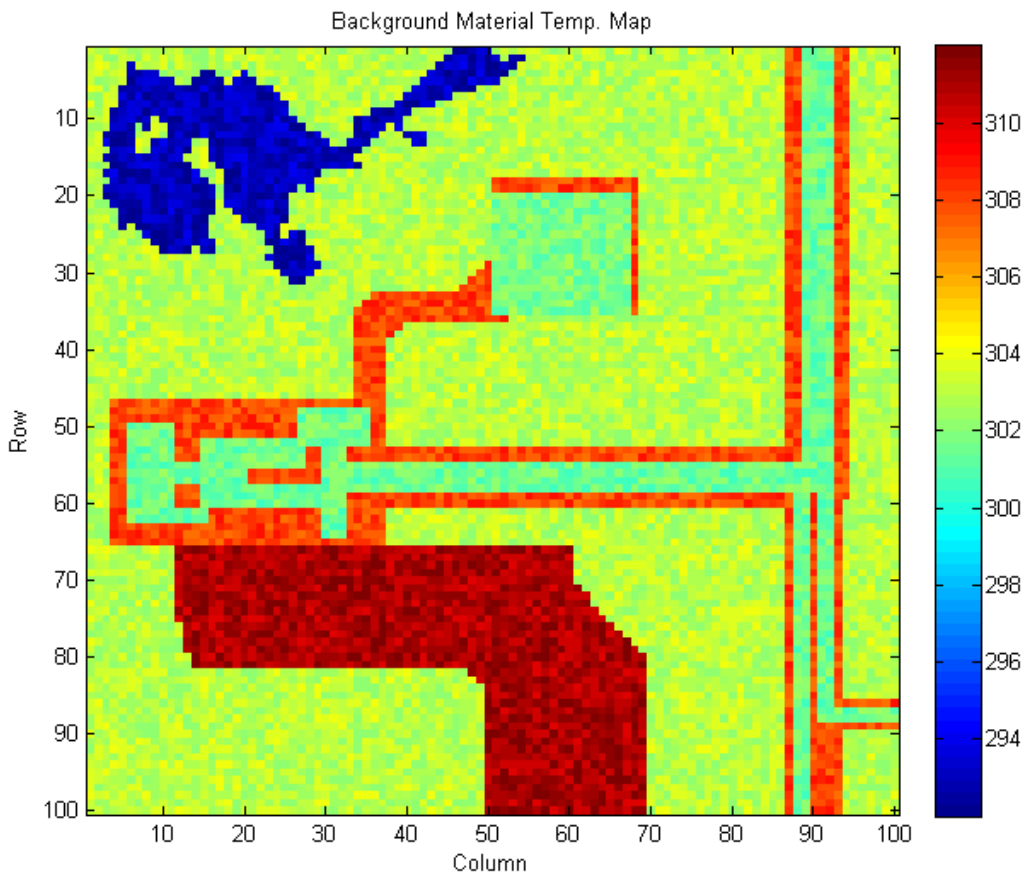

FIGURE 3.43 THERMAL VARIABILITY IN A SIMULATED SCENE

The thermal variability in the scene was accomplished by assigning a target temperature to each material class. Then each pixel within a material class was assigned a value from a uniform random distribution between 0 and 2, giving the variability among and between material classes.

Gas plumes were placed in the scene in multiple locations to provide for false alarm (FA) detection characterization as well as side-by-side gas plumes to determine the resolution requirement for successful confuser gas discrimination.

The scene actually consists of six to nine truth layers that when read into the model, dictate the specifications for each material, gas, and atmosphere at each ground element coordinate. The layers are shown in Table 3.2. It is important to note that this scene consists of mapped class information only and has no real three dimensional structure that would affect radiometric properties in determining the SRR.

Layer	Variable Defined
1	Gas CPL
2	Gas Temperature Map
3	Gas Type Map
4	Background Emissivity Map
5	Atmospheric Profile
6	Background Temperature Map
Optional Layers for Confuser Gas	
7	Confuser Gas CPL
8	Confuser Gas Temperature Map
9	Confuser Gas Type Map

TABLE 3.2 SCENE LAYER IDENTIFICATION TABLE

In this way, as the SRR cube is assembled from the given scene specifications, each parameter needed is specified on a per ground element basis, allowing substantial customization of the scene for maximum variability. Figure 3.44 shows the SRR for a narrow band around benzene from 1032-1042 cm^{-1}. While three plumes have been embedded in this scene, only two are visible in this band at this thermal contrast and cpl. The plume in the upper left corner is methyl chloride, which has significant spectral contributions in the same band as benzene. The plume source simulates a leak at the manifold in the center left of the scene. There is a plume of ammonia imbedded over the aluminum roof on the shed, but due to the minimal spectral features in this band, the plume is barely visible at the center, upper portion of the image. Finally, the benzene plume is visible to the right of the scene, emanating from the junction of pipes in the middle right. Note the changes in SRR as a function of the material and temperature of the surface beneath the gas.

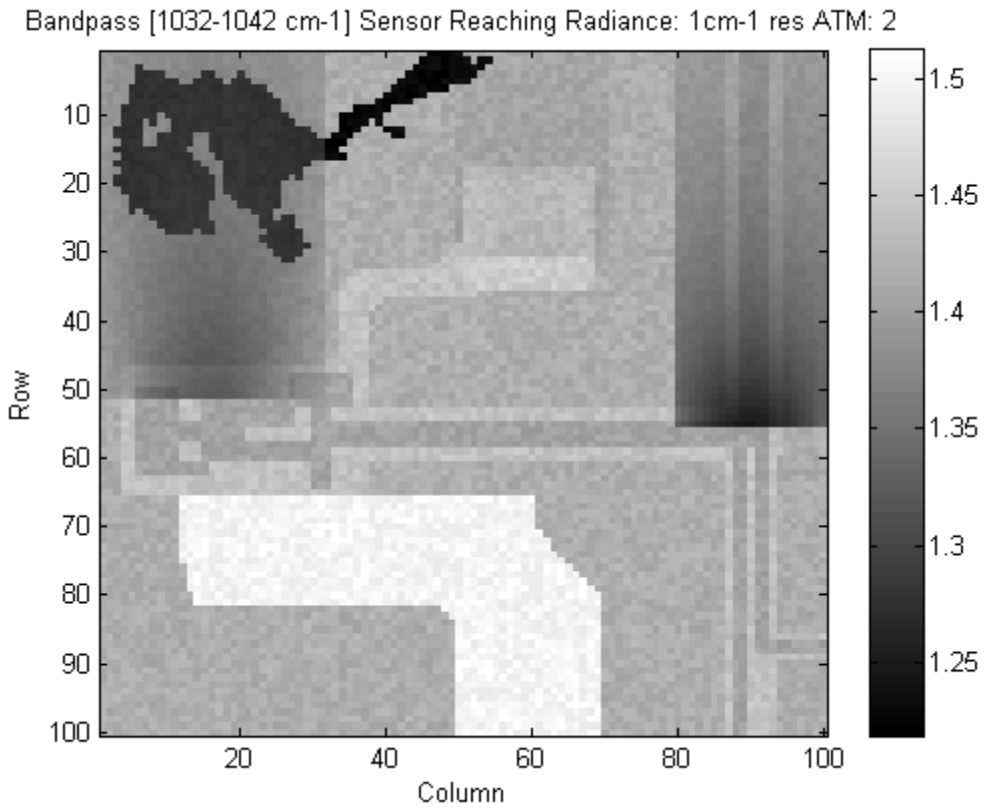

FIGURE 3.44 EMBEDDED GAS PLUMES IN MODELED SCENE

3.1.7 DIRSIG Modeling

DIRSIG was implemented in this work as a capstone test of sorts for the flight and detector model. DIRSIG produces extremely complex, radiometrically accurate representations of a scene for use in instrument evaluation. Scenes produced by DIRSIG contain far more spatial and spectral variability than scenes built by hand, contain three dimensional structure, and use ray tracing methods to develop radiometrically accurate SRR image cubes for input to the model.

The following excerpt, taken from the 2006 version of the DIRSIG User's Manual, provides a concise overview and introduction to the synthetic scene generation and system modeling capabilities accomplished by DIRSIG.

> The DIRSIG model is a complex synthetic image generation application which produces simulated imagery in the visible through thermal infrared regions. The model is designed to produce broad-band, multi-spectral and hyper-spectral imagery through the integration of a suite of first principles based radiation propagation sub models. These sub models are responsible for tasks ranging from the bi-directional reflectance distribution function (BRDF) predictions of a surface to the dynamic scanning geometry of a line scanning imaging instrument. In addition to sub models that have been specifically created for the DIRSIG model, several of these components (MODTRAN and FASCODE) are the modeling workhorses for the multi- and hyper-spectral community. All modeled components are combined using a spectral representation and integrated radiance images can be simultaneously

produced for an arbitrary number of user defined bandpasses. (Digital Imaging and Remote Sensing Laboratory, 2006)

The DIRSIG model has been well validated and is used widely throughout the remote sensing community as a cost efficient alternative to live field collects or in evaluating the efficacy of new detection algorithms or sensor system designs.

Figure 3.45 shows a panchromatic rendered image in the visible spectrum of a model of a synthetic residential area used as the background material source for the DIRSIG studies in this work, courtesy of another graduate student in the DIRS group, Michael Presnar. This region of the model is known as Tile 1 of Megascene 1 and is derived from actual field measurements taken in Rochester, NY. The image shows distance scales on each axis for perspective.

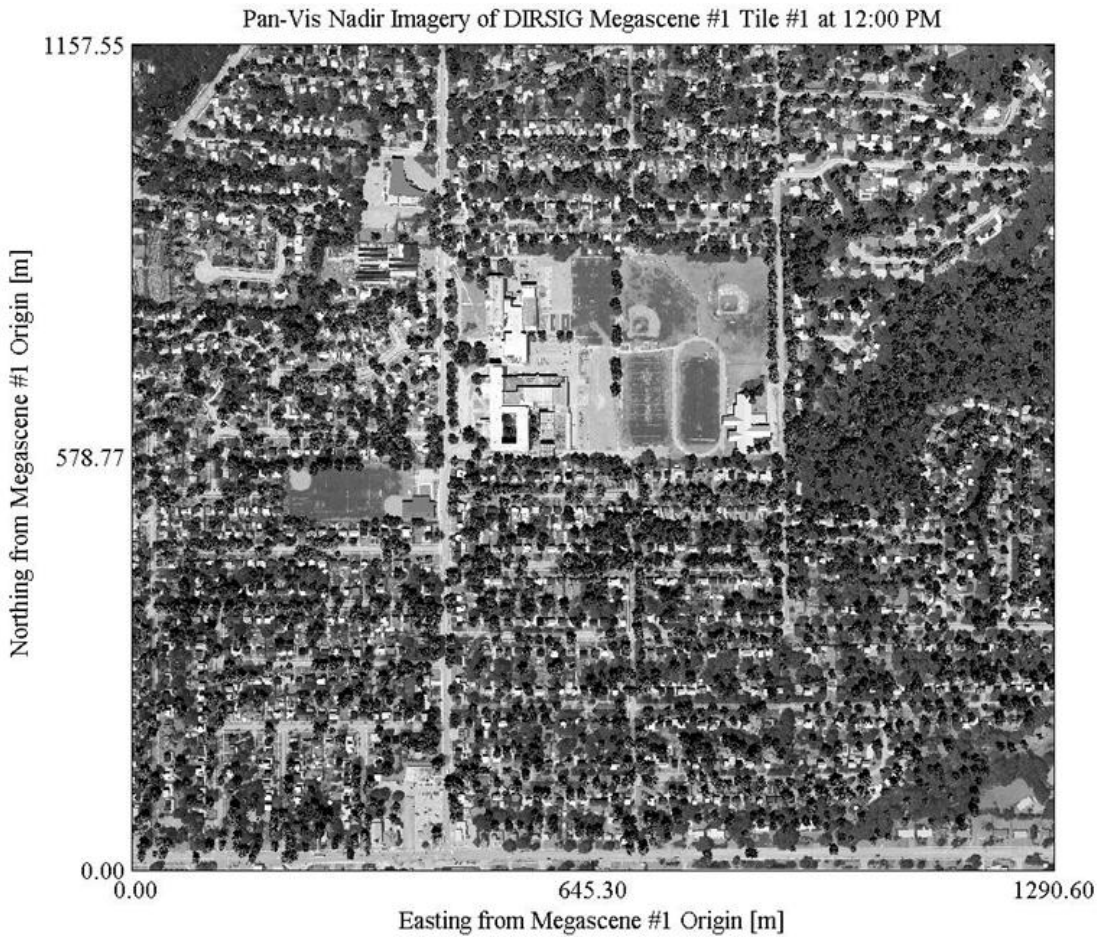

FIGURE 3.45 VISIBLE PANCHROMATIC IMAGE OF MEGASCENE 1, TILE 1 (COURTESY OF MICHAEL PRESNAR)

Figure 3.46 is another rendered image of the same scene in color with annotations marking the two locations gas plumes were released in various studies. One location is in a largely spectrally uniform field with little 3D structure, while the other is embedded in trees and shrubs and presents much more spectral and structural variability.

FIGURE 3.46 RGB COLOR IMAGE OF MEGASCENE 1, TILE 1 WITH ANNOTATED GAS RELEASE LOCATIONS

Figure 3.47 shows the thermal variability in a scene from Megascene Tile 1. The image is rendered from approximately 1km AGL with 0.2 m ground sample distance (GSD).

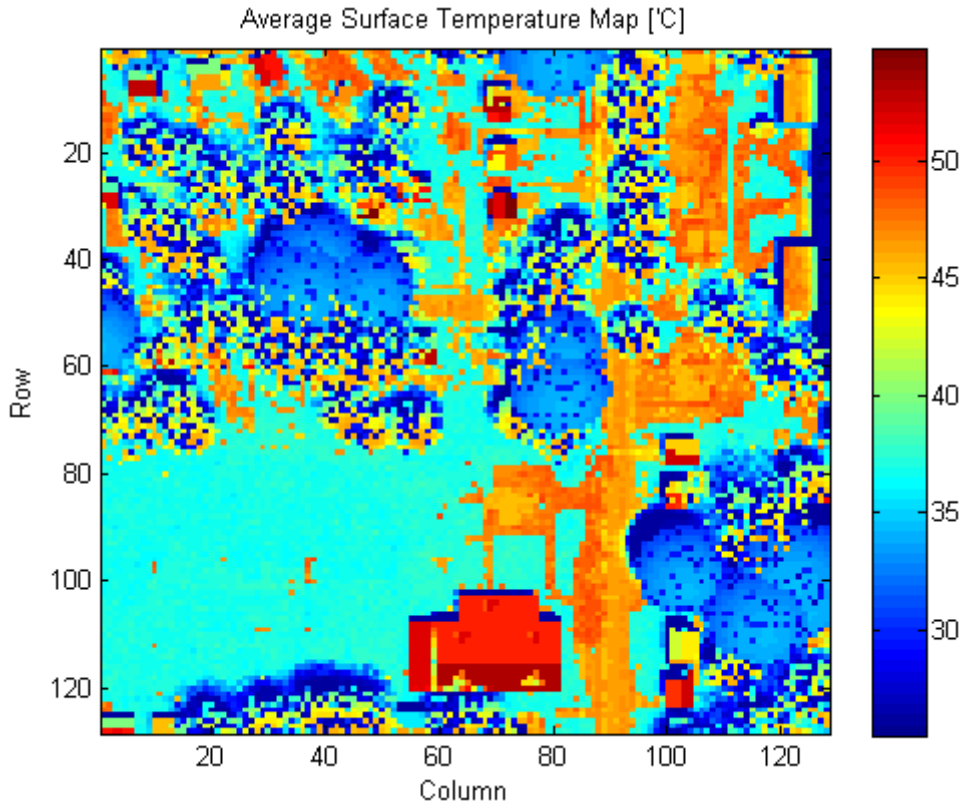

FIGURE 3.47 TEMPERATURE SURFACE MAP OF DIRSIG SYNTHETIC RESIDENTIAL AREA FROM 1KM WITH 0.2M GSD

The model shows a dramatic temperature range from approximately 25 C to 55 C which serves as an excellent testbed for both positive and negative thermal contrast in the plume. The image shows the thermal differences between multiple trees, bushes and shrubs, roads, sidewalks, and several buildings.

DIRSIG Blackadar Plume Model

The DIRSIG plume model that is currently implemented is not covered in the DIRSIG Users manual (Digital Imaging and Remote Sensing Laboratory, 2006). In general, a plume model seeks to describe the disbursement of a gaseous plume from a point source generation site in terms of spatial, temporal, and volumetric properties. Often, computational fluid dynamics are used to describe the motion of the particles in the plume based on the spatial variations of both pressure and temperature throughout a given volume. The ultimate goal would be to provide a user with knowledge of the spatial distribution of the plume such that remote sensing measurements would reveal accurate concentration path length products.

An example of an early plume simulation from DIRSIG showing a gas release overlaid on a synthetic residential area from Megascene1 and is shown in the figure below. The release, atmospheric, and instrument collection parameters are not known.

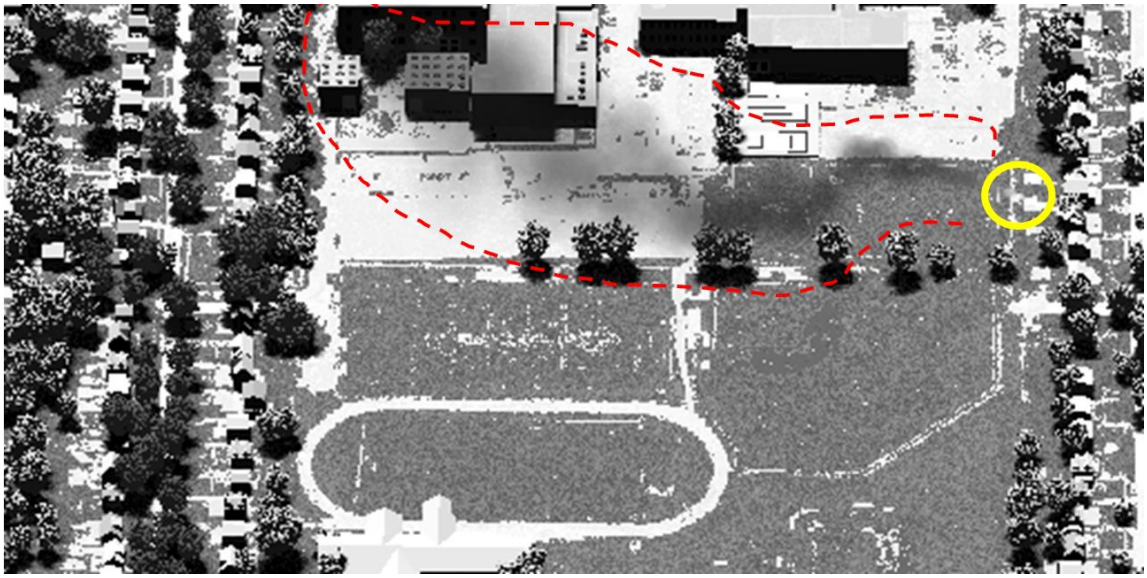

FIGURE 3.48 DIRSIG EXAMPLE GAS PLUME MODEL, SOURCE MARKED IN YELLOW CIRCLE, PLUME OUTLINED IN DASHED RED

The most current plume model used in DIRSIG is the Blackadar model and is used for modeling plume evolution from an exhaust stack. The Blackadar model uses computational methods to model the evolution of a plume over time and attempts to assign temperatures, CPLs, and spatial extent of plume constituents (Blackadar, 1997). The user sets a gas flow release rate in ppm/s and then sets how many 'puffs' per time increment the exhaust stack emits. The physical location, diameter and elevation of the stack are specified as well as wind speed and direction. Both the ambient and gas temperature at release are used in calculating the temperature values across the plume. Finally, the user specifies how long they want to allow for the plume to develop before the model is frozen and implanted in a scene. To this point, the plume model is static in the scene as the over flight is accomplished, which is unphysical, but the best that can be accomplished to date. The plume currently can only consist of one gas type, but in the near future, the ability to inject or overlay multiple plumes into a scene will be incorporated into the model. The plumes evolve as if they are in a flat open space, so there is no interaction with surrounding structures, nor consideration for disturbances to wind patterns as a result of those structures. Finally, the model used to determine the physical expansion in the plume is the same in both the vertical and horizontal dimensions, where in reality the plume is not spatially symmetric.

While there are some drawbacks to the model as mentioned above, it is still the best source of accurate plume modeling for DIRSIG users. Examples of the modeled plume follow.

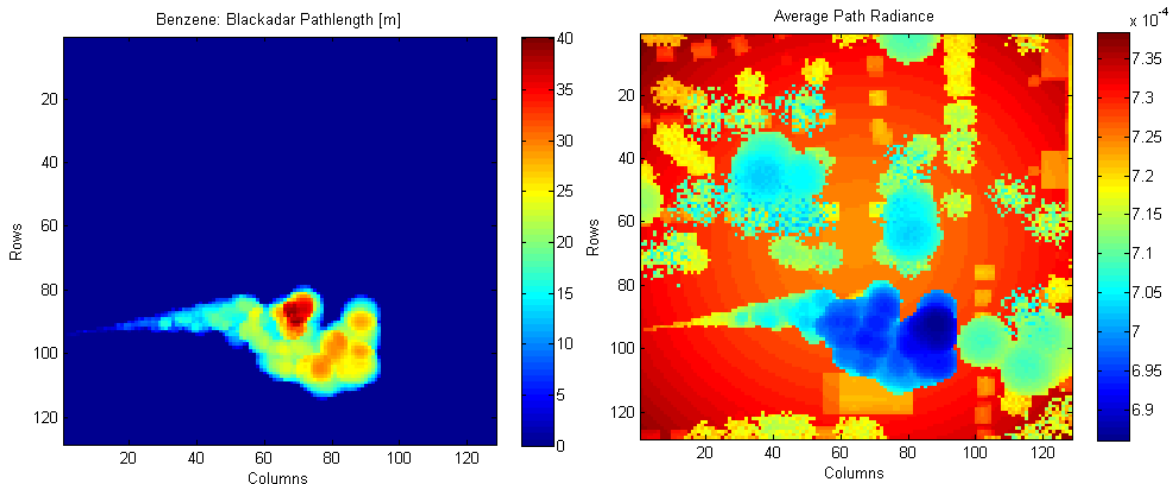

FIGURE 3.49 PATHLENGTH (LEFT) AND PATH RADIANCE (RIGHT) FOR SAMPLE BLACKADAR MODELED GAS PLUME

Figure 3.49 shows the path length and average path radiance for the Blackadar modeled plume in DIRSIG. The image to the left is an isolated truth mask that shows the rapid expansion of the plume in the spatial extent, particularly once it begins to eddy downrange. The image to the right is of the average path radiance over a synthetic residential area taken from Megascene Tile 1. It shows that the plume can be absorptive and reduces SRR from emitting materials beneath it. The concentric bands in the image are due to the cosine roll-off of the projected image energy from off-axis locations in the image.

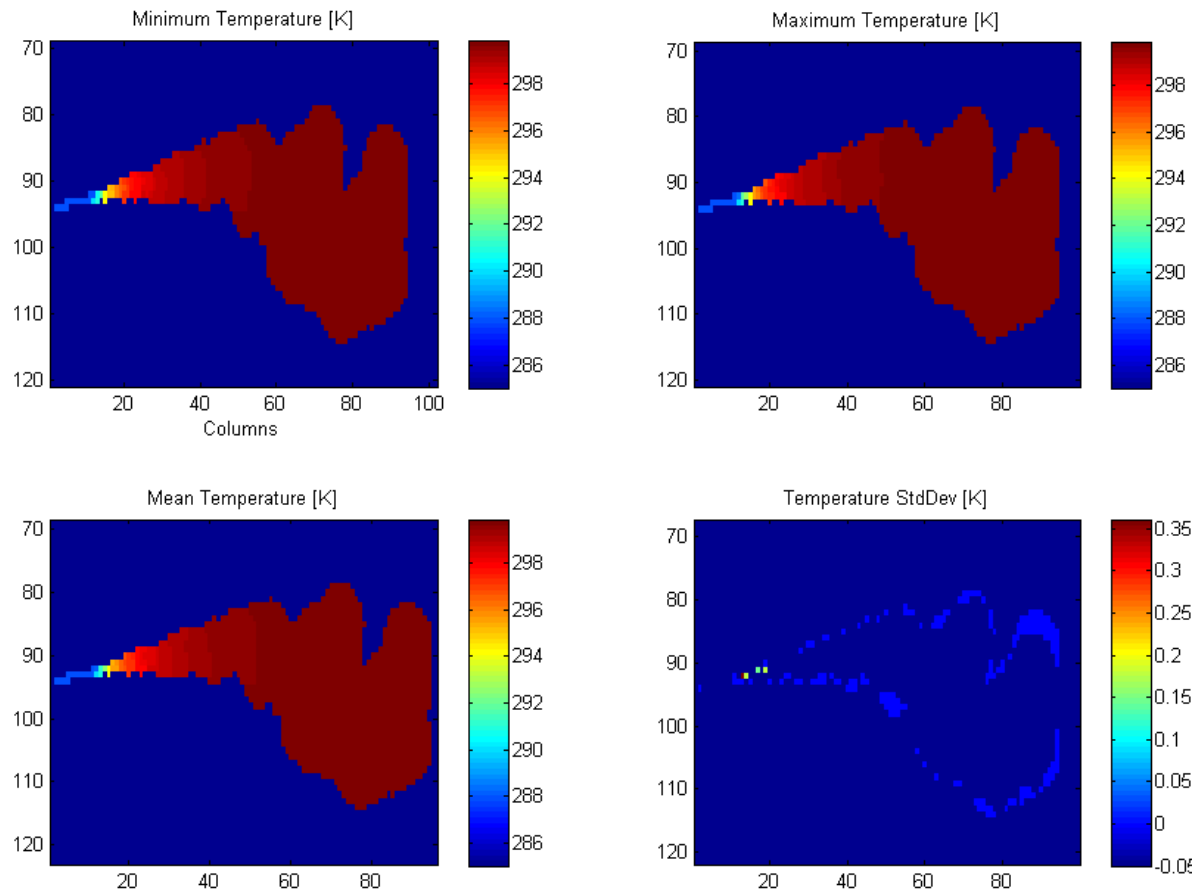

FIGURE 3.50 TEMPERATURE DISTRIBUTION IN BLACKADAR MODEL GENERATED GAS PLUME

Figure 3.50 shows the temperature distribution in the plume. Note that the plume rapidly thermalizes to the atmospheric temperature of 300K (27 C). The major variability in temperature occurs in the first quarter of the plume and the increased thermal contrast, as well as high concentrations, provide the best opportunity for detection. The panel on the lower right of the figure illustrates the variability within the plume and shows that it is at a maximum in the early parts of the plume and to a lesser extent, out at the boundaries where contact/mixing with the atmosphere is maximized.

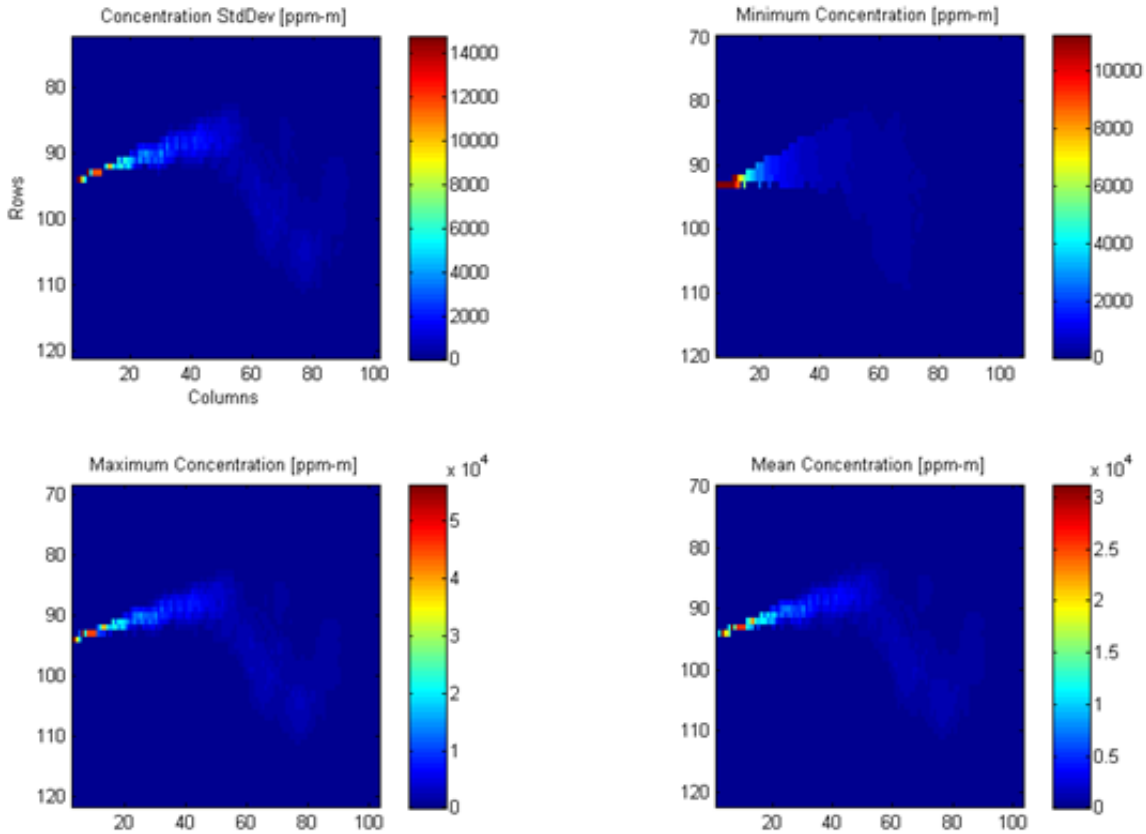

FIGURE 3.51 CONCENTRATION MAP FOR BLACKADAR MODEL GENERATED GAS PLUME

Figure 3.51 shows the distribution in CPL across the plume. As with temperature, the variability in concentration occurs early in the plume within meters of the release point. Interestingly, even though the pathlength approaches over 40 m near the end of the plume, the concentration is so dispersed that the CPL value is near zero. The wide range of CPL ensures a good test of not only the detection metrics for a given target vector, but also in performance trades that affect sensitivity.

The following images in Figure 3.52 are of the detection planes for the gas plume shown above first released in an open field, then the same plume released north of this location in the trees. The data were collected by the airborne model at 50 m/s and a scan rate of 60 Hz at maximum resolution. The gas shown is benzene in absorption with a 1km MLS atmosphere. The initial release concentration was 30,000 ppm-m and quickly decays to below 500 ppm-m over the length of the plume. The release temperature was 286K which quickly thermalizes to the atmospheric temperature of 300K within the first half of the plume extent as shown in Figure 3.50. The images are further analyzed in the Results section, 4.2.3.

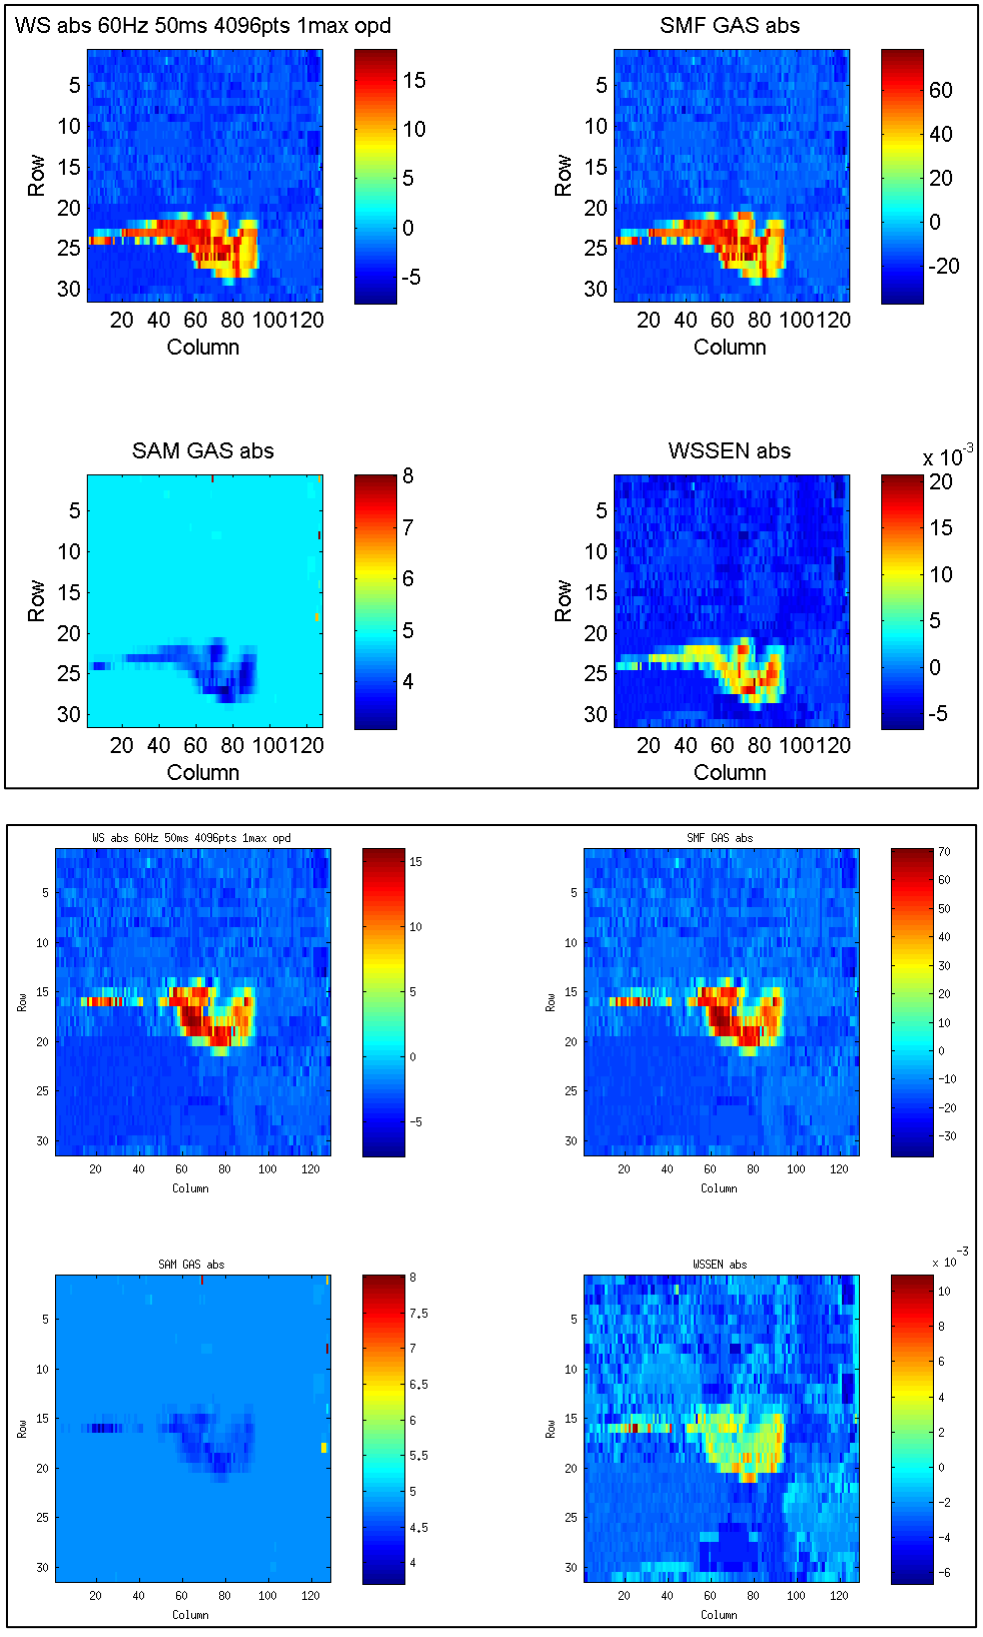

FIGURE 3.52 DETECTION PLANE IMAGES FOR GAS PLUME IN FIELD (TOP) AND IN TREES (BOTTOM)

3.2 Performance Trade Studies

This section details the approach taken to the various studies that were conducted in this work. The major goal of these studies is to explore how the process of interferogram formation, as unique to the FTS, responds to various changes and trades in the environmental, instrumental, and error regimes and how those changes affect the ability of an FTS to successfully characterize the presence of a target gas. In some aspects, one would expect the FTS to respond as any other spectrometer, such as target detection performance under different types of atmosphere. However, it is when errors are introduced that are unique to an FTS, such as moving mirror uncertainty, or in scanning across a scene that new information is gained.

These studies were conducted on a point basis where a 'scene' in this case is composed of generally only one background material, one atmosphere, and fixed gas and background temperatures, while other parameters are varied. In the case of a jitter study, the off-target viewing would be of another 'scene' of a specified composition. Each of these tests will be called an experiment, as distinct from actual scene-based studies. The result of each increment of a study is a single transformed spectrum that was then subjected to the detection metrics. Single values were derived from each metric, recorded, and the next increment was started. Scene-based studies, as distinct from the experiment-based performance trade studies, were conducted later and studied the interplay between highly variable backgrounds, temperatures, and atmospheres. This work is introduced in section 3.3.

The initial studies characterized the impact of the moving mirror and jitter errors in the form of a Monte Carlo simulation where 200, 500, and 1,000 trials of random seed variables were run for each effect and both effects together. The results of this study allowed for a mapping of the range of impact of each error source to determine a suitable envelope for testing the model under scene collection conditions with an appropriate amount of error.

The next major study pitted resolution against concentration under a variety of environmental conditions and confuser gas concentrations.

3.2.1 Monte Carlo Mirror and Jitter Studies

The purpose of this study was to map the range of influence that each error source had on the ability to successfully detect a gas as determined by the detection metrics. Since the application of the error to the interferogram is a random event in itself, it stood to reason that a Monte Carlo study would be a good approach to seeing the overall effect of the errors in terms of the mean and standard deviation of the resulting detection scores.

The main point to consider in these studies is that the accuracy of the resulting transformed spectrum is highly dependent on where in the spectrum the major errors occurred. The large number of random events helps to determine the maximum anticipated effect of the error. As mentioned in the detection section, each detection metric returns a higher score based on slightly different criteria. Thus, while an early interferogram interruption may cause wild bias errors in the

transformed spectrum, the SMF score would come back high due to the erroneously increased energy, but SAM would generally still register a good match spectrally.

Due to the computational burden of accomplishing this volume of runs, the Large Memory Cluster (LMC) at the Research Computing Center at RIT was employed for this work. As an example, accomplishing 1,000 trials for 5 jitter run lengths at each of 5 mirror error uncertainties took 18 hours per gas on this multi-core resource.

The code essentially executed as one large loop, incrementing mirror and/or jitter error magnitudes after the specified number of random events had been run. The scenario was executed with one background material type, one gas type, fixed resolution, fixed atmosphere, fixed concentrations and temperatures. Thus, the only difference from run to run was in the magnitude and location in the interferogram of the mirror errors. The standard deviation of the error level limited the total number of errors in the interferogram to a value that varied less than 2% between runs.

3.2.2 Resolution versus Concentration

The resolution against concentration trade studies served as the foundation for testing the impact of multiple instrument parameters, errors, and environmental effects on detecting a gas both with and without the presence of a confuser gas. The resolution range for these studies spanned 1 to 5 cm^{-1} which corresponded to a MOPD range of 1 to 0.2 cm. It was found that anything below 0.2 cm OPD resulted in substantial FA detects as multiple spectral features from the atmosphere, gas, and confuser gas blend together and become indistinguishable. In addition, concentrating on the higher resolution sets this study apart from prior work.

Referring back to Figure 3.38, it shows an example of a detection metric image plane from a concentration versus resolution experiment. The image plane is populated by conducting a run on the specified resolution and concentration for the gas while all other variables are fixed for the duration. The resulting score metric is recorded and the next increment is run. So when testing the impact of mirror error on the ability to find benzene with substantial ozone contribution from a 5km atmosphere with the presence of methyl chloride at 60,000 ppm-m, all temperatures and materials in the experiment are fixed and only the target gas concentration would be varied at each resolution value. Generally, 20 values for resolution and 20 for gas concentration would be compared, yielding a 400 pixel detection plane. As can be seen from the example, multiple study types can be executed with such a large number of scenarios to run against.

3.2.3 Confuser Gas Evaluation

As mentioned in the target gas section, the confuser gas chosen for masking benzene in this work was methyl chloride. The peaks from the Q branch of benzene and that of the neighboring methyl chloride peaks were separated by only 4 and 3 cm^{-1} on each side. This allowed the peaks to be totally unresolved at the lowest maximum OPD setting in the trade study and for the peaks to gradually separate as the resolution was increased. Three concentrations of confuser gas were injected into the experiment which represented ½, 1, and 1.5x the strength of the benzene target vector. This translated to 30,000, 60,000, and 90,000 ppm-m concentrations of the gas. Both gases

were given the same temperature profile. While one could easily distinguish methyl chloride in the presence of benzene due to the larger bandwidth of the complete set of methyl chloride spectral features, the test is to see if benzene can be distinguished from methyl chloride.

3.2.4 Atmospheric Impact Studies

As ozone was the biggest culprit in masking the presence of target gases in the region, it was of interest to vary atmospheres between 1 and 5 km pathlengths and with variable emissivity backgrounds. The altitude difference significantly increased the total column content of ozone and the variable emissivity served to increase the presence of the down welled radiance. This effort would characterize detection performance both as a function of the error introduced, but also as a function of the presence of ozone in the spectrum.

The difference between the MLS and tropical atmospheric models is essentially isolated to the total column water content, which is nearly double in the tropical atmospheric model. The increased presence of the water rotational continuum provided many similar sized peaks as those being matched by the detection metrics and provided substantial differences in spectral energy in various portions of the SRR. In particular, this effect would cause major problems in the detection of methyl chloride.

All trade studies were begun with no atmosphere to isolate the error effect, then continued through each of the atmospheric models.

3.2.5 FTS Instrument Studies

Early studies were conducted to verify the accuracy of the instrument model and determine limitations. These studies included presenting the model with various combinations of delta functions to determine transform accuracy, sampling studies to determine optimal sampling strategies and associated effects, and apodization and scaling effects on radiometric accuracy and spectral feature recovery, to name a few. Most of these studies were nominal, with the exception of the point sampling study which showed how sampling accuracy impacted the spectral feature shape and therefore affected the resulting detection score in single sided interferograms without phase correction. Some scenes were run in both double and single sided interferograms and the impact to detection performance was reviewed.

3.3 Modeled Scene Detection Studies

In order to study the interplay between increased variability in surface emissivity, thermal contrast, and to test a rudimentary scanning model, scenes were constructed. The complexity, realism, and accuracy of the scenes increased with each evolution of the scene model. The initial scene consisted of various Gaussian-based plume models inserted around a scene comprised of regions of the same surface material and temperatures, with different atmospheres in bands over the scene. This progressed to a materials map resembling a gas distribution center with various pipelines and simple structures that were assigned their own emissivity values and temperatures. This scene provided substantial improvement in thermal and surface variability. Finally, the DIRSIG modeling program was used to produce realistic scenes modeled from an actual residential area in Rochester,

New York with field collect-derived emissivity files from the site assigned to each structure, and forward modeled thermal profiles for each material based on solar loading estimates. The DIRSIG modeled scenes provide unmatched realism as well as thermal and spectral variability and represent the best test cases for the model to date.

3.3.1 Initial Scene

The initial scene built for the scanning model consisted of the Gaussian gas plume over a scene divided in quadrants by background material type, and rows by atmosphere type, to review differences in detection performance by scene composition. Figure 3.53 gives an example of this study as shown in the four panel series below.

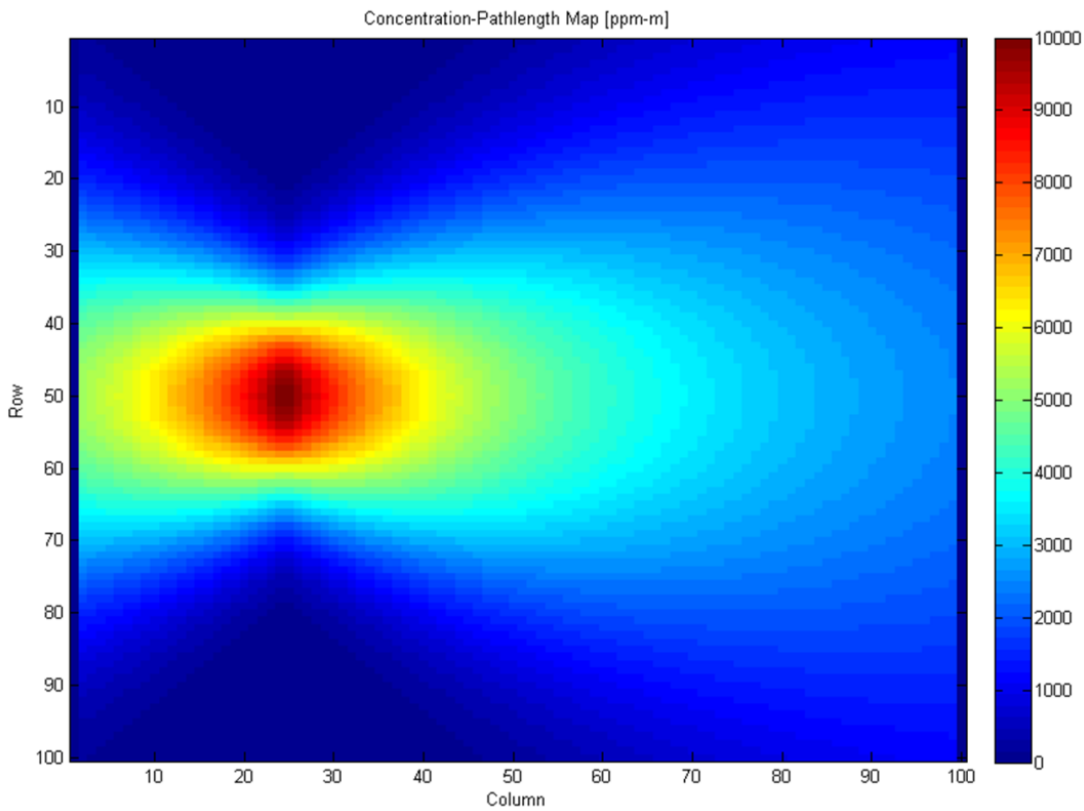

a) Concentration-pathlength map

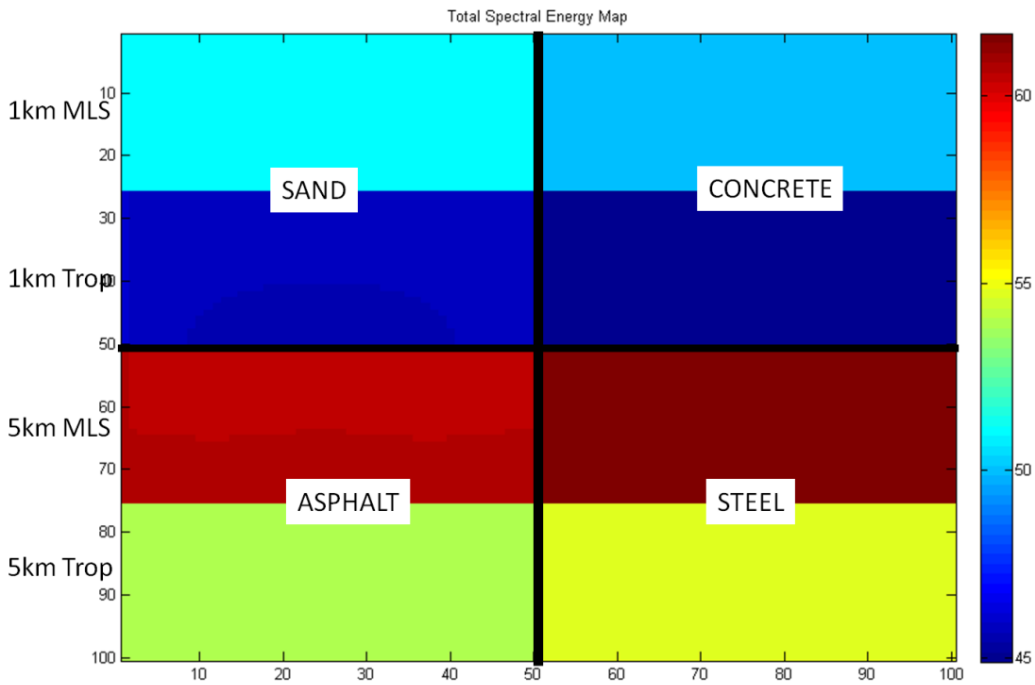

b) Total spectral energy map with material by quadrants and atmosphere by rows

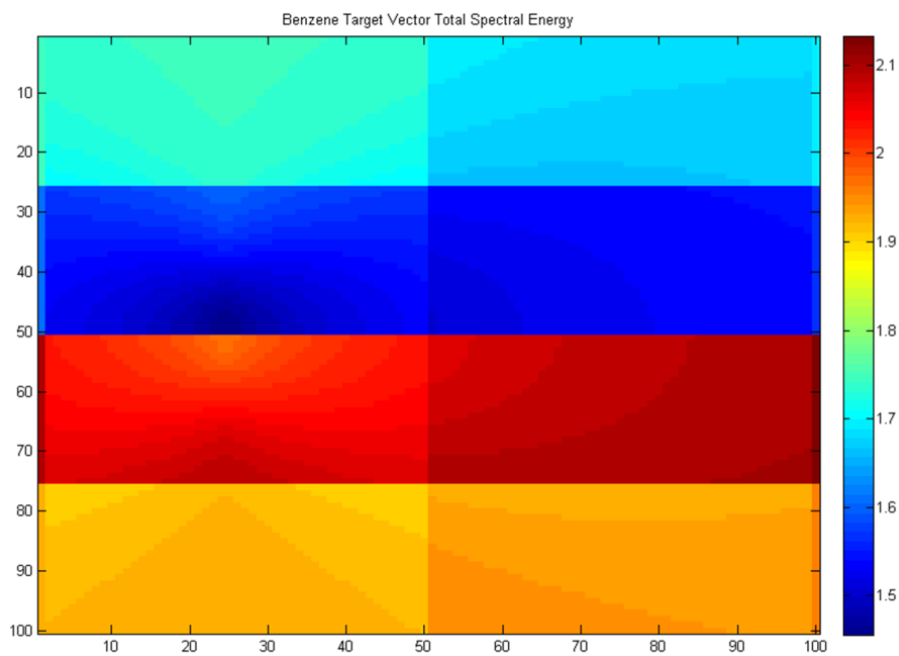

c) Benzene target vector total spectral energy summation

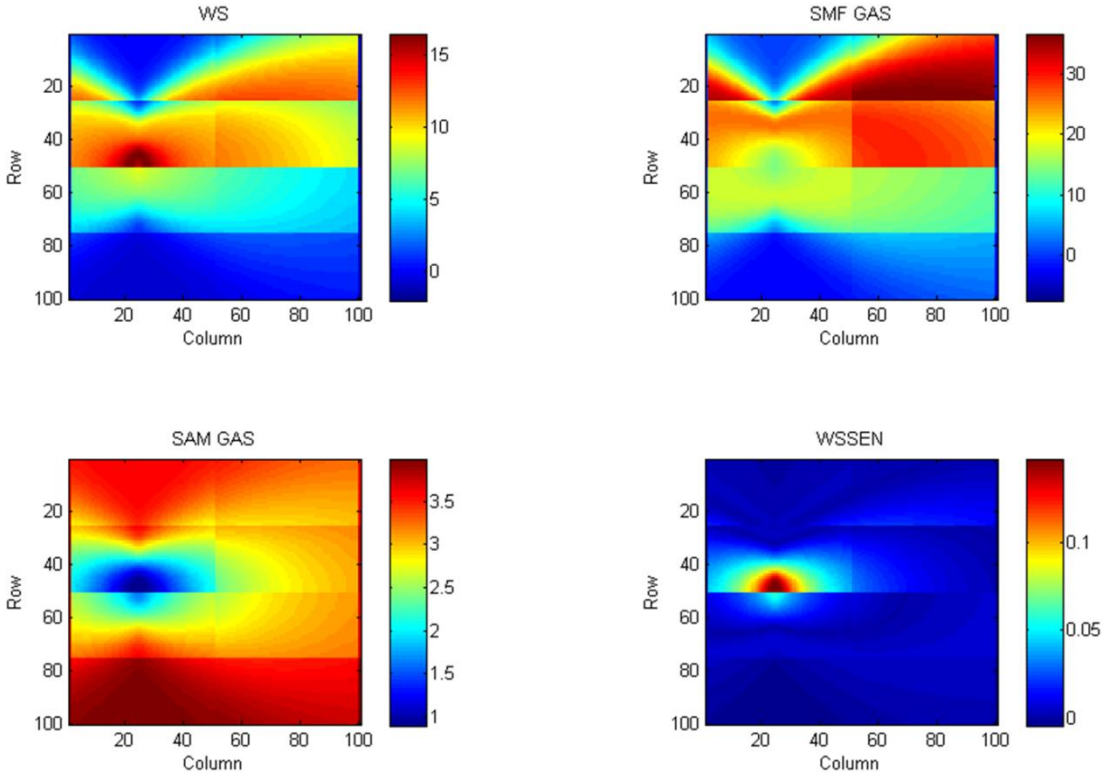

d) Detection plane images for segmented gas plume spectral image

FIGURE 3.53 FOUR PANEL SERIES OF GAS PLUME/BACKGROUND/ATMOSPHERE STUDY

The first panel in the series, panel *a*, shows a Gaussian model gas plume consisting of benzene as if it were emanating from a pipe with a length-wise crack in it as seen from above. The maximum CPL is found at the center of the source of the plume and is 10,000 ppm-m and decays from there. Note that the outermost columns contain no gas and serve as control strips that span all background materials and all atmospheres.

The next panel, *b*, shows how the remaining environmental variables in the scene were divided amongst the plume as previously described. This was not a scientific test in its truest sense in that not all CPL regions of the plume experience the same range of atmosphere and background variability, so no direct quantitative comparisons can be made between quadrants. The third panel, *c*, shows the total summed spectral energy over the target vector band. It can be seen from this example that atmosphere has a greater impact on SRR over this range than does the background for this particular target gas example.

Finally, the last panel, *d*, shows the detection metric image planes. This essentially shows how the sensitivity of each metric is dependent on not only the CPL values, but also atmosphere and to a far lesser extent, background. Note how SAM is essentially insensitive to changes in background material within an atmospheric band, while SMF shows much more sensitivity. Another way of looking at the detection results is to group by atmosphere in that the only difference between the environment of the first two strips is a doubling of water content, which results in an overall lower

transmission. The second half differ from the first half in the presence of much more ozone and then the increased path length resulting in decreased transmission. Lastly, the two lower half strips differ from each other in water content. Note the control strips at either end of the plume correctly show a zero detection score and very high SAM score.

3.3.2 Increased Variability Scene

The increased variability scene was constructed to do just that, while adding a sense of realism to the detection scenario. Various Gaussian modeled gas plumes were injected in this scene. As there was no wind model, all plumes faced the same direction for consistency. Plumes were overlaid to determine confuser gas impacts. Temperatures assigned to various structures were varied to get the thermal contrast to a level where a plume would have regions in absorption as well as emission. These scenes presented the first attempt at realism and variability to test the effect on forming an interferogram from multiple disparate scene elements in the scanning model and determining the effect. Figure 3.54 shows an example detection image plane from this study. This detection image was collected at 60 Hz, 50 m/s, with a critically sampled double sided interferogram, at maximum resolution.

The area of interest is the upper right quadrant of the scene where a benzene gas leak has occurred at the junction of pipes. The plume is blown north over the large pipe that sits above a concrete pad. The concrete pad was assigned a much cooler temperature as compared to the pipe. Only two narrow strips of concrete are visible in the scene where the pipe and sand comprise the remaining surface area. In following the plume over the strips, initially the gas is cooler than the concrete at the release point and absorption is evident, but as the plume progresses, it warms and a clear transition from absorption to emission can be seen. The negative score values for SMF show that the plume has gone into emission and validate that SMF can provide reliable detection metrics simultaneously for absorption and detection events without a change to the target vector. SMF and therefore, the combination metrics, correctly ignore the presence of the two other plumes in the scene. The SAM metric however, has highlighted the anomalous nature of ammonia coupled with the comparatively low emissivity of the aluminum roof over the shed and keyed on this region, not as a target match, but as a clear anomaly. This particular scene lacked the thermal variability within material classes as employed in later versions, but is still illustrative of how much information can be gleaned from a carefully designed scene.

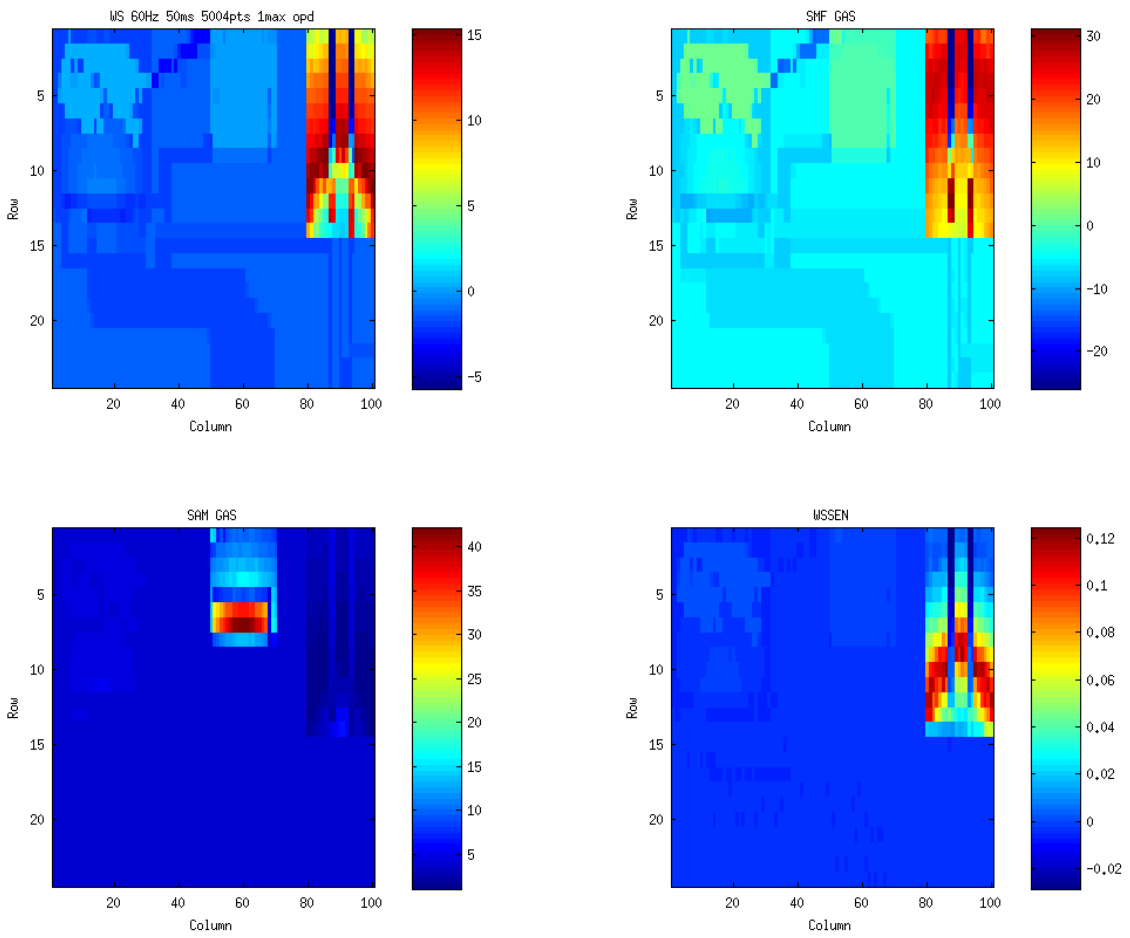

FIGURE 3.54 DETECTION PLANE IMAGE OF SCENE WITH PLUME IN SIMULTANEOUS ABSORPTION AND EMISSION

3.3.3 DIRSIG Scenes

The DIRSIG scenes were developed by a Staff Research Scientist with the DIRS group at RIT, Erin Ontiveros. Erin developed these scenes with the author's input and specifications on gases, temperature profiles, locations, atmosphere, altitudes, desired CPL profiles, and sampled spectral and spatial resolution. The model uses the same PNNL gases used in the earlier scenes and incorporates one of the first gas detection characterizations of the Blackadar model derived gas plume in a complex scene such as that derived from the Megascene Tile 1 resource (Brown, 2010).

The DIRSIG produced scenes provided a more realistic surface and plume model for characterizing the FTS performance at various scan rates. The scenes used all consisted of the same 128x128 image from Megascene Tile1 and contained a gas plume of either methyl chloride or benzene, as two plumes cannot be generated in the same scene as of this writing. The plumes are released either in the open field or in the trees, directly north of the field as shown previously. The image cubes were generated from a nadir viewing geometry with an atmosphere and IFOV set to match each of the atmospheric models used in this work.

Both benzene and methyl chloride gas plume containing scenes were 'flown' against scan rates consisting of 2, 5, 10, 20, and 40 Hz for each atmosphere at 1 and 5 km altitude and then all runs were completed again with mirror error present.

The results of these runs were displayed as detection plane images as done previously, but also in the form of Receiver-Operating Characteristic (ROC) curves given by a detect rate and FA rate. Truth masks were developed for each scan rate so that the appropriate binning of the plume would match that of the data. The truth masks were then used to determine if a detection score was associated with a pixel inside or outside of the plume. The plume was considered to comprise any area in the scene where the CPL was nonzero – the most stressing detection case. Next, each pixel score in the scene was rank ordered from greatest to smallest and subjected to a filter that counted the pixel as a detection or a FA as a function of the total number of elements in the scene given a detection or FA rate. Note that these are rates and not the usually probability of detection or FA in that these values only represent the results from one image and not a distribution of images. This ROC curve is not meant to characterize the efficacy of the detection metric, but rather show change amongst the MOPD and scan rates that were varied between 'flights.' Figure 3.55 shows an example of a ROC curve covering the entire range of values.

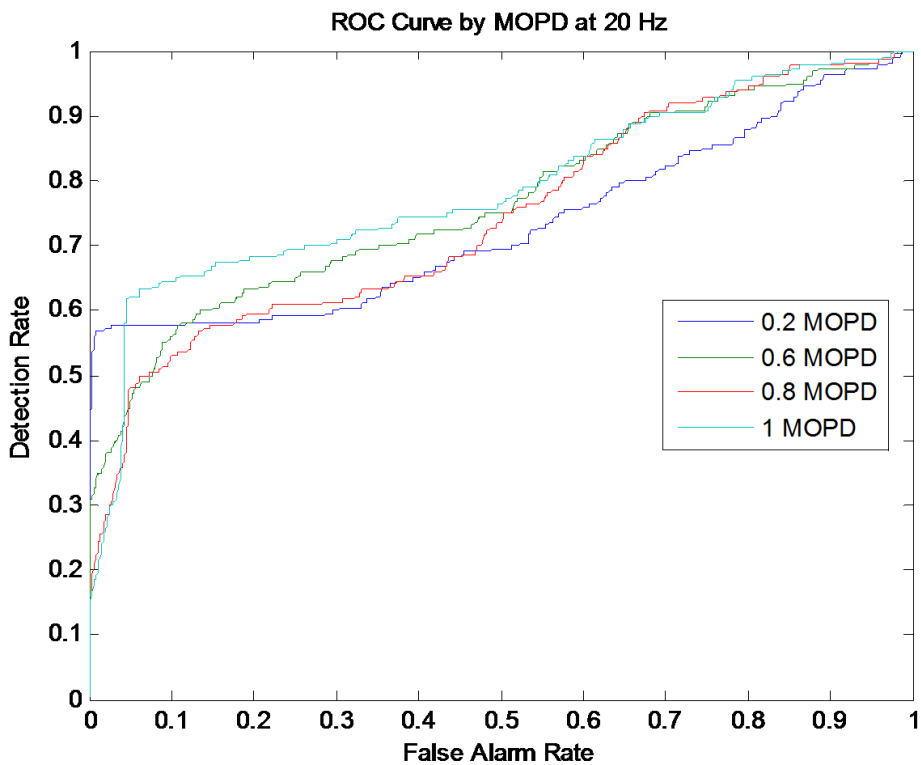

FIGURE 3.55 SAMPLE ROC CURVE FOR WSSEN METRIC AT 1KM MLS, BENZENE GAS IN TREE LINE AT 20 HZ SCAN RATE, VARIED BY MOPD

4 Results

This chapter will provide selected results from each of the major studies completed during this work. They will begin with environmental, instrumental, and error impacts to spectral detection scores to understand the magnitude of the effect of each parameter. Once an operational envelope for each parameter is determined, it is employed in a performance trade study. Finally, the performance trade study results, which illustrate the most important parameter impacts to overall system performance, are further evaluated under complex scenes to measure performance under more realistic conditions.

As a point of departure for this work, the current advanced airborne FTS systems state-of-the-art performance limits were used in shaping the envelope of the performance parameters examined. These systems are covered in section 2.6.1 and 2.6.2. The TELOPS system essentially shows a predicted maximum spectral resolution of 4 and 1 cm^{-1} at 1 and 5 km altitudes, respectively. They also predicted a 0.1 Hz datacube rate at 1 cm^{-1} resolution. The TurboFT system reports 100 Hz rates at 4 cm^{-1} resolution. The rates and resolutions considered in this work span the 1 to 5 cm^{-1} resolution range from 2 to 60 Hz at both 1 and 5 km. However, it must be noted that neither real system has been shown in the literature to achieve meaningful results at their suggested high resolution and scan rate performance limits in the airborne environment.

As mentioned previously, it is also important to note that these results reflect only ideal optical conditions, a noise-free environment (apart from mirror and jitter error, if present), no losses to quantization, a uniform system response of unity across the spectral band of interest, phase corrected interferograms unless stated otherwise, and no thermal contributions from the instrument. These effects can easily be added to the model in the future if specific systems or system components need to be modeled.

4.1 Performance Trade Study Selected Results

In building up to the performance trade studies, the impact of individual parameters needed to be established. One important question that needed to be examined was the effect of background and atmosphere on the collected target spectrum. It was posited in this work that the background is largely invariant in a spectral sense at high resolution over a region of a given gas spectral feature and has very little impact to the detection of the target gas spectrum. However, as previously stated, the variation in emissivity associated with the varying backgrounds can have influence over the spectral shape of the target gas in effectively shifting the thermal contrast. This shift results in a muted absorption feature, a loss of thermal contrast, or a transition to emission – all adversely affecting the resulting metrics.

4.1.1 Atmospheric Profile and Background Characterization Study

To determine the full impact that the background and atmosphere have on a target gas spectrum, a CPL versus gas temperature trade study was used to evaluate the dependence of the detection metrics on the influence of these variables. In this study, benzene and methyl chloride were each run over a range of temperatures from 280 to 300K and a CPL range of 0-30,000 ppm-m for benzene and 0-50,000 ppm-m for methyl chloride. This tradespace was repeated over each surface material and atmospheric model used in this work. The target vector was a 290K target gas over a 300K black body background at 10,000 and 20,000 ppm-m for benzene and 30,000 ppm-m for methyl chloride. These target vector CPLs were derived to give approximately equal absorption strength between the target gases. The peak score from the gas detection metrics was also recorded to show the variability in the metrics. Figure 4.1 is an example of a series selected from the study which shows the progression of this work through the various atmospheres for the gas plume over only the 300K asphalt background material using the 290K benzene 20,000 ppm-m target vector. The same series were reproduced for every atmosphere and background used in this work using various target vectors.

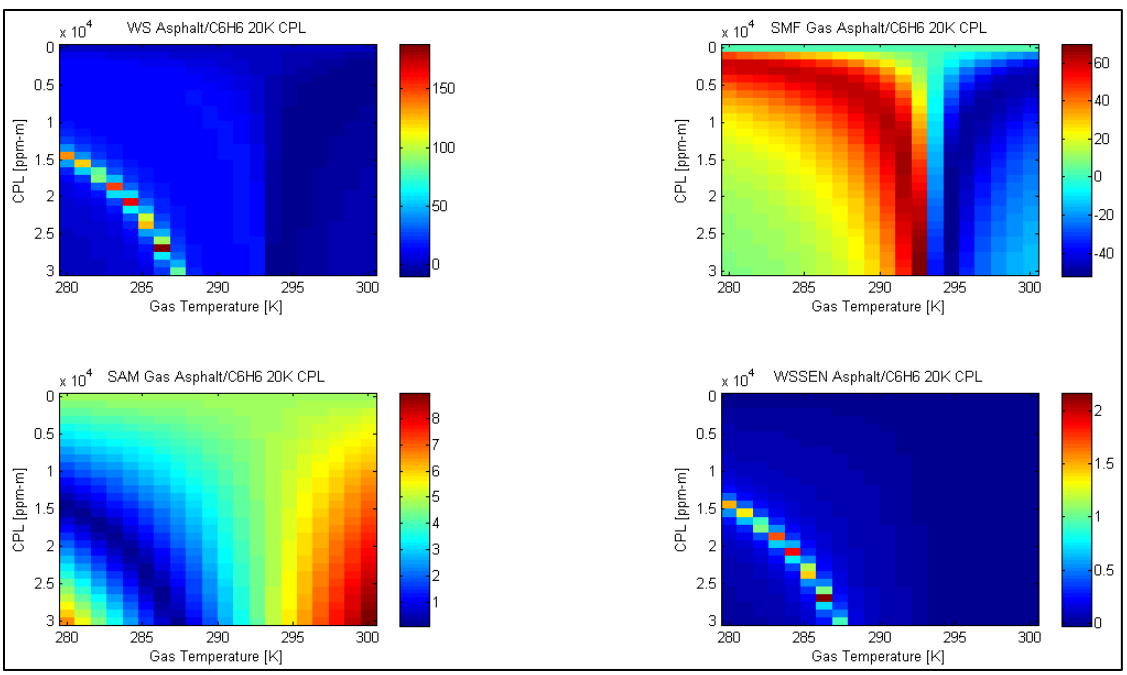

a) No atmosphere

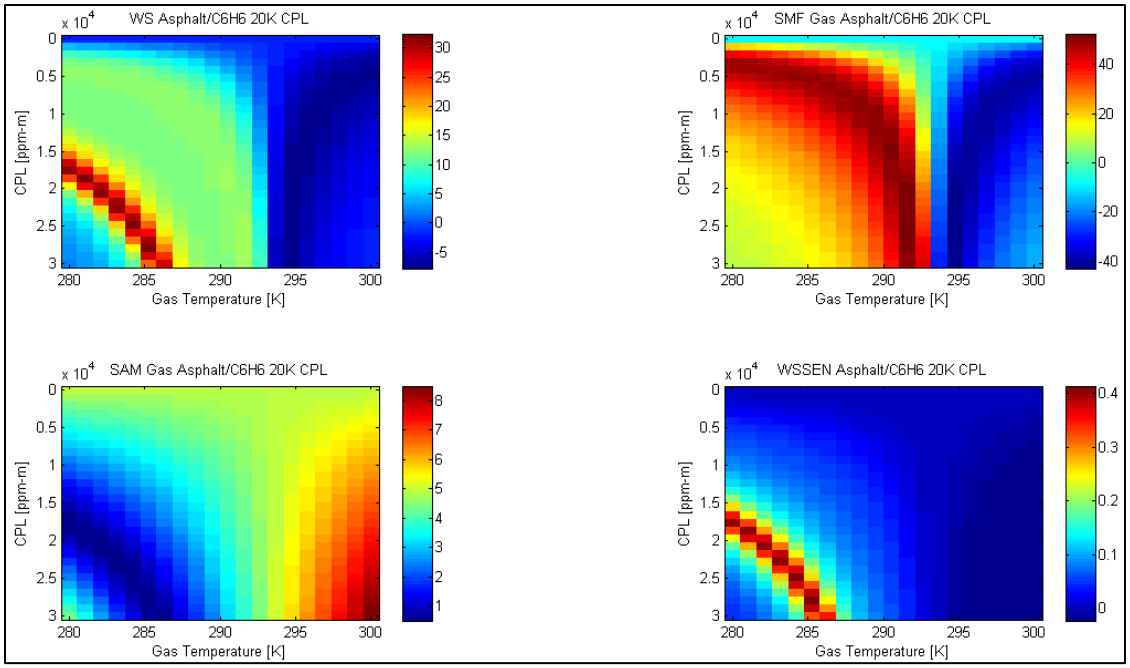

b) 1km MLS atmospheric profile

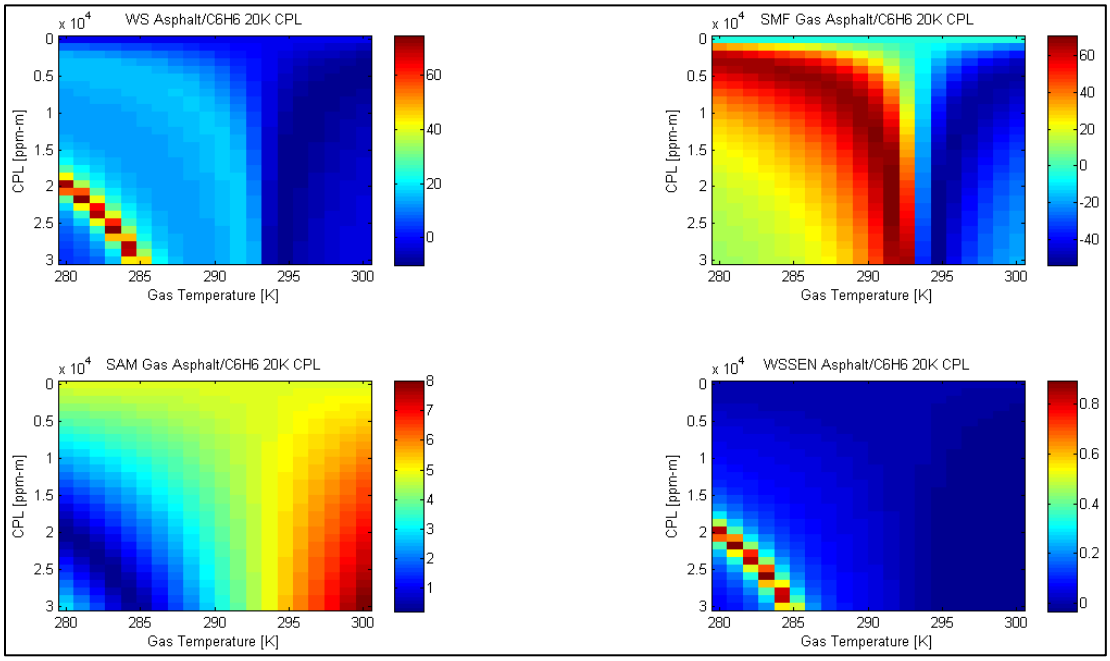

c) 5km MLS atmospheric profile

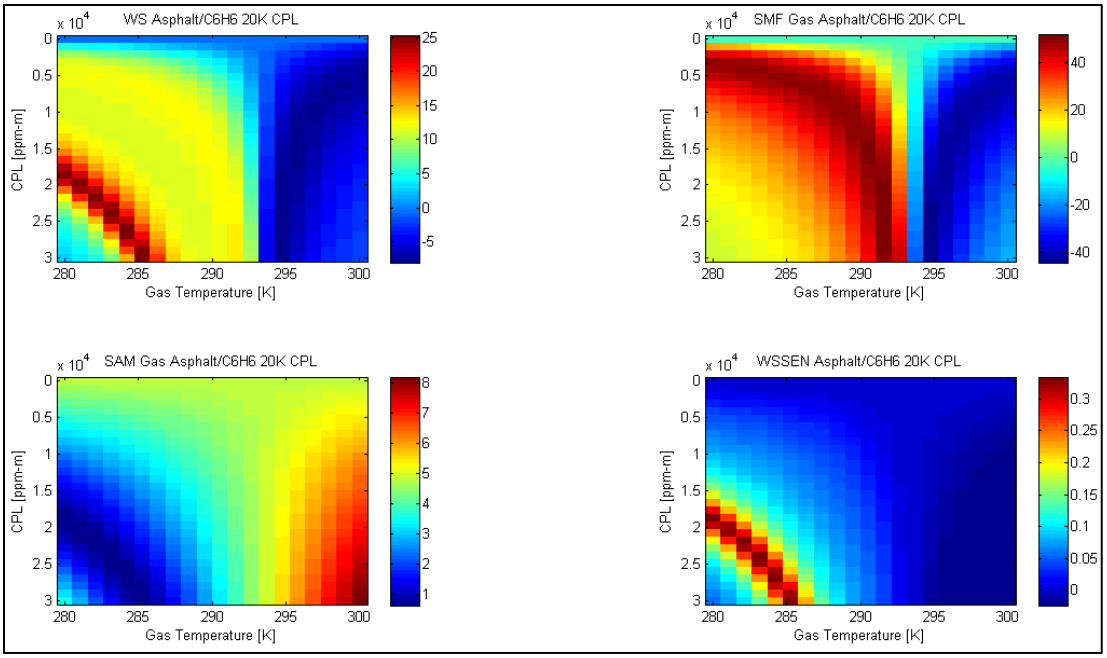

d) 1km tropical atmospheric profile

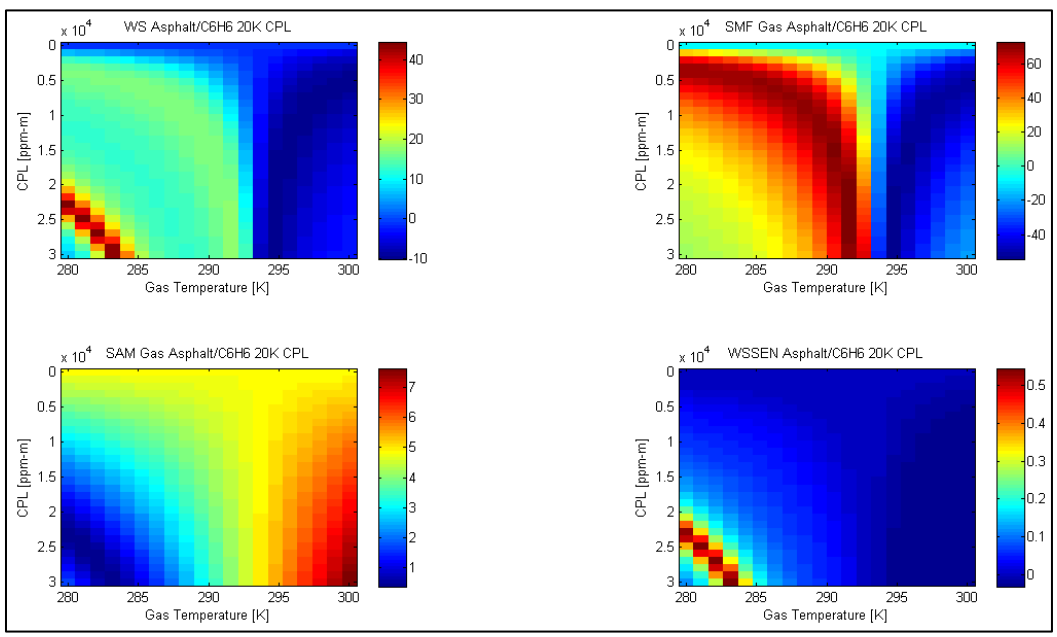

e) 5km tropical atmospheric profile

FIGURE 4.1 CPL VERSUS GAS TEMPERATURE SERIES BY ATMOSPHERE FOR 20,000 PPM-M BENZENE

The first panel, labeled *a*, of four detection planes shows the results with no atmosphere, the second, *b* is with a 1 km MLS, the third, *c* with a 5km MLS, the fourth, *d* with a 1 km tropical, and the final panel, *e* with a 5km tropical atmospheric profile. The spectral ratio (WS) detection plane can be used to follow the trend through the series.

In the no atmosphere panel, the identification of the proper match to the target vector is quite clear and shows a relationship of varying spectral strength to achieve similar spectral feature shape and magnitude to arrive at similar detection scores despite widely disparate collection conditions. This is an important point for future gas plume studies. There are many competing factors that affect the shape of the target spectral feature. As the thermal contrast decreases, the plume moves from absorption to emission as is most evident in the SMF image in the top right panel of each series. The plume slowly decreases in absorption strength with increasing temperature for a given concentration, reducing the spectral feature. Increasing the concentration has the opposite effect and sets up a competing relationship. Also, it should be noted that the target vector was collected over a blackbody background. A blackbody has a constant emissivity of one, meaning that the transition point from absorption to emission occurs when the thermal contrast is 0. However, the emissivity of asphalt at the benzene spectral feature region is about 0.9, meaning the transition point is moved to a cooler gas temperature than the background. Thus, detecting a target much warmer than 290K in this scenario is difficult in that the absorption feature rapidly diminishes with increasing temperature. Thus, following the suggested rule of thumb, greater concentration is required to achieve a similar spectral feature at a lower gas temperature and each of the panels in the series reflects this. This is an important point to consider when determining where in the plume one would want to detect, which has a temperature dependence on how well mixed it has become with the atmosphere, and requires a target vector adjusted accordingly. Finally, the increase in score between 1 and 5 km atmospheres is due to the increase in ozone, inflating the SMF score, which in turn, affects the WS and WSSEN score to a smaller extent.

As the atmosphere is thickened, the ambiguity between spectral feature and atmosphere causes a greater spread and general overall reduction in the detection scores over the same range. The “T” shape seen in the detection planes is a combination of the 0 CPL case registering no score in the horizontal, while the vertical strip on the right shows the transition from absorption to emission for the gas and explains the negative scores in the SMF and WS metrics.

Figure 4.2 shows the summary of the results of running each of the detection metrics against a target gas plume over several of the surface materials used in this work coupled first with no atmosphere, then each of the atmospheric models used in this work. The surfaces are at 300K with the scores derived from benzene gas at 290K and 10,000 ppm-m. The target vector used for the detection metrics is the same temperature and CPL as the simulated gas. Figure 4.3 shows the same series with a 290K and 20,000 ppm-m benzene target vector. Note from the detection planes shown above, the target vector values of CPL and gas temperature used to make the following charts are not necessarily where the maximum detection takes place.

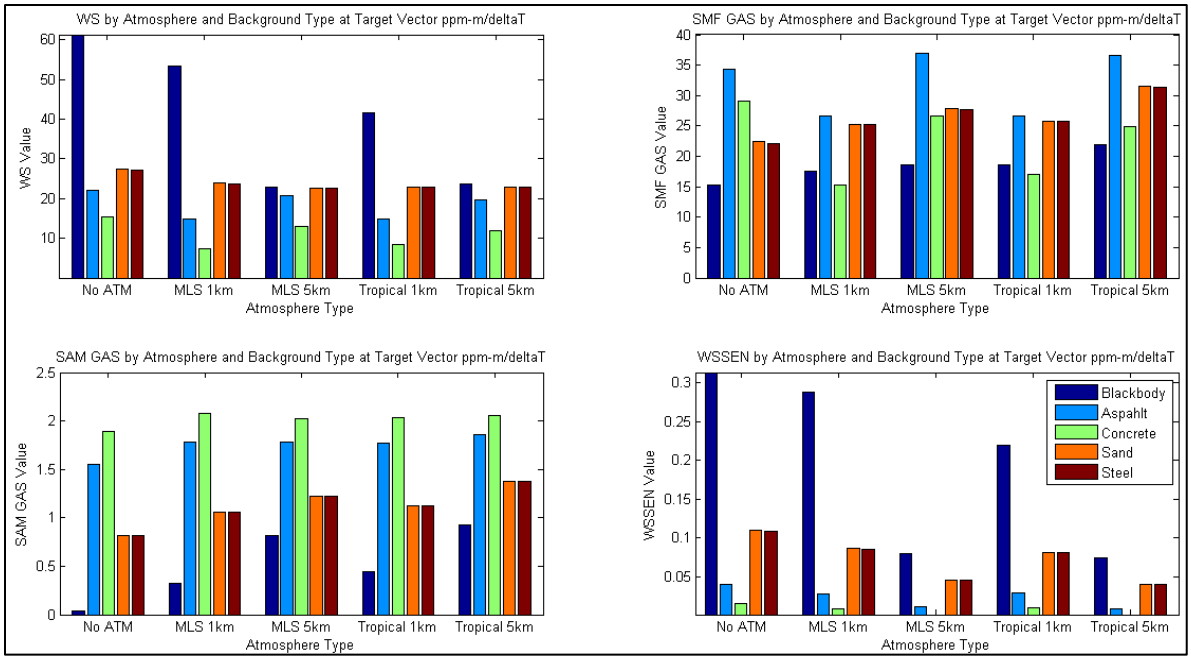

FIGURE 4.2 DETECTION SCORES BY BACKGROUND EMISSIVITY AND ATMOSPHERE FOR 10,000 PPM-M 290K BENZENE TARGET

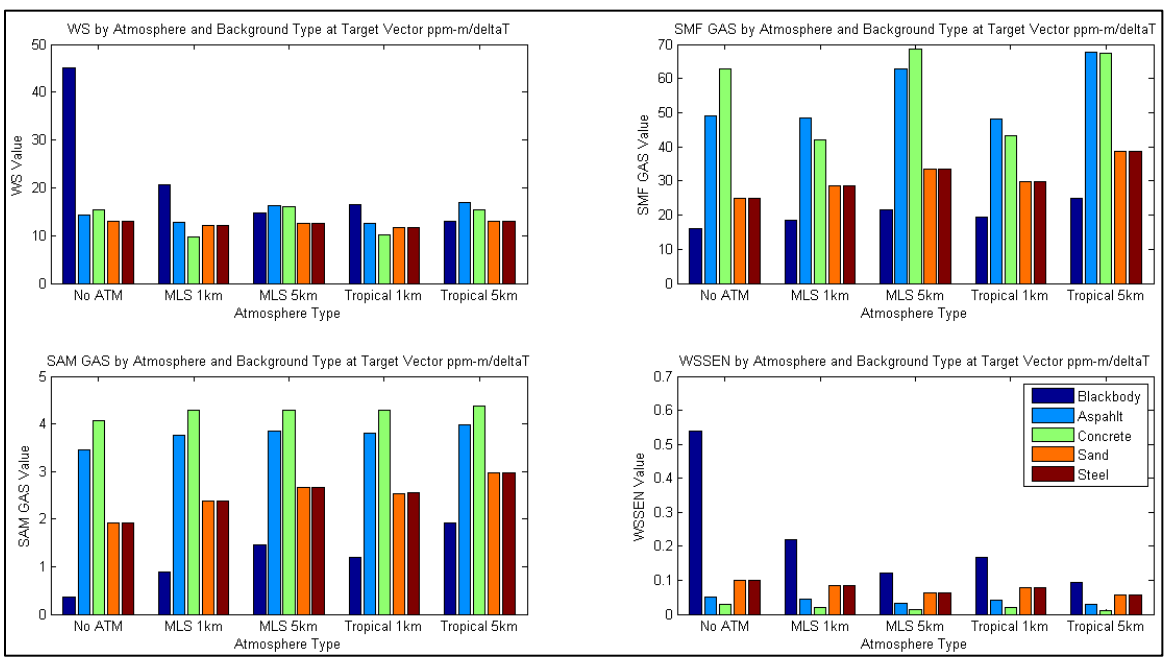

FIGURE 4.3 DETECTION SCORES BY BACKGROUND EMISSIVITY AND ATMOSPHERE FOR 20,000 PPM-M 290K BENZENE TARGET

The trends are probably best captured by the bar chart from the Spectral Ratio (WS) detection metric in the upper left. The bars consist of the following materials at 300K from left to right: a blackbody, asphalt, concrete, sand, and steel. The first series, on the left, shows the score results with no atmosphere – essentially the gas leaving radiance. The result of the WS metric shows that the background, and even atmosphere at this target level, are largely invariant. Note that as the CPL

of the gas increases, the background material and atmospheric effect are slightly muted, as shown by the spectral ratio.

The results in general show that the background is almost as important as the atmosphere in gas detection in this scenario as the variability between atmospheres is nearly the same as that between background types, depending on score metric. The general atmospheric trend for benzene detection at these CPLs is that the lower the ozone contamination, the better the score. As ozone concentration increases from 1 to 5 km, the corresponding effect is reflected in the scores by atmosphere. The impact of increased water vapor between the MLS and tropical models causes only a minor reduction in score at this CPL level and that is due to the absence of major water absorption lines in the spectral feature region.

While benzene is only present over a narrow range of the total spectrum, methyl chloride is found across a much larger portion of the spectrum and serves as a good contrasting test case to benzene. Figure 4.4 illustrates the results under the same conditions.

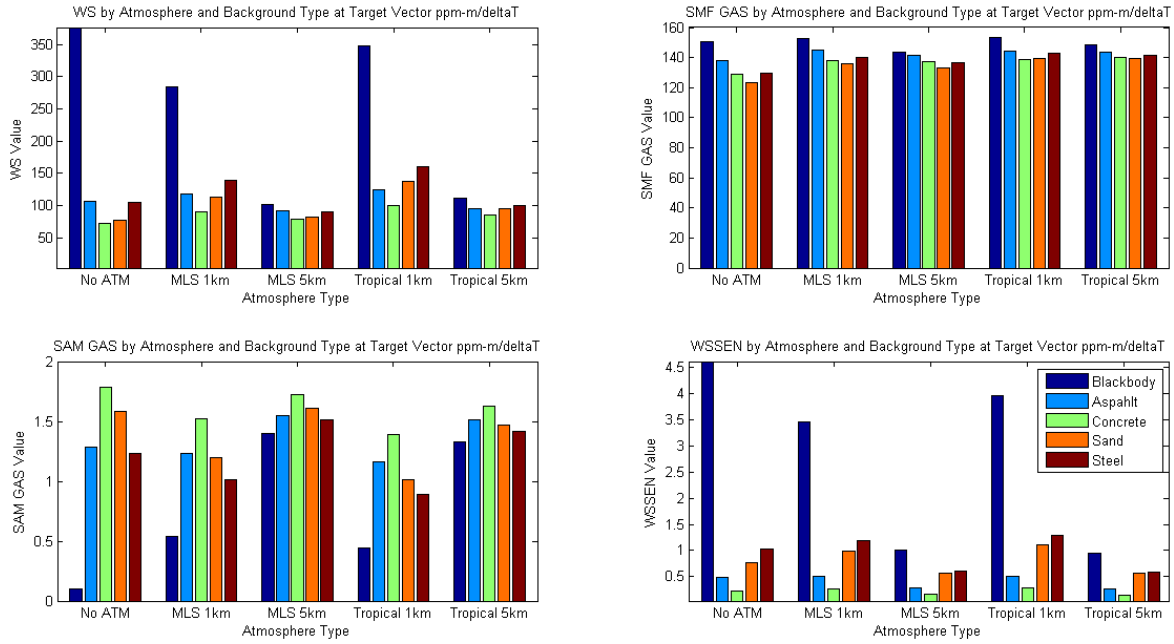

FIGURE 4.4 DETECTION SCORE TRENDS BY MATERIAL AND ATMOSPHERE FOR METHYL CHLORIDE

The increase in score with the tropical profile over MLS is due to the increased presence of the rotational lines from water that are interspersed with the rotational manifold of methyl chloride. The detection scores in general are larger in magnitude as compared to benzene due to the greater amount of spectral energy, and they essentially follow the same trend, even when the spectral features span a much larger portion of the collected spectrum. The charts also show that the reduction in detection score is commensurate with the overall transmission of the given atmosphere, reducing the total spectral energy between 1 and 5 kilometers.

Figure 4.5 summarizes the maximum score results for benzene. Unfortunately, with these kinds of studies, telling the whole story in one plot is not possible. While the plot follows the peak scores for each detection plane across the atmosphere and surface material series, it does not show how the peak value coordinates changed across the same metric in different atmospheres. The peak value is useful in that it shows one's best bet at detecting a gas of the specified strength in those conditions, but it does not show the trend of the strength with increasing atmospheric contamination. The general trend is that greater values of CPL and/or thermal contrast are required for the spectral feature to overcome the presence of the masking contaminant.

Figure 4.5 shows results for the same study taken at the maximum score position at each detection metric. Note that each score metric value does not necessarily correspond to the same CPL and gas temperature coordinate in each plane. This chart essentially reflects the required conditions to give maximum detection performance for each surface material and atmosphere.

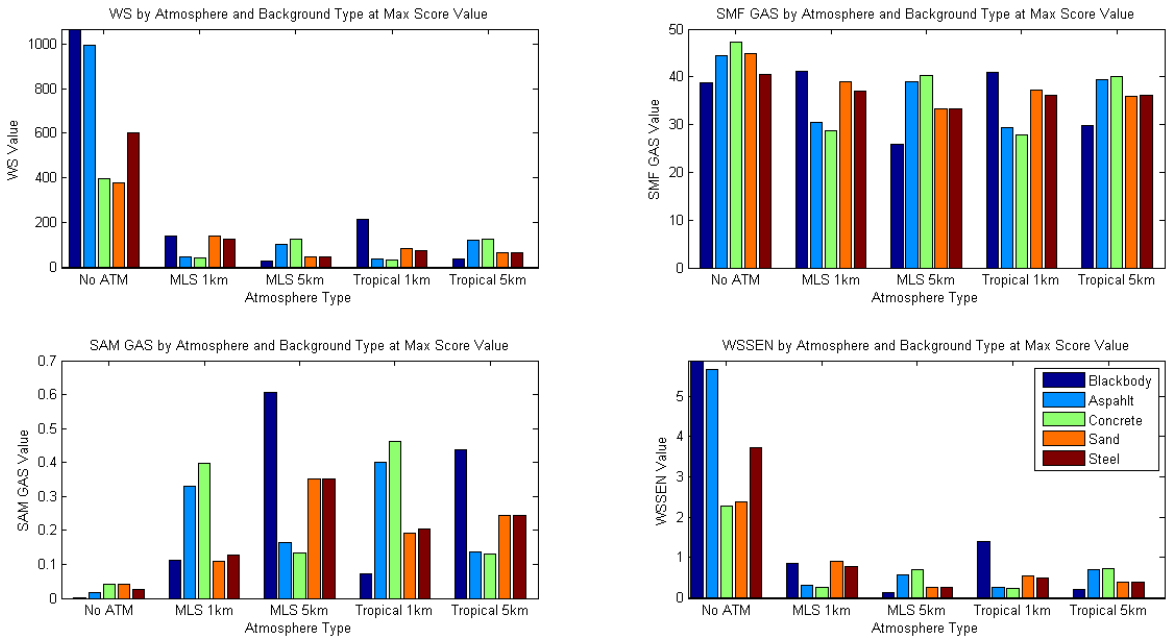

FIGURE 4.5 DETECTION SCORE TRENDS FOR BENZENE BY ATMOSPHERE AND SURFACE MATERIAL, PEAK SCORE METRIC

The biggest result from this plot is further evidence that water vapor has only minor impact on benzene detection. This can be seen in the comparison of both the 1 and 5 km atmosphere models between the MLS and tropical profiles. The major impact for benzene comes from the presence of ozone and that is increased with increasing altitude. Hence, the jump from 1 to 5 km makes the most impact, regardless of water content. Figure 4.6 shows the chart from the same study for methyl chloride which tells a different story.

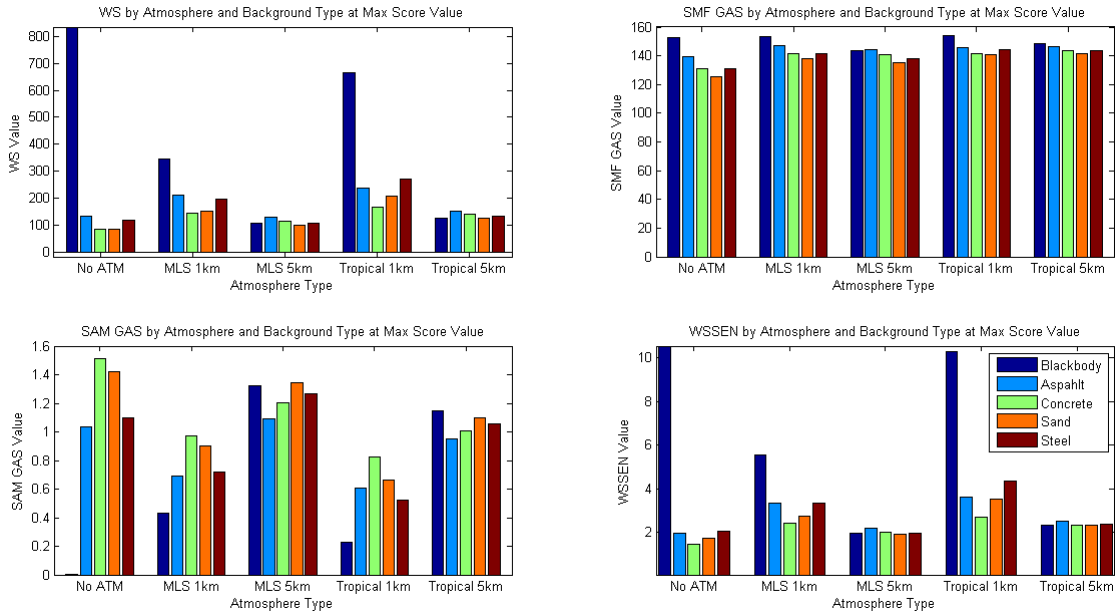

FIGURE 4.6 DETECTION SCORE TRENDS FOR METHYL CHLORIDE BY ATMOSPHERE AND SURFACE MATERIAL, PEAK SCORE METRIC

The results here show that methyl chloride, which appears over a much larger portion of the collected spectrum, is much more susceptible to water vapor content than benzene. The water rotational lines look very much like methyl chloride and are essentially interspersed among the spectral features. The trend shows that as the water vapor content increases from MLS to tropical profiles, then from 1 to 5km where total column content is increased, the scores actually improve. Figure 4.7 shows a comparison of the gas target vector and the actual transformed SRR at a 5km Tropical atmosphere with the gas at 11,500 ppm-m at 290K over a 300K asphalt background. It can be seen that the rotational lines of water that overlap the spectral features change their shape and affect the resulting detection scores.

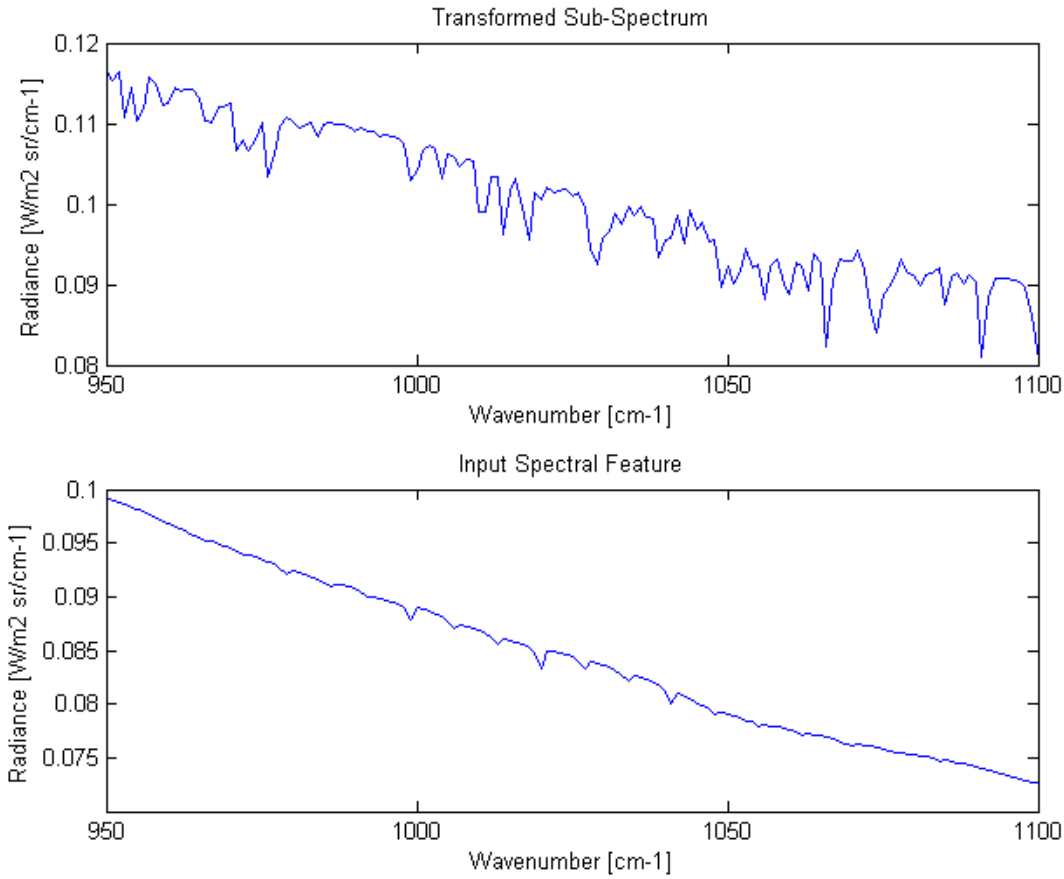

FIGURE 4.7 METHYL CHLORIDE SPECTRAL FEATURE CONTAMINATION BY ROTATIONAL LINES OF WATER (TOP), TARGET VECTOR SPECTRUM (BOTTOM)

4.1.2 Atmospheric Component and FTS Error Series Study

This study completed a series of benzene target gas detections by spectral score for a progression through the major atmospheric components, then through both jitter and mirror error. The first series is the gas leaving radiance, then atmospheric transmission only, followed by the transmission with down-welled radiance only, then transmission with up-welled radiance only, then all three – essentially the full atmosphere. Next the error series was executed with full atmosphere and jitter, full atmosphere and mirror, then all three. The simulation was run with a 290K, 20,000 ppm-m benzene target vector over a 300K blackbody. The simulation settings were set for full resolution, a critically sampled interferogram with 290K benzene gas at 20,000 ppm-m over 300K asphalt and a 1 km MLS atmospheric profile. The SAM score is for the entire spectrum to show how the overall spectrum is affected, while the remaining scores are for the target vector region. While the 'no atmosphere' results can be compared as the baseline to the trials varied by atmosphere, the 'full atmosphere' trial scores should be used as the baseline for the FTS error series.

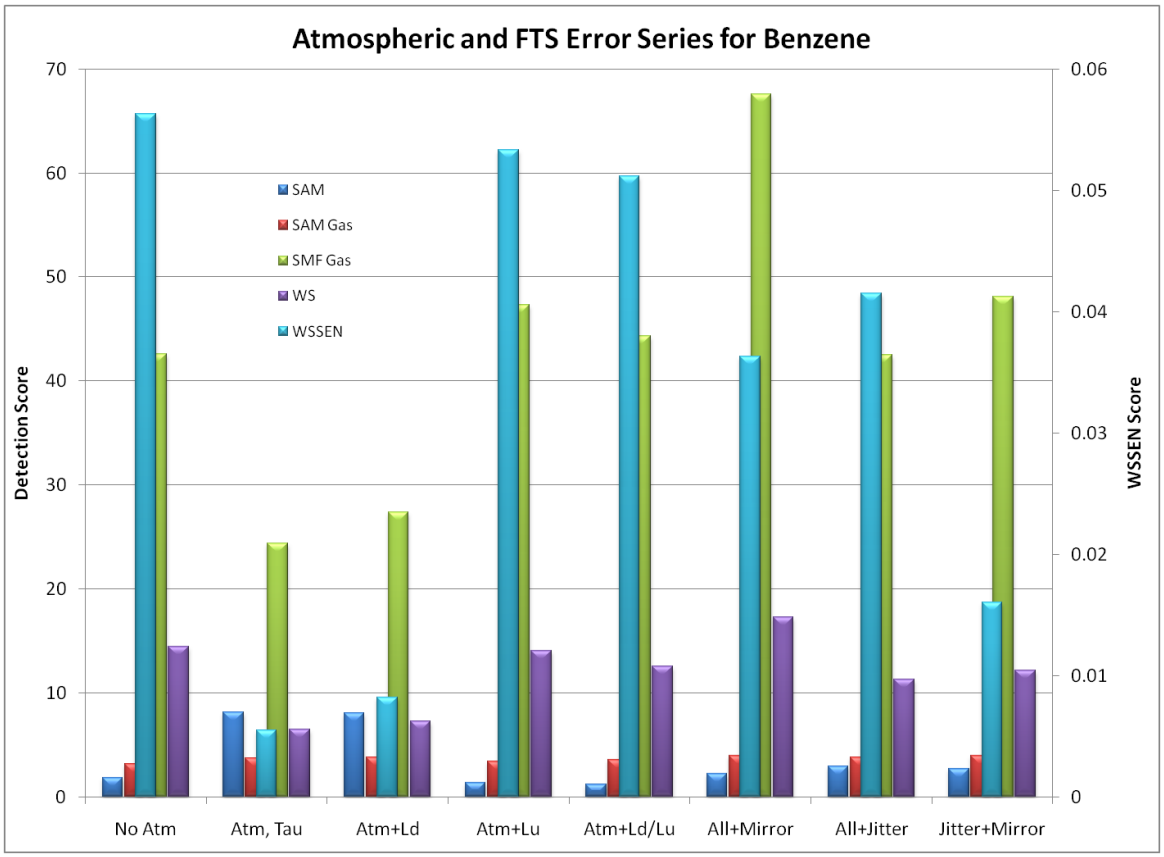

FIGURE 4.8 ATMOSPHERIC AND FTS ERROR SERIES FOR BENZENE

Figure 4.8 shows that the atmosphere in general has a large influence on the detection scores and further shows the magnitude of the up and down-welled radiance separately in terms of their impact to the detection score. While up and down-welled radiance are a natural component of real data collects, this series shows the minor effect in a given detection scenario of getting the down-welled estimates correct, or included at all for that matter, in the detection of benzene. If there was a significantly reflective surface material in the scene, the down-welled term would become much larger, contributing significantly to the benzene target region due to the reflected ozone bands. Any slight improvement of score in the atmosphere and down-welled only trial is from the slight addition of ozone to the target region. Conversely, the up-welled radiance has a major impact to the detection scores as the emissive nature of this term effectively cancels a large portion of the atmospheric absorption effects along the transmission path making the atmosphere more blackbody-like.

The mirror error was set to give a mean of 10% uncertainty in the OPD position recorded for the moving mirror. This was approximately 0.4 µm over a 4 µm sample interval. The effects of the mirror are apparent on the WSSEN score and SMF score. These scores show that while additional spectral energy was injected into the target region, the distribution of that energy did not match that of the target. The SAM score shows a factor of two disruption to the spectrum.

The magnitude of the jitter error resulted in 25% of samples off-target to a scene with no target gas and a 5km MLS atmosphere with 290K sand. The SAM score for jitter shows a disruption of the spectrum of almost a factor of three. When both error types are combined, a major loss to the WSSEN score is seen. Jitter and noise errors, being random in nature, make it difficult to show repeated trends. They do serve to show the weaknesses of the various detection metrics and that the most reliable indicator, the WSSEN score, shows an overall decrease in detection performance.

Figure 4.9 shows the same series for ammonia.

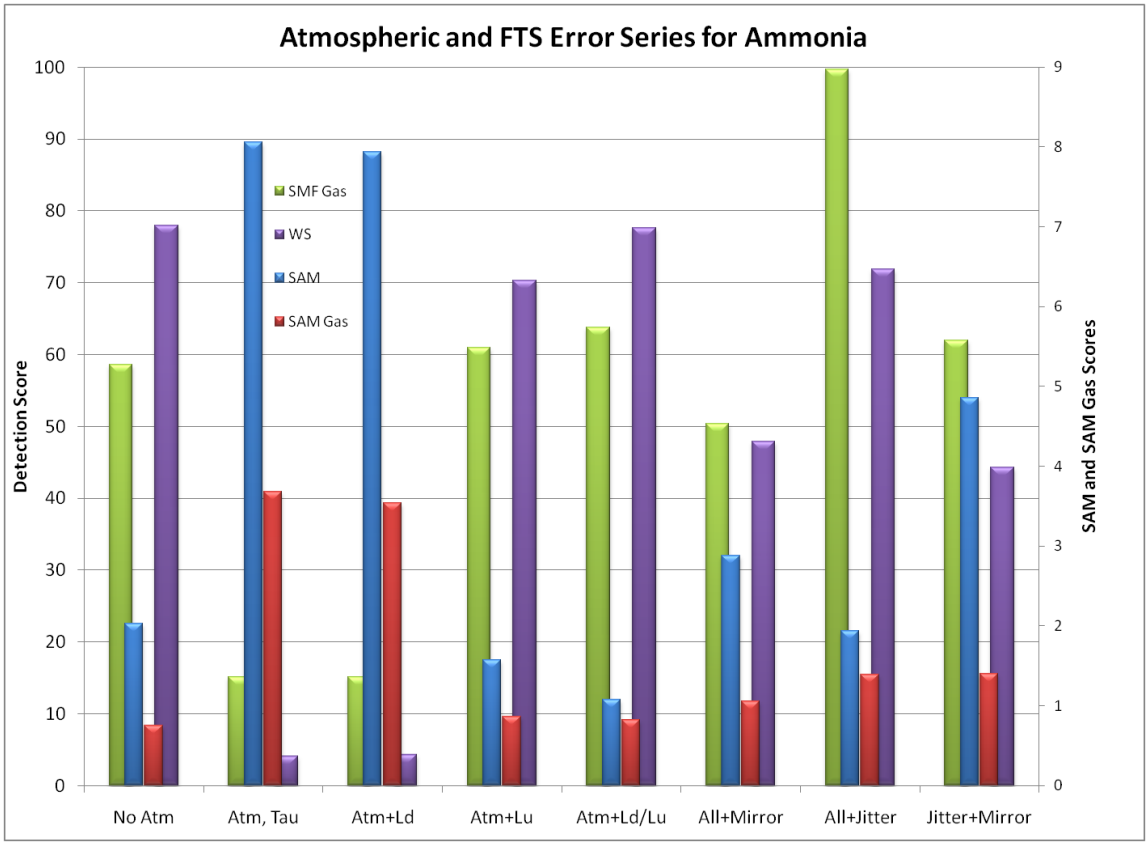

FIGURE 4.9 ATMOSPHERIC AND FTS ERROR SERIES FOR AMMONIA

In this case both SAM scores are assigned to the right axis for better differentiation from the other scores. The WSSEN score is not included in this series as there are too many spectral features to account for. Ammonia is far more sensitive to the atmosphere than the other gases because its spectral features are dispersed across the entire LWIR window and encounter multiple atmospheric lines. The atmosphere series shows the minimal impact of the down-welled radiance, but a much more significant contribution from the up-welled. The detection score here improves some and that is due to the same up-welled radiance effect as described previously.

In the error series section the errors produced significant changes to the scores, particularly for the jitter error. The combination of both errors gave the lowest detection scores.

4.1.3 Confuser Gas Presence Impact to Spectral Feature Detection

As previously mentioned, methyl chloride was chosen as a confuser gas for benzene due to the 3-4 cm^{-1} peak proximity of similar spectral features that stressed the MOPD ranges used in this work, reference Figure 3.23. The selectivity of the detection metrics as well as the effect of varying MOPD could be determined simultaneously from this study. Figure 4.10 shows the improvement in spectral discrimination between the peaks of methyl chloride and benzene at equal strength through a series by MOPD with a 1km MLS atmosphere to minimize the confusion with ozone. The top trace in both panels is the transformed SRR at first 0.2 cm MOPD in the top panel, then 0.5 cm MOPD in the bottom panel. The lower trace in each panel is the SRR at 1 cm MOPD. The peak for benzene can be seen at 1038 cm^{-1}, while the two immediate neighboring peaks on either side are due to methyl chloride. It can be seen that at 0.2 cm MOPD, the three peaks blend into one, while at 0.5 cm MOPD, the ability to detect the presence of all three peaks becomes possible.

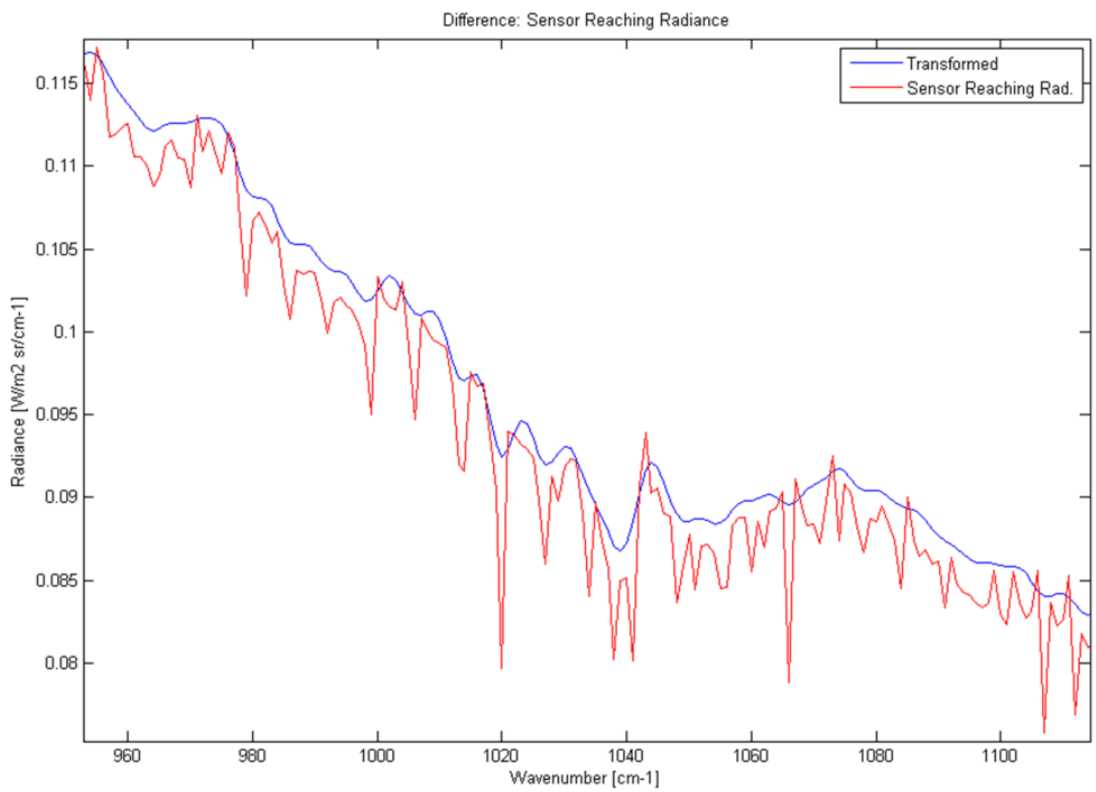

a) 0.2 cm MOPD spectrum (top trace) with 1 cm OPD SRR (bottom trace)

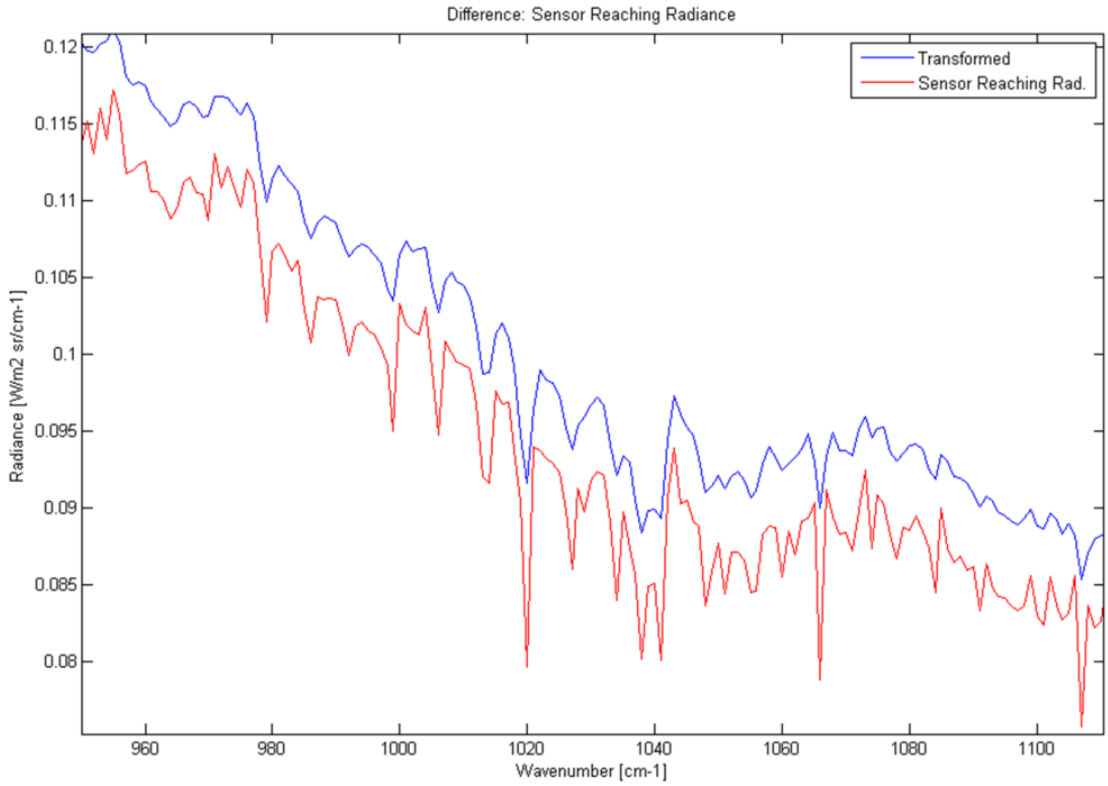

b) 0.5 cm MOPD spectrum (top trace) with 1 cm MOPD SRR (bottom trace)

FIGURE 4.10 BENZENE WITH METHYL CHLORIDE CONFUSER GAS AT 0.2 (A) AND 0.5 CM MOPD (B)

The next series of panels shows the results of injecting methyl chloride into the spectrum at 30,000, 60,000 and 90,000 ppm-m CPL with benzene, yielding absorption strengths approximately equal to ½, 1 and 1.5 times that of benzene. Figure 4.11 shows the first panel top left as a benzene only detection plane series with a 1km MLS atmosphere. Moving to the top right the 30,000 ppm-m methyl chloride gas is added. The bottom left panel shows the effect of adding the 60,000 ppm-m confuser and the bottom right is the 90,000 ppm-m confuser. The target vector for the detection metrics was 4,800 ppm-m benzene at 290K with a 300K black body background.

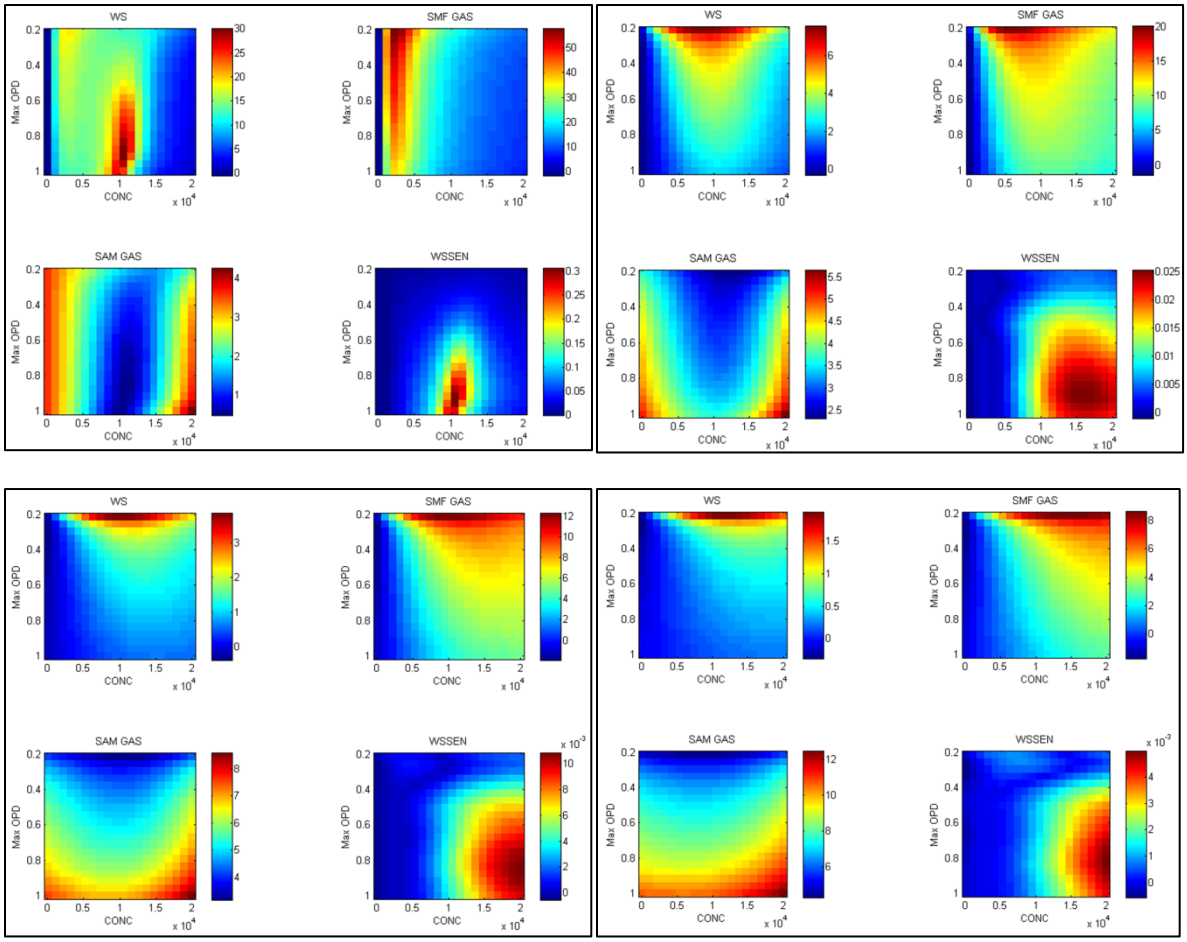

FIGURE 4.11 IMPACT TO DETECTION FROM CH_3Cl CONFUSER GAS SERIES: 4,800 PPM-M BENZENE WITH NO CONFUSER (TOP LEFT), 30,000 PPM-M CONFUSER (TOP RIGHT), 60,000 PPM-M CONFUSER (BOTTOM LEFT), AND 90,000 PPM-M CONFUSER (BOTTOM RIGHT)

The panels illustrate many points. Initially, without confuser gas, the detection metrics due a good job of correctly characterizing the benzene peak at the appropriate resolution as distinct from the neighboring ozone peak at 1039 cm^{-1}. The spectral energy normalized spectral ratio (WSSEN) detection plane shows an isolated detection region that gives the best match to the target vector, requiring a higher CPL value due in part to the reduced emissivity of the background as compared to the target vector.

As one follows the progression of confuser gas CPL through the series it can be seen that greater CPL values must be attained for benzene to emerge above the confuser gas. Additionally, note that virtually all of the remaining detection planes are corrupted by the presence of the confuser gas while the WSSEN metric continues to correctly identify the benzene peak presence.

Figure 4.12 shows a bar chart that captures the peak WSSEN score as a function of atmosphere, FTS error types, and confuser gas. Note that the confuser gas has the largest impact to benzene detection as compared to all of the other variables. The 0.075 in the chart refers to the standard deviation of the mirror error. The jitter 16 and jitter 1024 refer to the jitter periods and represent the high and low frequency jitter errors, respectively. Jitter 16 has a period of recurrence every 16

OPD samples in the interferogram with a mean of 4 samples and standard deviation of 2, giving a mean of 25% of interferogram OPD points representing an off-target pixel. The jitter 1024 has a period of recurrence every 1,024 OPD samples with a mean jitter run length of 300 and standard deviation of 70 samples. This gives a mean of approximately 30% of interferogram OPD samples off-target.

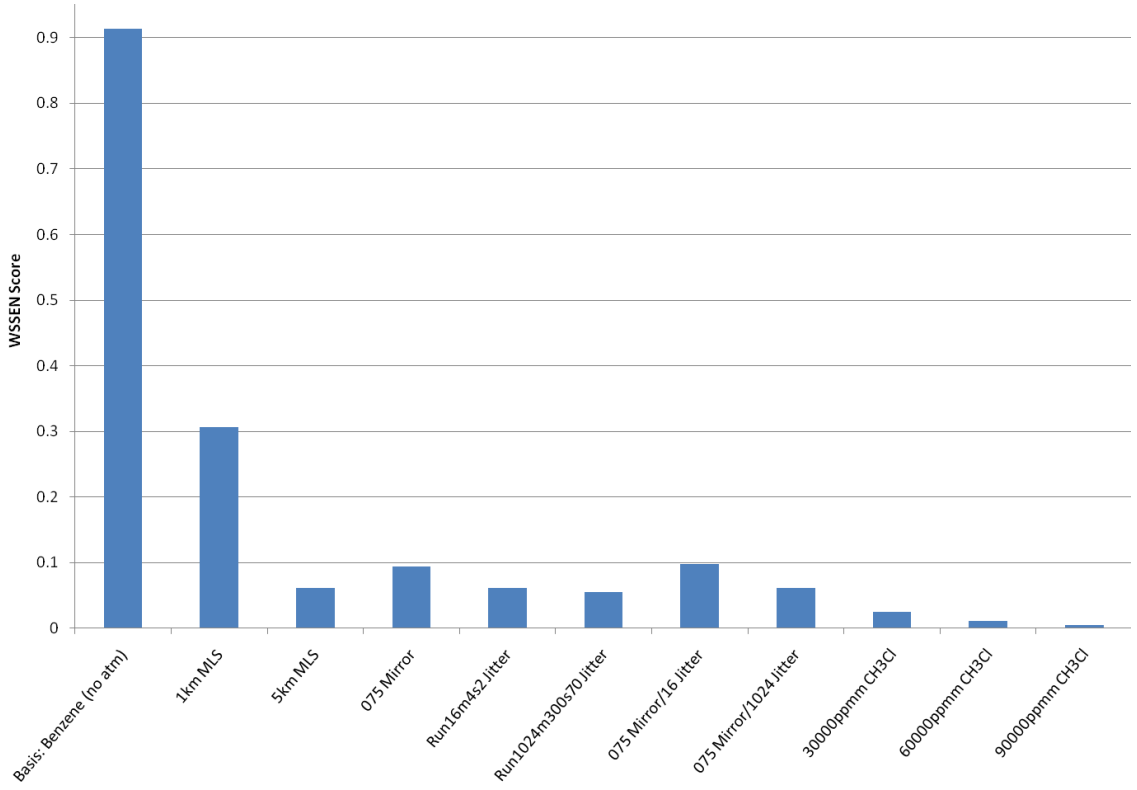

FIGURE 4.12 WSSEN SCORE SERIES RESULTS FOR BENZENE WITH CONFUSER GAS

Unfortunately, tracking the peak WSSEN score does not give information about how the requirement for improved resolution or additional CPL for the gas changed as well. Figure 4.13 shows a vector representation of the WSSEN score given by the width of the vector and the end point of the vector indicating the required MOPD and CPL required to produce the score. Note that the vector has no meaning for values between the origin and the end point; the only components with meaning are the thickness of the vector and the end point. This chart captures both the bar chart shown previously as well as the performance trade information in producing a given score. For example, the thickest vector corresponds to the baseline case of benzene gas with no atmosphere. The large thickness corresponds to the high detection score and the terminus of the vector indicates that this maximum score was achieved with an 8,500 ppm-m CPL and 0.92 cm MOPD. The next largest vector shows the case with a 1km MLS atmosphere. The much thinner bar indicates a much smaller detection score, but the terminus shows an increased CPL requirement of 10,500 ppm-m and resolution of 0.93 cm MOPD to achieve the maximum score under the given conditions. The cases generally show an increased requirement for target gas CPL to compensate for the noise or confuser source introduced.

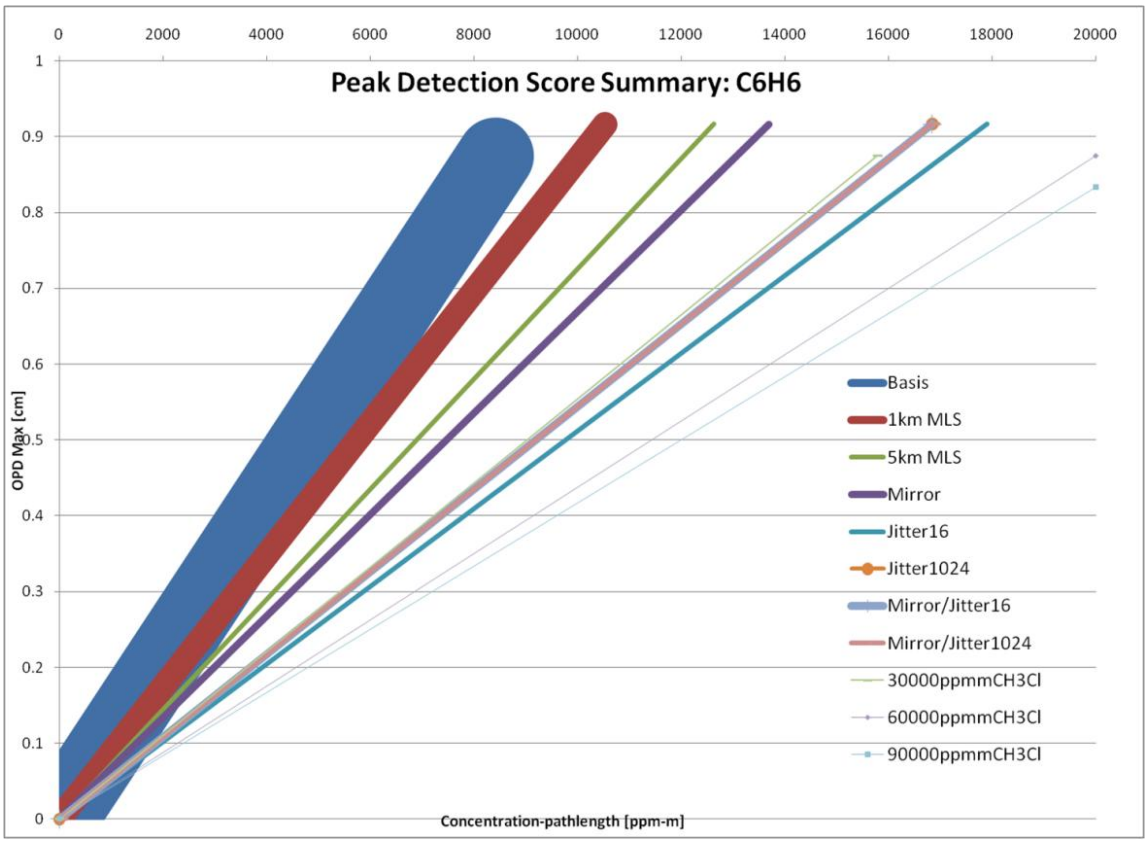

FIGURE 4.13 PEAK DETECTION SCORE VECTOR SUMMARY FOR BENZENE

Figure 4.13 also shows how much more target gas is required to produce the comparatively weak detection for a given amount of confuser gas. Figure 4.14 shows the same vector-style chart for methyl chloride with benzene confuser gas. The chart shows that as expected, the comparatively narrow spectral feature that comprises benzene in the LWIR has minimal effect on the detection of the much more expansive methyl chloride, even at significant values of CPL. The major detractor from methyl chloride detection in this scenario appears to be the high frequency jitter coupled with mirror error. The high frequency jitter causes significant portions of the periodic spectral feature points in the interferogram to become overwritten with OPD radiance values from another pixel without the gas causing malformations to the transformed spectrum and reducing the detection score.

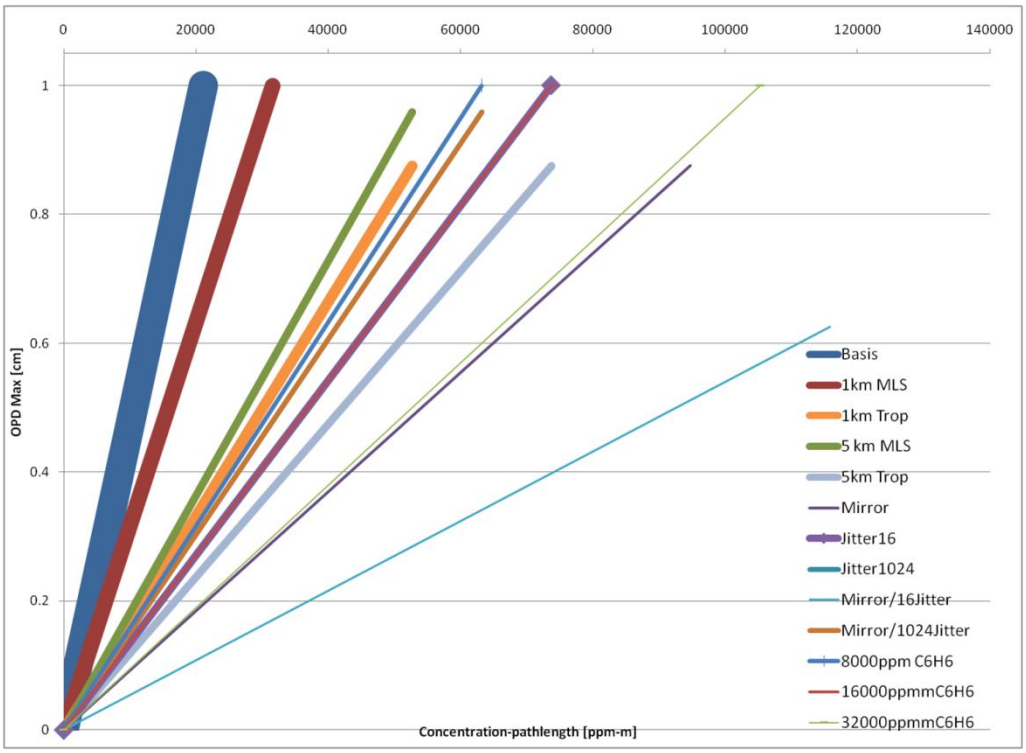

FIGURE 4.14 PEAK DETECTION SCORE VECTOR SUMMARY FOR METHYL CHLORIDE

4.1.4 Spectral Shape Dependence on OPD Sample Points for Single-sided Interferograms

The next study looked at how the number of sample points used in constructing a single-sided interferogram without phase correction would impact the resulting spectral feature shape. This study is somewhat confusing in name in that the instrument lineshape, which is determined by the maximum OPD, is what determines the resolution of the sampled spectral feature. Any additional points added to the sampling should essentially fall on a line that is a linear fit between the critically sampled points. This is true in that the number of sample points do not affect the instrument resolution, but there was minor impact to the shape of the spectral feature due to differences in the sampling interval. The smaller interval with a larger number of sample points over the same distance resulted in generating different OPD radiance samples, which when transformed, resulted in slightly different peak shapes when no phase correction was applied. For some spectra no difference was noted, while others showed unpredictable trends. This study helped determine the impact of increasing the sampling points and to find the point where no further improvement could be made. The plot below shows a comparison in target spectral feature shape with a successively larger number of sample points.

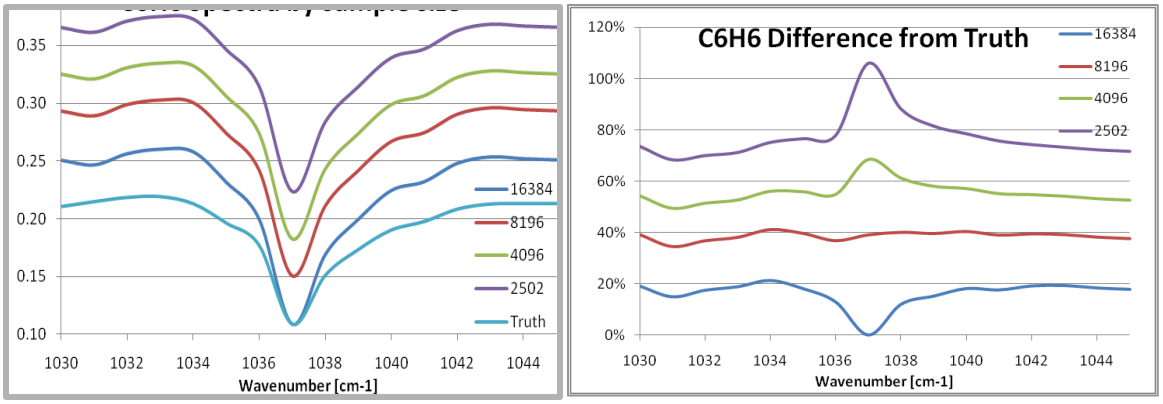

FIGURE 4.15 SPECTRA SAMPLE SIZE COMPARISON FOR BENZENE (LEFT), TRANSFORMED SPECTRA DIFFERENCE FROM TRUTH (RIGHT)

Figure 4.15 shows that as the number of sample points change, the shape of the spectral feature changes, but the FWHM is maintained, as limited by the maximum OPD. This particular series shows that 8,196 sample points results in the closest match to the actual input target spectral feature, with increasing points actually decreasing the match. Figure 4.16 shows that this result was seen in a spectral feature from ammonia with a similar progression.

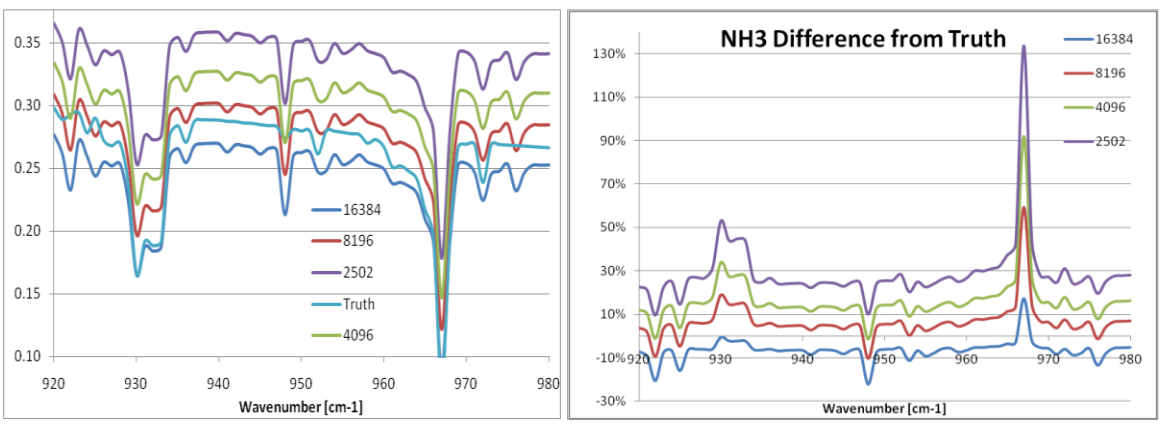

FIGURE 4.16 SPECTRA SAMPLE SIZE COMPARISON FOR AMMONIA (LEFT), TRANSFORMED SPECTRA DIFFERENCE FROM TRUTH (RIGHT)

4.1.5 Interferogrammetric Sensitivity to Identifying Spectral Feature Study

During a series of interferogram studies looking at the sensitivity of various regions in the interferogram to corresponding components of spectral features in the spectral domain, an experiment was run to determine how much information from the interferogram is required to be able to identify a spectral feature as a match for a given species. The idea was to determine if there was a point of diminishing returns where additional OPD pathlength would fail to create a proportionate addition to the spectral information. This study consisted of a performance trade to determine the range of MOPD required to characterize each of the gases in the work. SAM scores for the gas target absorptivity spectra were calculated at varying MOPD values for each gas used in the work. The PNNL absorptivity spectra were sampled at an interval of approximately 0.06 cm^{-1} (0.6

nm at 10µm), which would require an MOPD of 16.67 cm to fully resolve using the instrument model in this work.

The absorptivity spectra were run directly and sampled by the instrument model using the varying MOPD settings. The simulation was run out to a maximum MOPD of 2 cm, which is over 8x under-resolved as compared to the input spectra, and explains why the SAM scores are still quite large at this MOPD. The vertical lines that run through each of the plots show the 0.2 and 1 cm MOPD points, which were the minimum and maximum MOPD values chosen for this work. Figure 4.17 shows the results of the experiment -- note the differences in vertical scale between panels.

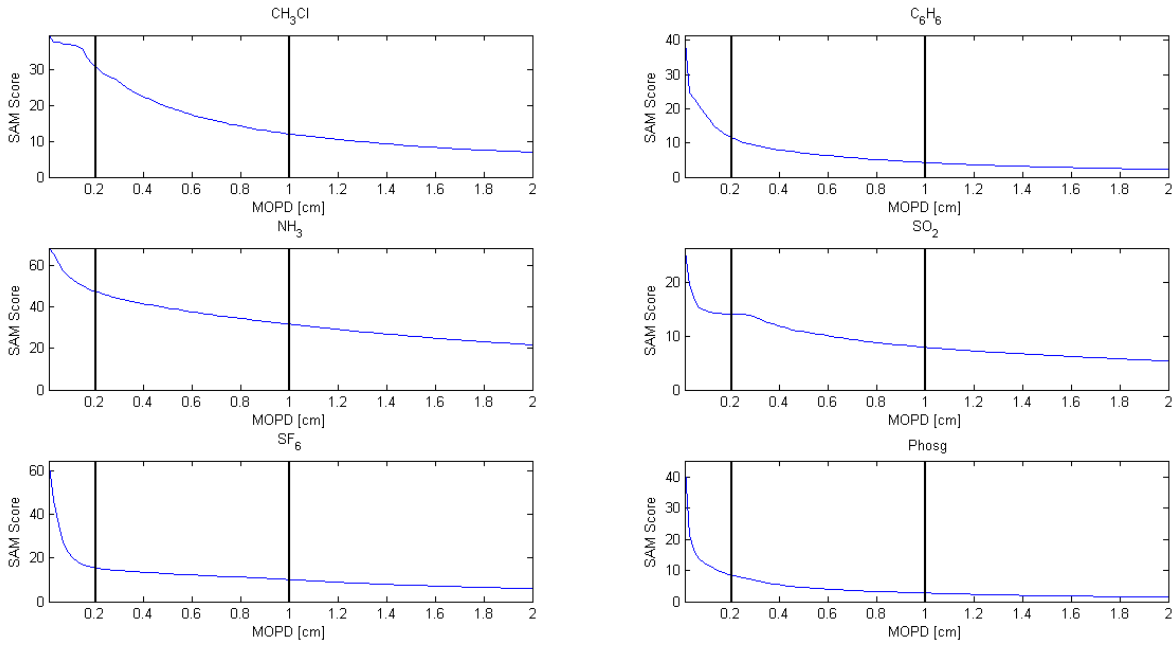

FIGURE 4.17 SPECTRAL ANGLE VS. MAXIMUM OPD BY GAS SPECIES

The plots show that phosgene, sulfur hexafluoride, and to some extent benzene, contain most of their spectral information in the low frequency portions of the interferogram. This concept makes sense when reviewing the interferogram series shown in Figure 3.28 and Figure 3.29. Conversely, ammonia, methyl chloride, and to some extent sulfur dioxide, show more of a linear progression with spectral accuracy at increasing MOPD. This is due to the fact that the spectral feature information in the interferogram is significant at nearly all frequencies across the range through 1 cm MOPD. Figure 4.18 shows how much more spectral information is contained in ammonia with increasing spectral resolution due to the abundance of rotational lines across the LWIR band.

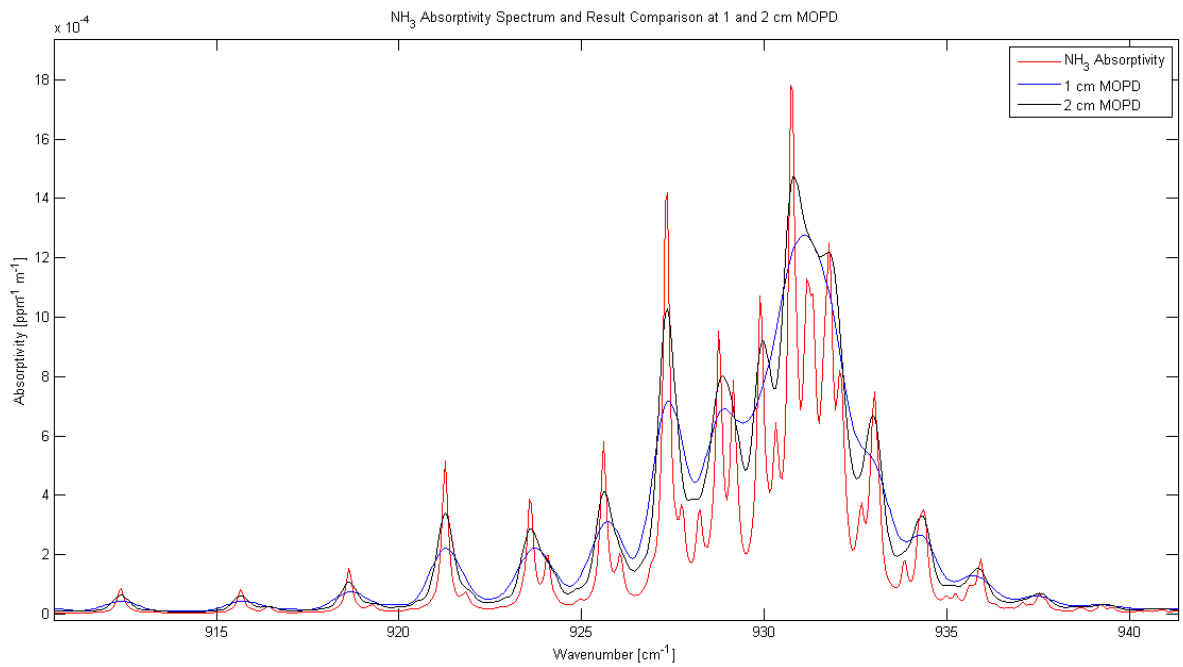

FIGURE 4.18 NH_3 COMPARISON AT 1 (BLUE TRACE) AND 2 (BLACK TRACE) CM MOPD TO FULL RESOLUTION PNNL DATA (RED TRACE)

This study was a quick, quantitative way to determine how susceptible a gas would be to interferogram corruption in given locations as well as how much resolution is required to resolve the spectral feature to a level where identification is possible with some confidence.

4.1.6 Interferogram Filtering and Gas Identification Study

Early on, the potential to accomplish gas detection and identification through the use of only a narrow region in the interferogram was explored. As previously mentioned, this would allow spectral feature detection regardless of background, total radiance, atmospheric effects, and noise. The target region in the interferogram is the one that responds to changes in the shape of the spectral feature of interest. The following series of plots illustrates the progression of this study.

Figure 4.19 shows the spectrum at 1 cm MOPD where the presence of ammonia and SF_6 in approximately equal spectral strength are shown in emission between 930-970 cm^{-1}.

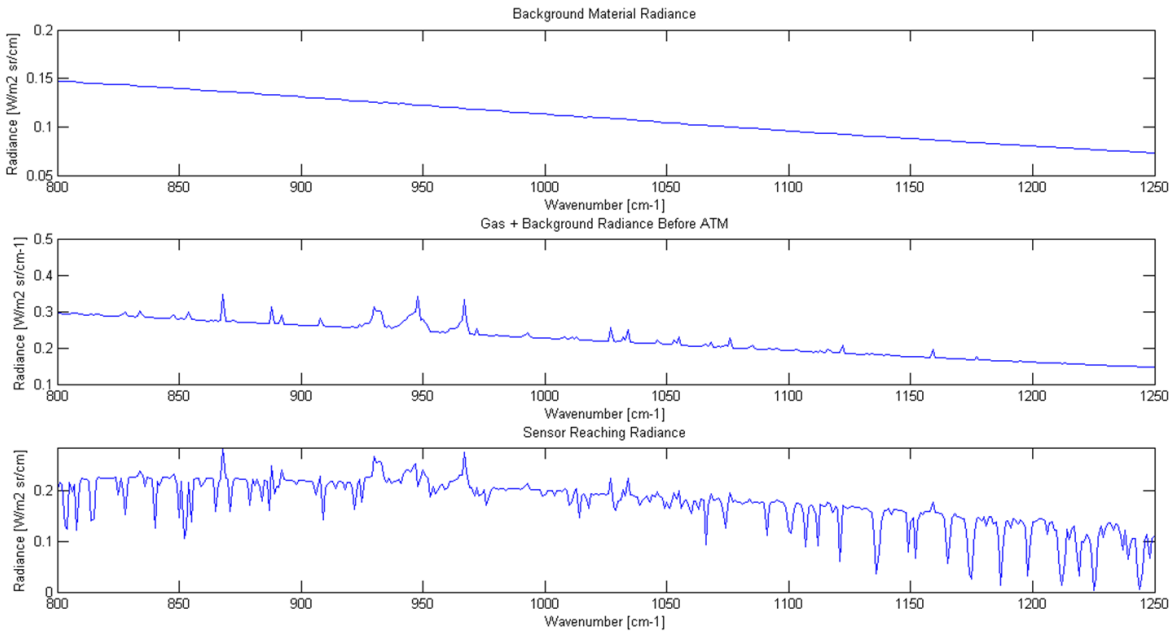

FIGURE 4.19 SPECTRAL PLOTS OF BACKGROUD MATERIAL RADAINCE (TOP), PLUME-LEAVING RADIANCE (MIDDLE), SRR WITH SF_6 AND NH_3 IN EMISSION (BOTTOM)

Figure 4.20 consists of a four-panel series that shows the region in the interferogram most sensitive to the target gases in emission. The second panel shows the notch filtered interferogram, while the third shows an overlay of interferograms that contain and do not contain the target gases. The fourth panel is a difference plot between the interferograms in the third panel.

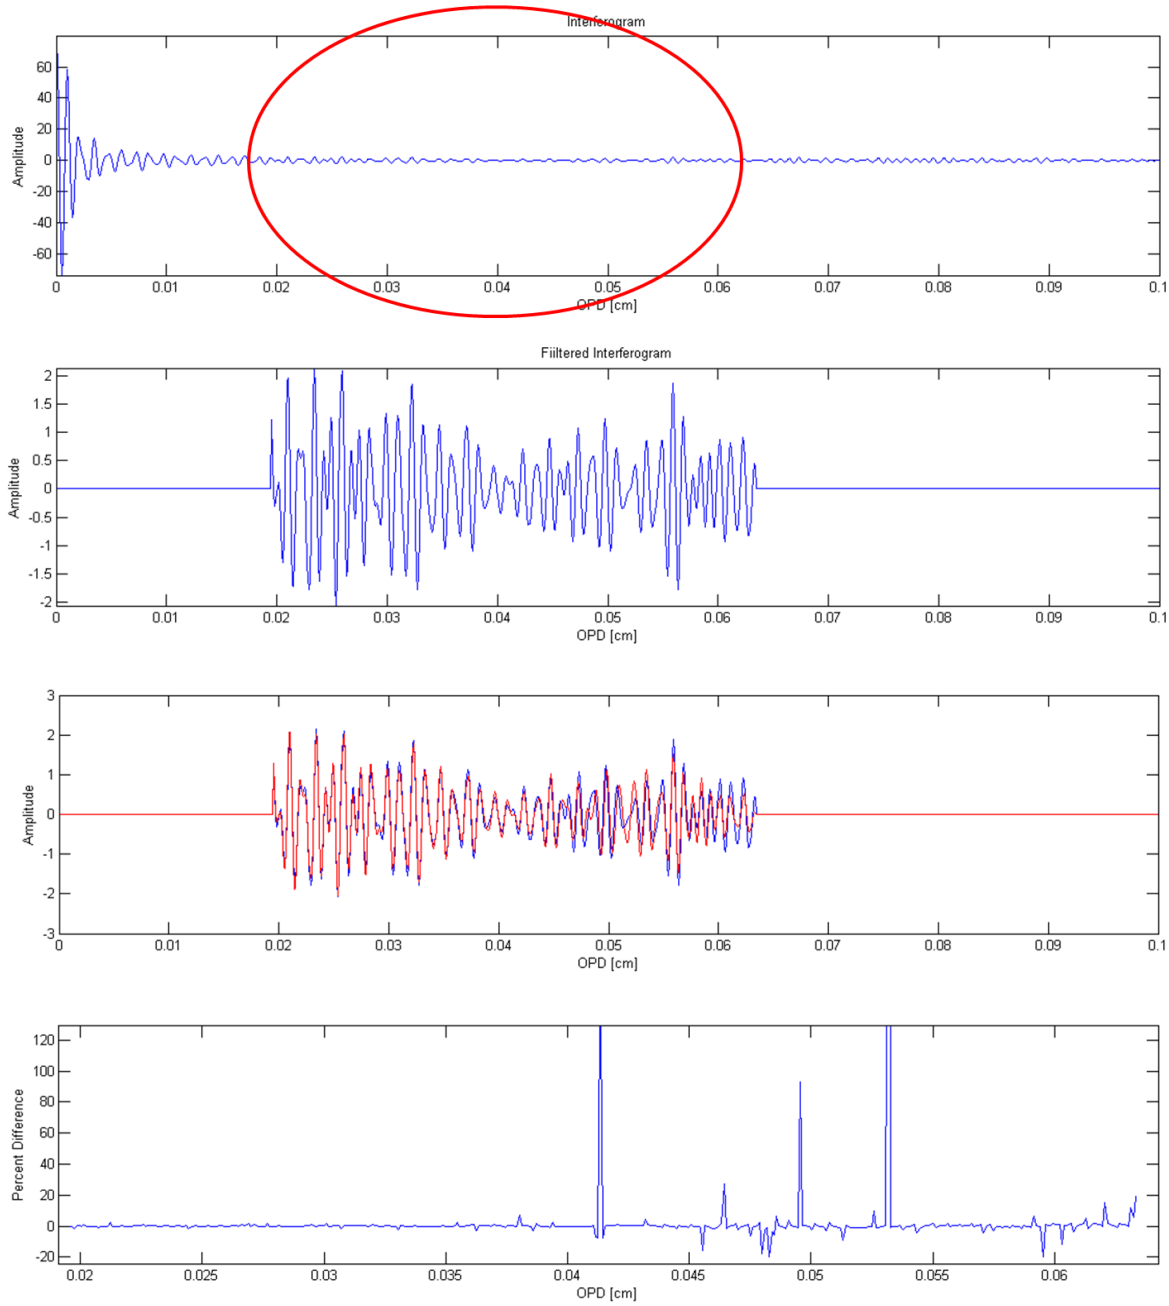

FIGURE 4.20 INTERFEROGRAM AT 0.1 CM MAXIMUM OPD WITH FILTERED SENSITIVE REGION

The fourth panel shows there are clearly three points in the interferogram that are most sensitive to the presence of the gases over this region. It further shows that the interferogram filter could be compressed further from 0.04 cm to 0.6 cm OPD. Figure 4.21 is a two-panel series that shows the transform of the notch filtered interferogram. The trace in blue is the transform of the interferogram with the target gases, while the red trace does not contain targets. The second panel shows the transform of the smaller, 0.04 to 0.06 cm OPD region suggested from the difference plot above, which further exploits the peak positions.

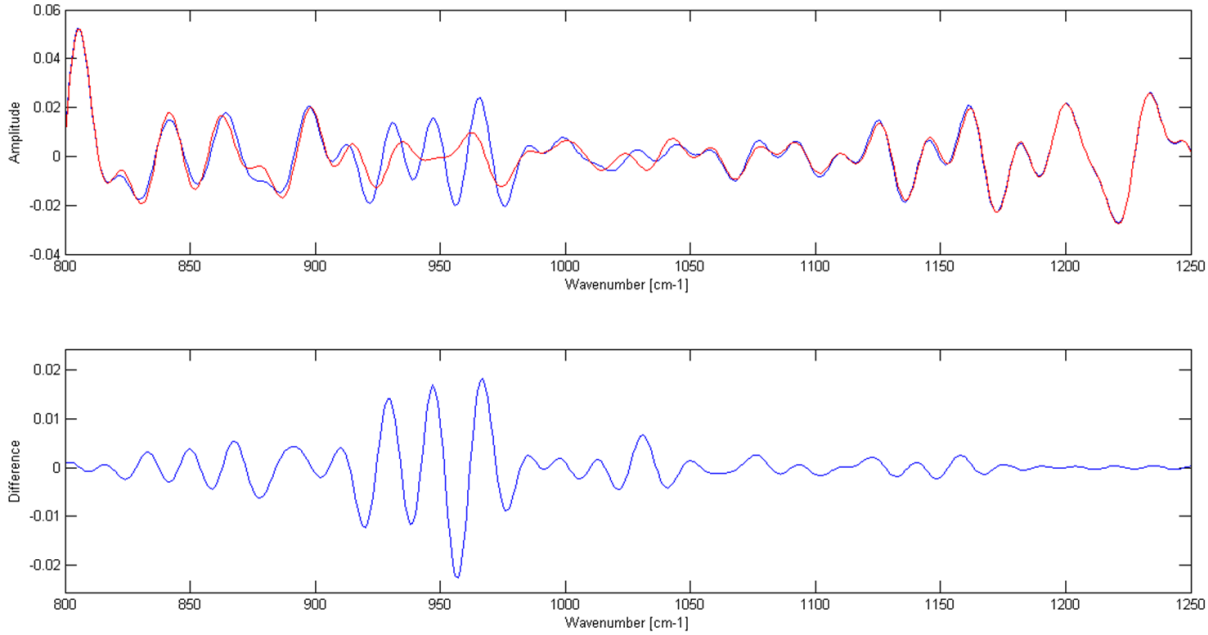

FIGURE 4.21 SPECTRAL TRANSFORMS OF FILTERED INTERFEROGRAM, BLUE TRACE WITH TARGETS, RED TRACE WITHOUT (TOP), DIFFERENCE BETWEEN TRANSFORMS WITH AND WITHOUT GAS (BOTTOM)

The series clearly shows the possibilities in isolating and monitoring the interferogram for portions of the gas spectral feature that are considered unique and identifiable. This approach was ultimately abandoned as described previously when it was determined that sources of interferogram error could easily corrupt such a narrow portion of the interferogram region. Spectral band monitoring proved to be a more robust method for spectral feature identification.

4.1.7 Identifying Gases by the Fine Spectral Detail Portion of the Interferogram

In a final interferogram study, the ability of a gas to be detected by its fine spectral detail alone was determined when a unique result from SF_6 was found. The earlier interferogram and SAM score study in Figure 4.17 indicates that most of the spectral feature information in SF_6 is front loaded in the interferogram. An experiment, known as interferogram 'blocking' was being explored as a precursor to developing a jitter error module. The concept was to determine the extent of interferogram replacement effects by combining interferogrammetric information from multiple scenes. In this experiment, the first half of the interferogram was built from a 290K sand background with no gas target present, then switched to 320K SF_6 gas at 100 ppm-m over a 300K asphalt surface at maximum resolution. Figure 4.22 shows results for both apodized and unapodized versions of the interferogram.

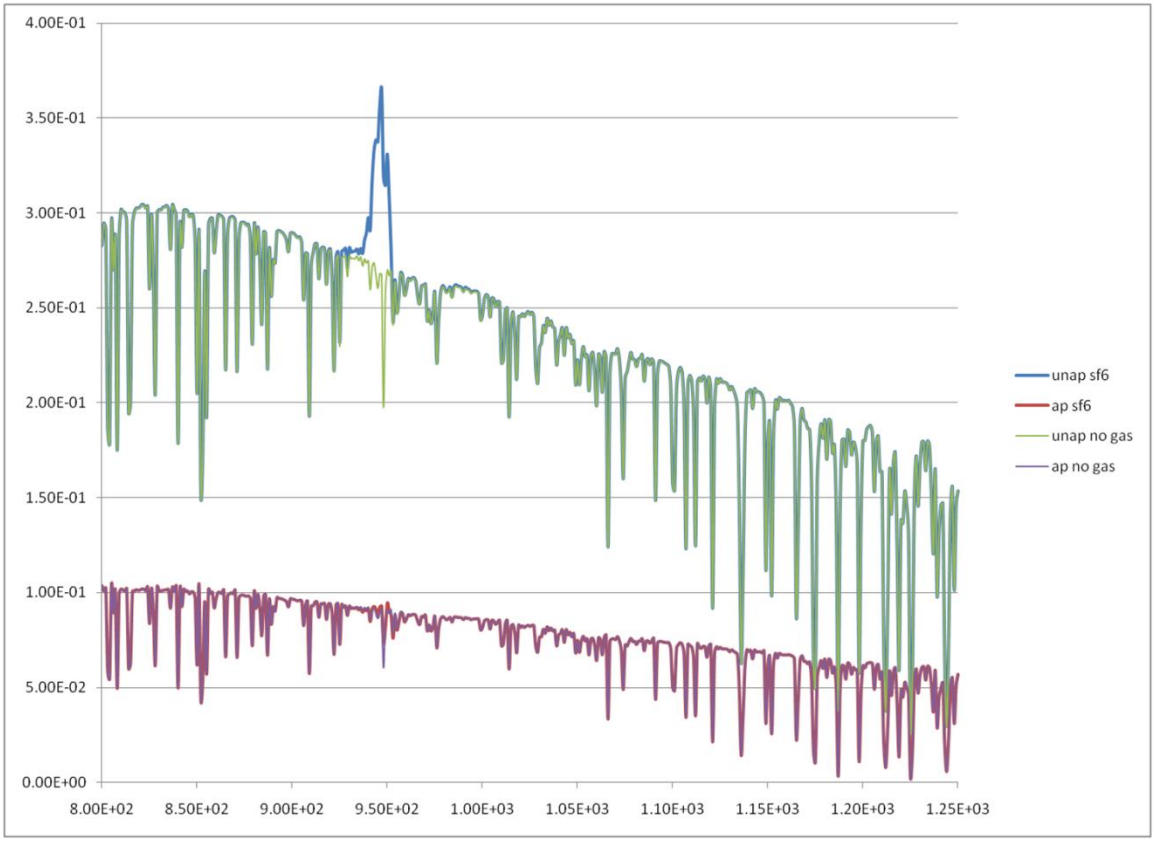

FIGURE 4.22 APODIZED (BOTTOM) AND UNAPODIZED (TOP) SPECTRA OF SF_6 (BLUE AND RED TRACES) AND CONTROL (GREEN AND PURPLE TRACES)

The plot shows two unscaled versions where the top series is a comparison of the unapodized series with and without the gas, and the bottom series are both apodized. The bottom series barely shows any indication of SF_6 gas in the spectrum and that is because all of the bias information for the spectral feature is lost from the first half of the interferogram giving no form for the fine spectral detail to modify. The unapodized result shows substantial evidence of SF_6. The full version of SF_6 bares only slight resemblance to this spectral artifact as this peak is grossly compressed with almost all bias information lost.

What this experiment demonstrated is that in some cases, species can be identified by their fine spectral detail alone, as long as the gas is present in sufficient spectral strength. Here, 100 ppm-m at a +20K thermal contrast is a substantial strength for SF_6 and the peak from a full interferogram would have been much larger. This plot suggests that one possible approach to finding a gas in the presence of significant low frequency, long duration jitter would be to run a detection metric against both the apodized and unapodized interferogram either real time, or off-line, if a provision is made to retain both interferograms.

4.1.8 Monte Carlo Studies for Mapping the Effect of Mirror and Jitter Error

The Monte Carlo studies gave an important insight into understanding the impact of mirror and jitter error on the transform of the spectral feature. The following plots showed why this study was needed.

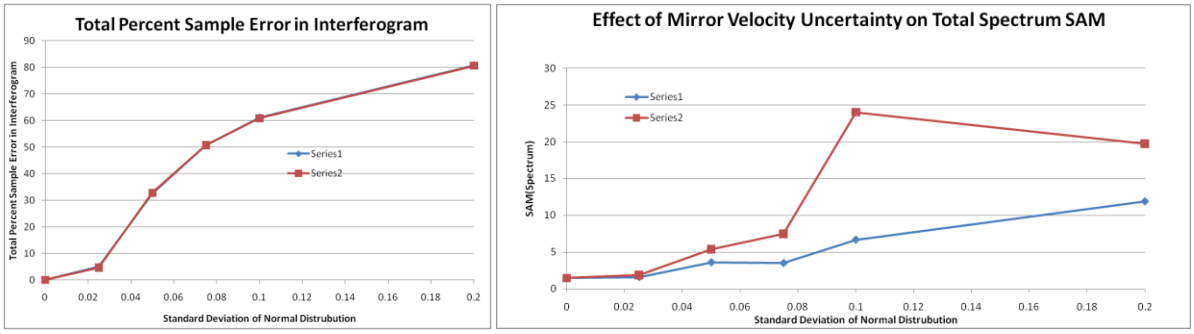

FIGURE 4.23 TOTAL SAMPLE ERRORS (LEFT) AND EFFECT OF MIRROR UNCERTAINTY ON TOTAL SPECTRUM SAM SCORE (RIGHT)

Figure 4.23 shows that while nearly identical numbers of interferogram samples were subjected to error in both trials, the resulting transformed spectra could yield a detection score over 5x worse as compared to another. Two points became evident when reviewing the data. First, the overall impact to the transformed spectrum was dependent on where in the interferogram the error took place, and second, the random nature of this event has highly variable effects. The target was 268K benzene at 18,000 ppm-m over 300K asphalt with a 1 km MLS atmosphere. Figure 4.24 was the first result from a 200 trial Monte Carlo mapping of mirror error uncertainty.

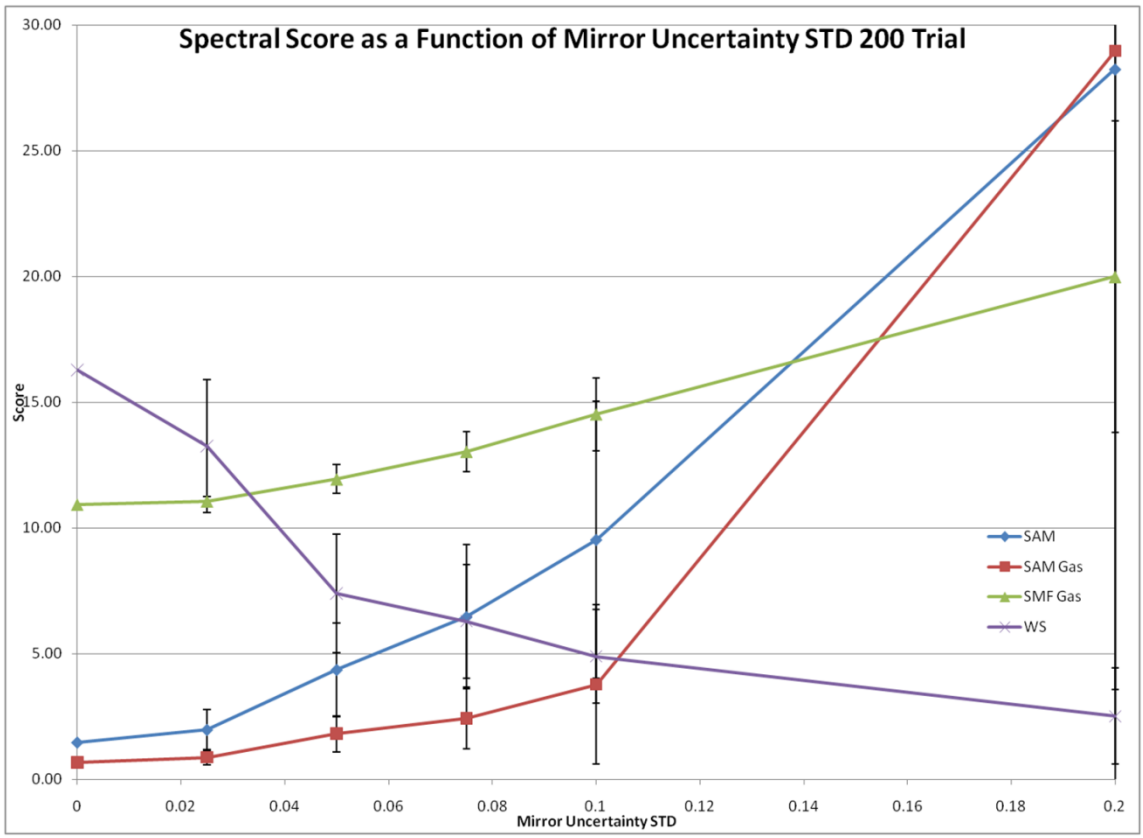

FIGURE 4.24 200 TRIAL MONTE CARLO DETECTION SCORE RESULT FOR MIRROR ERROR UNCERTAINTY

The chart shows the mean values for the detection scores with bounds that show the minimum and maximum scores that resulted from each error series. While additional trials resulted in slight differences to the score boundaries, the mean score values were invariant. The spectral ratio showed the most sensitivity and the clearest trend of the effect of the errors among the detection metrics.

The mirror uncertainty standard deviation is the actual standard deviation applied to the normal distribution responsible for generating the mirror errors. The resulting values from the distribution are scaled and rounded to integers and applied to the OPD sample points. No real OPD displacement can be applied to the standard deviation value as it represents a range of sample error values that result from the expanded distribution. It is acknowledged that using the standard deviation to describe this type of error magnitude is somewhat confusing, but so are the alternatives. The largest magnitude error could be listed in terms of displacement, but these would occur infrequently in the distribution, creating misunderstanding as to how much of an impact the error value really had. In short, the standard deviation will be used as a measure of the magnitude of the error through the series.

Figure 4.25 shows a series of results illustrating the effects of mirror errors, jitter errors, then the combination of both.

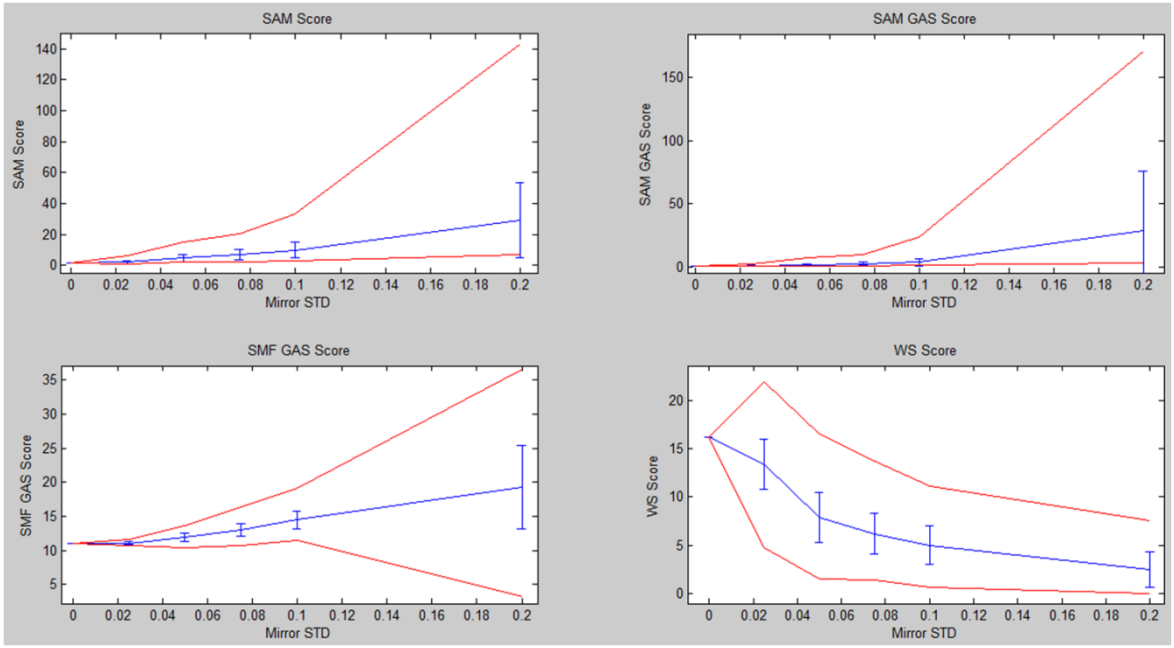

a) Mirror error results

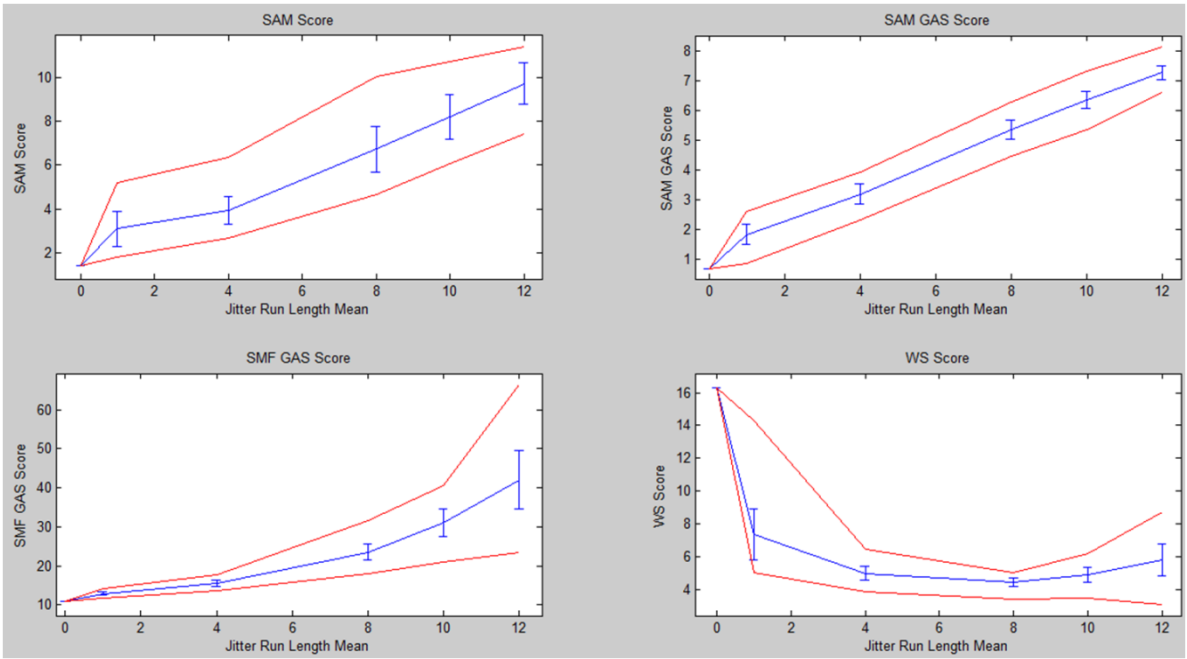

b) Jitter error results

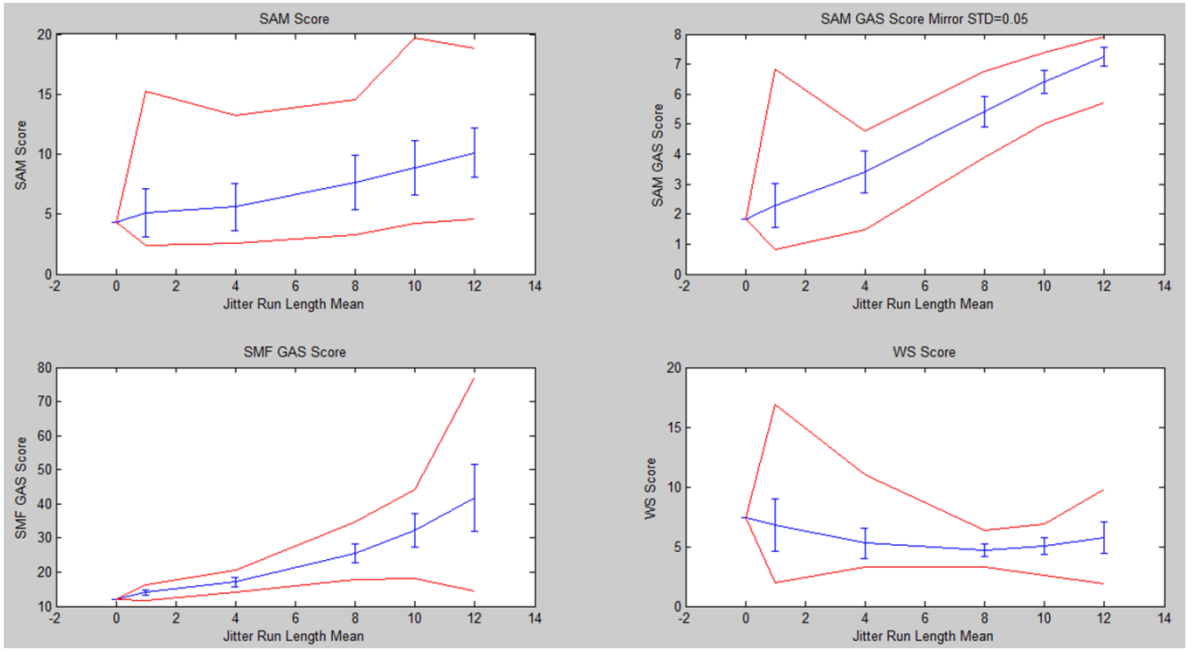

c) Combined jitter and mirror error results

FIGURE 4.25 MONTE CARLO SERIES: MIRROR ERROR (A), JITTER ERROR (B), COMBINED ERRORS (C)

In the second panel, *b*, the jitter run length mean is the mean value assigned to the number of OPD samples that view the off-target scene without the gas. The jitter run length is essentially the duration of the jitter error measured in OPD sample points. The third panel shows one example from the series of combined jitter and mirror error trials. This trial was for a mirror error standard deviation of 0.05. Figure 4.26 shows the next series of plots containing the results by mean detection score for each of the mirror error magnitudes at each jitter run length in the combined error trials.

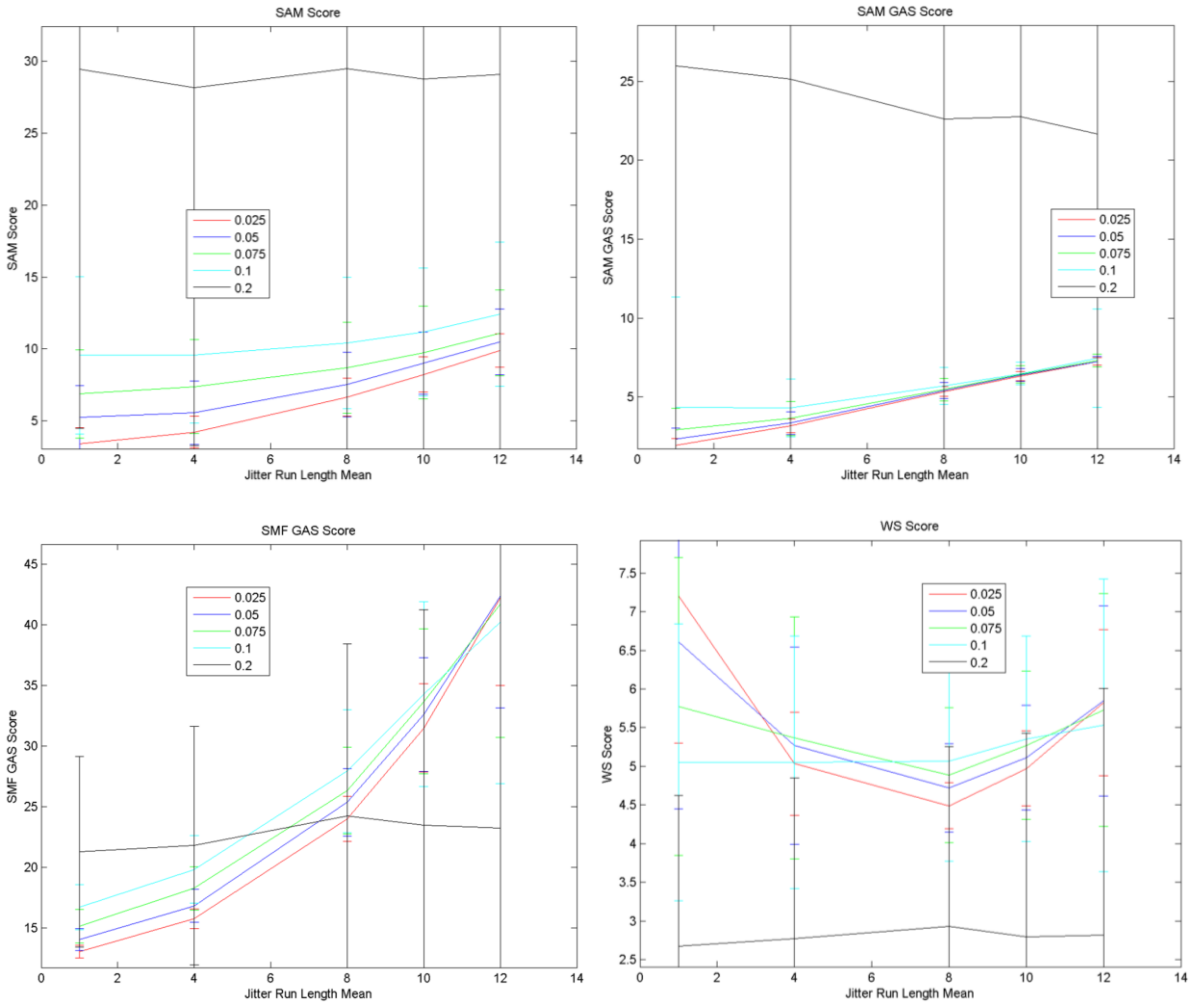

FIGURE 4.26 DETECTION SCORE SERIES BY JITTER RUN LENGTH AS A FUNCTION OF MIRROR UNCERTAINTY FOR BENZENE

The plots seem to indicate that after an initial influence from the mirror error, the jitter error becomes dominant and the scores tend to converge with increasing jitter run length. The increasing spectral ratio score, after the initial decrease, has been explained previously and is due to the increased spectral energy injected into the target band from bias errors from low frequency corruption in the interferogram. The 0.2 standard deviation level of mirror error introduces a significant amount of critical sampling errors and therefore no recognizable spectrum is recovered. This level of mirror error affects 80% of the OPD samples with 12% resulting in an OPD displacement of 9 µm and 25% out to 3.66 µm, where 4 µm is the critical sampling interval. The resulting spectra are essentially noise and average to the same detection score across the target vector band at each jitter error value. This is illustrated on essentially every score chart where the scores from the 0.2 mirror series are nearly flat across the jitter error series. Figure 4.27 shows a benzene spectrum with a 0.2 standard deviation mirror error compared to the SRR.

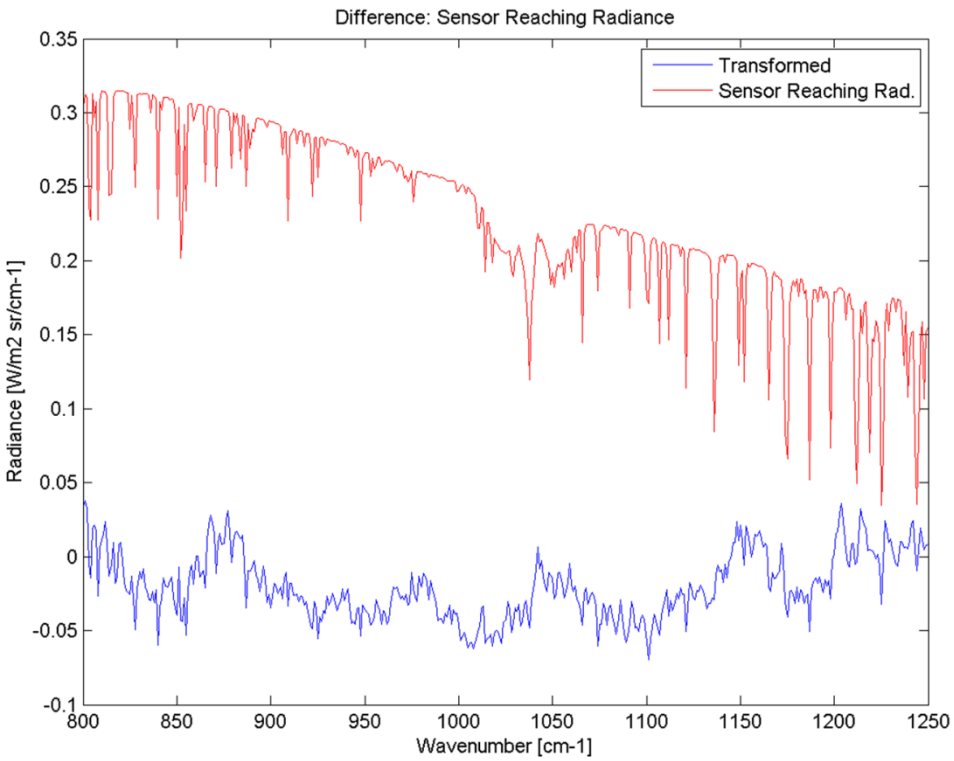

FIGURE 4.27 TRANSFORMED (BLUE TRACE) AND SRR COMPARISON (RED TRACE) FOR BENZENE AT 0.2 STANDARD DEVIATION MIRROR ERROR

The plot shows a minor resemblance to benzene on the high wavenumber side of the spectral feature, but not enough to elicit a reasonable detection score.

The results for ammonia were somewhat different from benzene and are worth noting here. Figure 4.28 shows the same series for ammonia as was just shown for benzene.

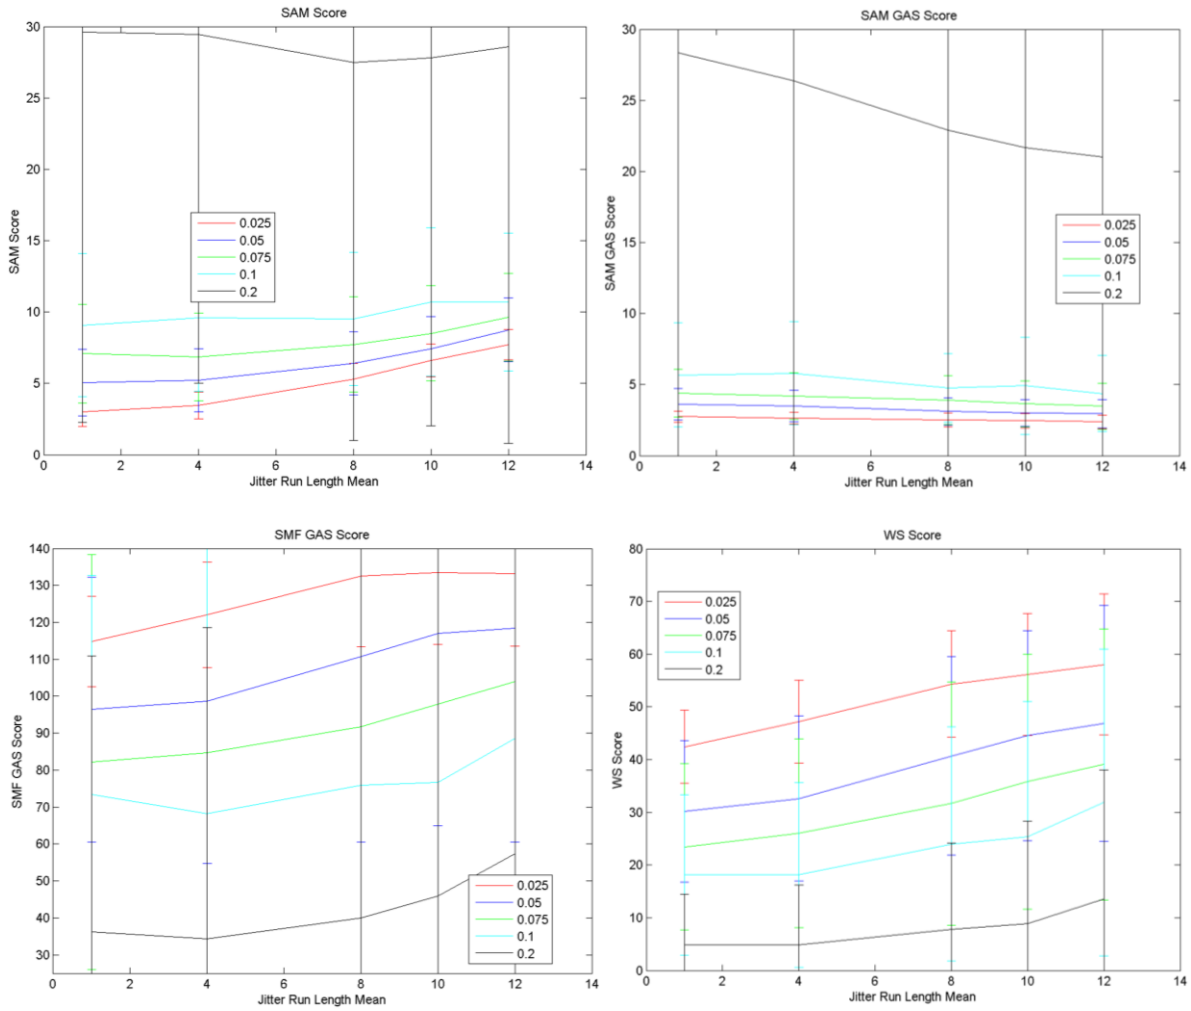

FIGURE 4.28 DETECTION SCORE SERIES BY JITTER RUN LENGTH AS A FUNCTION OF MIRROR UNCERTAINTY FOR AMMONIA

The score results for ammonia appear to show more resilience to both mirror and jitter error as compared to benzene. This effect can be understood from the interferogram of ammonia. As ammonia spectral features populate the entire spectrum with peaks of various shapes and magnitudes, evidence of these spectral features appear nearly continuously in the interferogram. Thus, even at the 0.2 level of mirror error, some of the spectrum is still recovered and detectable. Figure 4.29 shows ammonia at a 0.2 standard deviation level of mirror error as compared to the SRR and the second plot shows the target vector as compared to the transformed radiance over the same band.

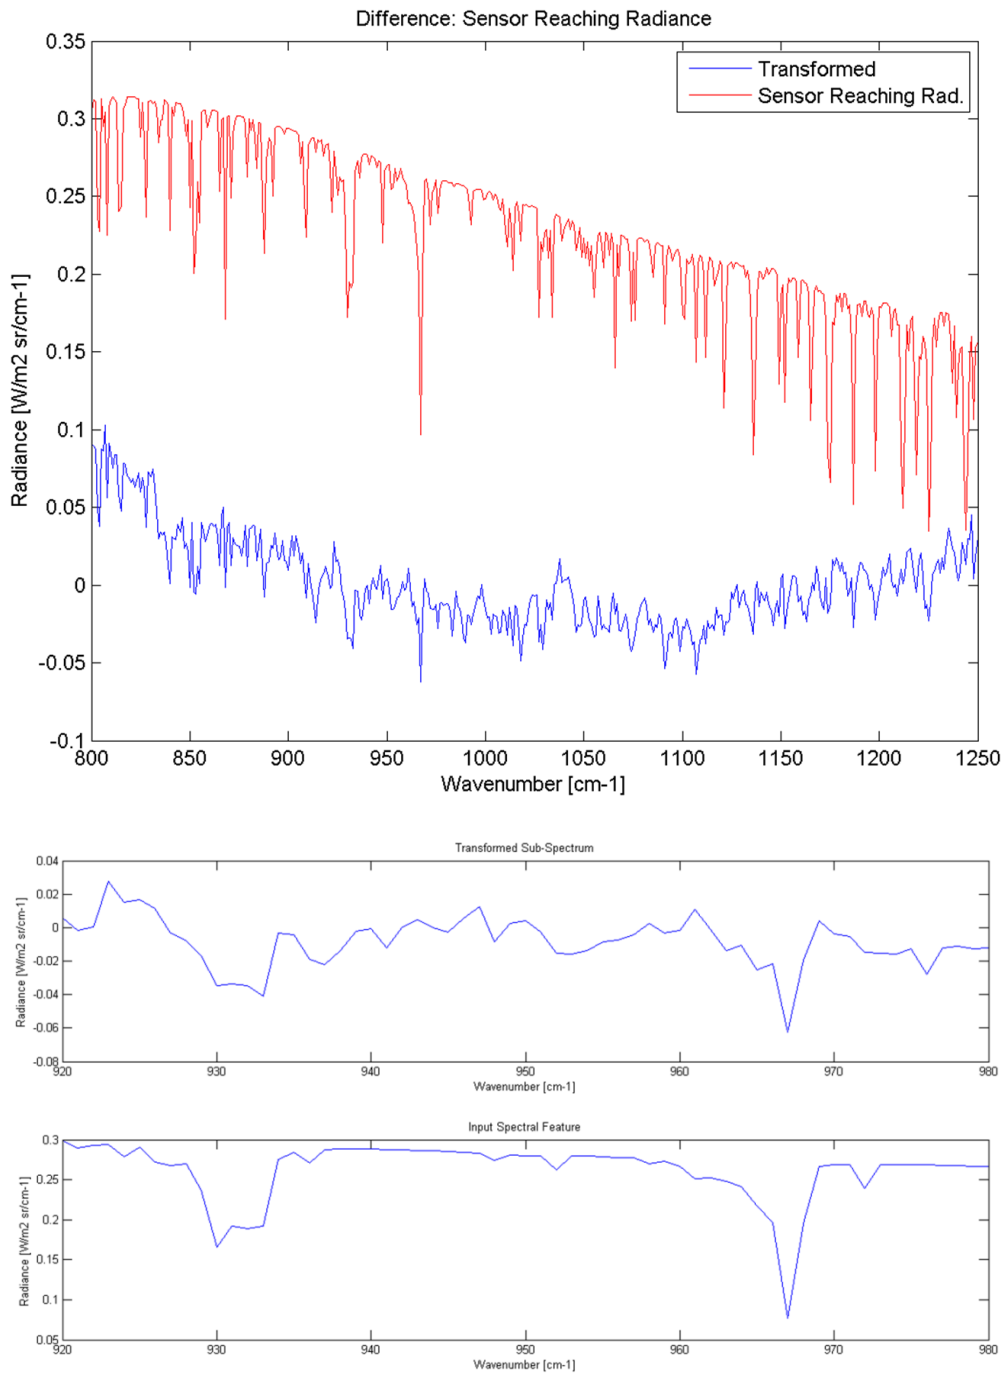

FIGURE 4.29 TRANSFORMED AND SRR COMPARISON FOR NH_3 AT 0.2 STD. DEV. MIRROR ERROR (TOP) WITH TARGET VECTOR COMPARISON (2- PANEL, BOTTOM)

As can be seen in the plot, particularly in the lower panel showing the transformed spectrum and target vector spectral bands, respectively, the main spectral features from ammonia are still largely visible after significant disruption of the interferogram.

4.2 Scene-based Study Selected Results

The scene-based study results differ from the previous section in that rather than run against a range of varying MOPD and CPL or some other parameter trade, a scene was constructed to test the ranges of parameters that were determined from the performance trades. These scenes consisted of first, a simple Gaussian plume model with appropriately varying thermal and CPL properties over a single background material to characterize system performance over the extent of a gas plume. Next, a material emissivity map with increased thermal and material emissivity variability with embedded gas plumes was used to characterize detection performance in a more realistic setting under a wider range of conditions. Finally, various DIRSIG scenes were constructed and 'flown' against to give the most realistic characterization of system performance.

4.2.1 Plume Model Detection Study

The Gaussian model plume was subjected to the scanning airborne FTS model to determine how well the system would do at collecting data from the plume at varying scan rates and MOPD in terms of binning and detection score. The scan rates for this study included 5, 10, 20, 40, and 60 Hz. Figure 4.30 shows a few selected detection planes from this study that show the improvement in spatial plume characterization at 1 cm MOPD for the 20, 40, and 60 Hz scan rates. While the plume was characterized more completely in the spatial extent, at 1 cm MOPD, the increased scan rates did not provide a proportionate increase in detection score. Thus, the maximum scan rate for future scene studies was limited to 40 Hz.

The target vector was again, 4,800 ppm-m benzene at 268K. The target vector was chosen as a 'best guess' for the leak source if the actual flow rate was unknown. The platform velocity was 50 m/s with 0.2 m GSD at nadir, from a flight altitude of 1 km. The atmosphere was a 1km MLS with 270K benzene gas emanating from the leak site over a 300K background of sand. The gas CPL in the plume ranged from 10,000 ppm-m at the source down to zero. A critically sampled, double-sided interferogram was used for this collect.

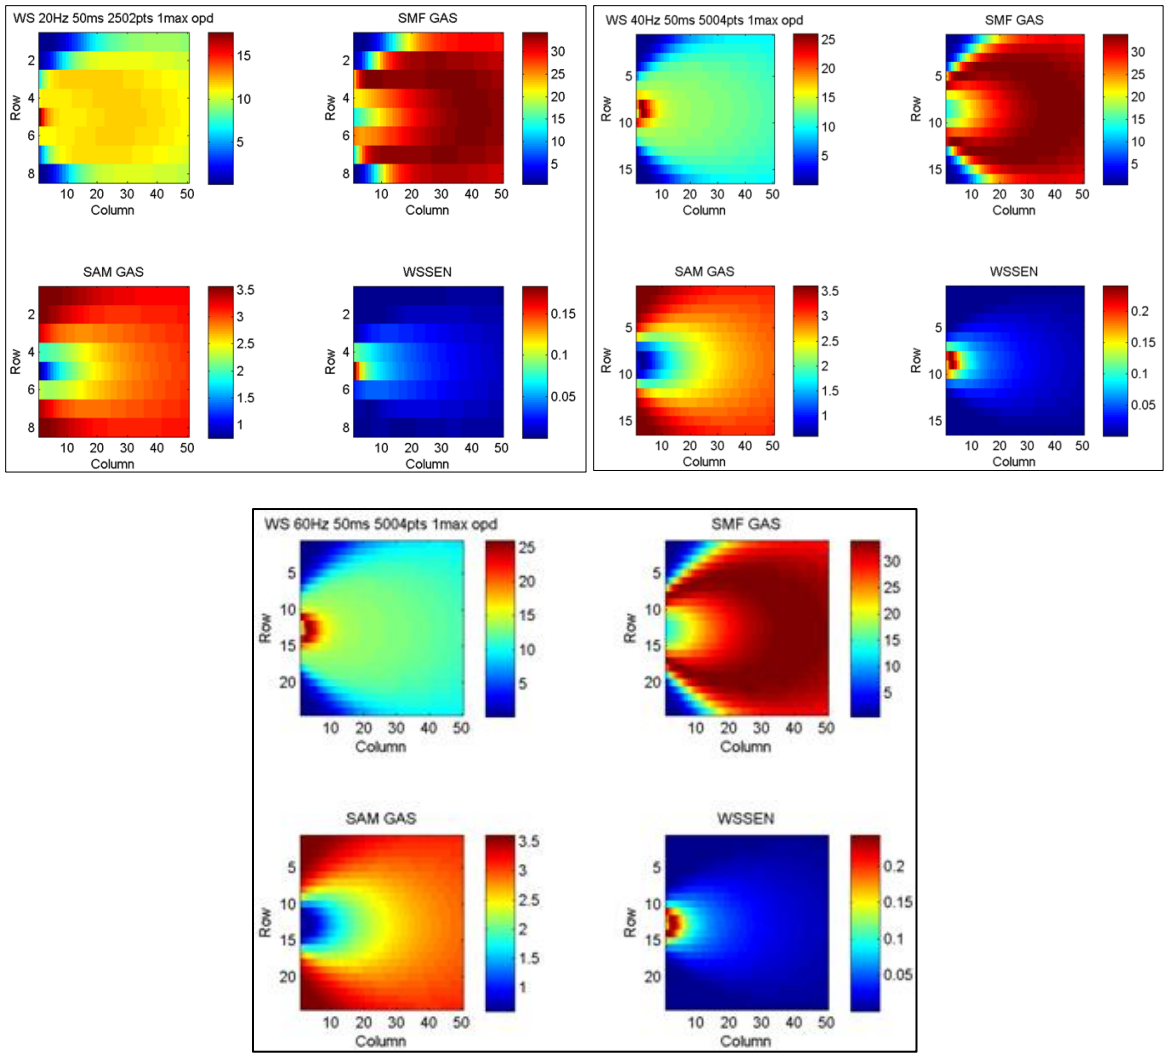

FIGURE 4.30 DETECTION PLANES OF 20/40/60 HZ SCAN RATE SERIES AT 1 CM MOPD FOR GAUSSIAN PLUME MODEL OF BENZENE GAS

The detection planes show the effect of binning at low scan rates. The detection metrics correctly identify the target vector in the plume spatially, as a function of both CPL and thermal contrast. The WSSEN metric is the most selective in the fit for the closest match to the target vector.

A 3D surface plot was used to summarize the results of this study. Figure 4.31 shows the MOPD and scan rate trade space for the aforementioned scan rates and MOPD values of 0.2, 0.4, 0.6, 0.8, and 1 cm. The peak WSSEN score was taken from each of the detection planes for each of the runs in making this chart. The chart shows that in the absence of FA and FTS errors, increasing scan rate and increasing MOPD provide better detection scores as one would hope. Note the sharp reduction in score improvement over the 40-60Hz scan rate range, as well as the 0.8 to 1 cm MOPD range. This result gave indications of the required parameter ranges for future studies.

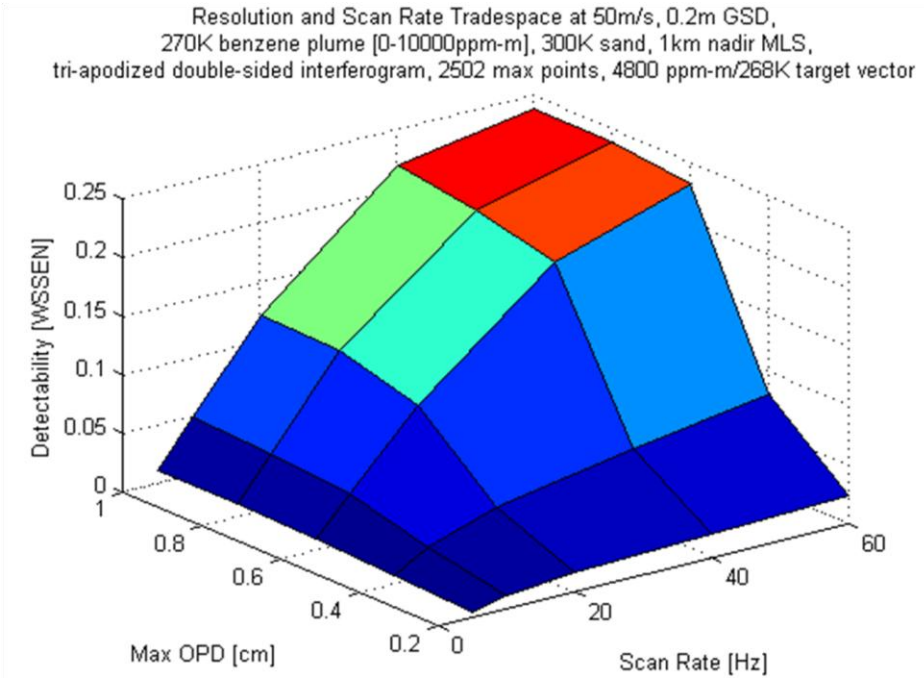

FIGURE 4.31 MOPD VS. SCAN RATE TRADESCPACE FOR GAUSSIAN PLUME MODEL DETECTION SERIES

4.2.2 Material Emissivity Mapped Scene Study

To test the fidelity of gas plume detection in a more realistic scenario, a scene was produced consisting of a material map of emissivities of varying temperatures. Gaussian modeled gas plumes were injected into the scene as overlays on the material map and allowed for characterization of the plume over a wide range of thermal contrast conditions. This scene was discussed in detail in section 3.3.2.

Figure 4.32 shows a sample series of detection planes from a 60 Hz collect at 0.2, 0.6 and 1 cm MOPD. Three plumes, each with a different gas, were injected into this scene. The target vector was for benzene and the benzene plume is at the far right of the scene, at the midpoint, flowing north. The plumes were each given the same temperature distribution profile, equivalent spectral strengths in terms of CPL, and were approximately the same spatial extent in width. This series shows that the system correctly identified benzene as distinct and separate from the other gases across all detection metrics. Further, the 'outlier' nature of the other gases can be seen in the widely desperate detection scores for the benzene and other plumes. Note the ability to 'see' differences in the material classes under the plume, or rather through the plume, due to the change in detection score due to the change in thermal contrast in the form of differences in emissivity.

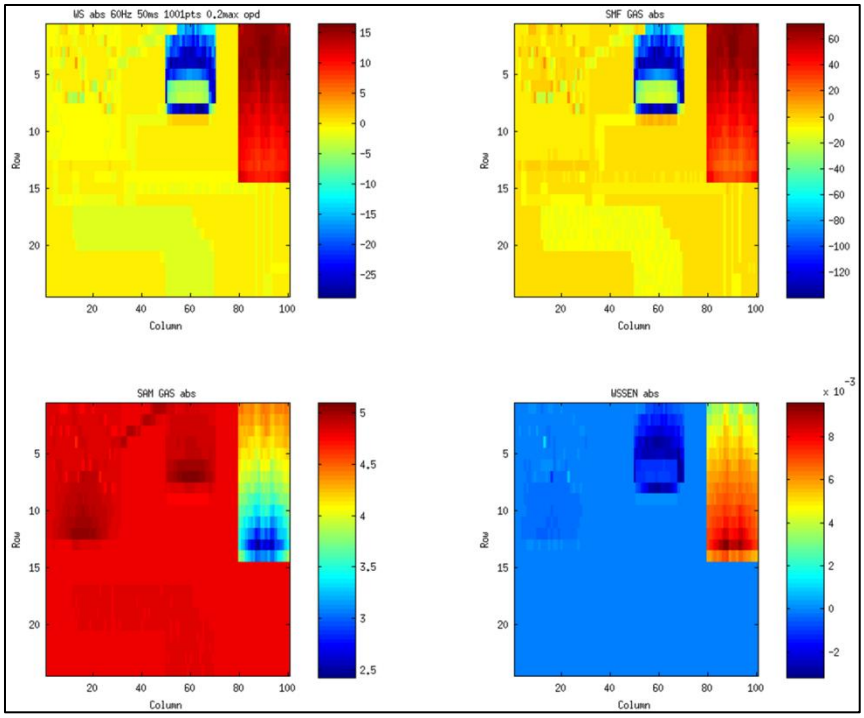

a) 0.2 cm MOPD

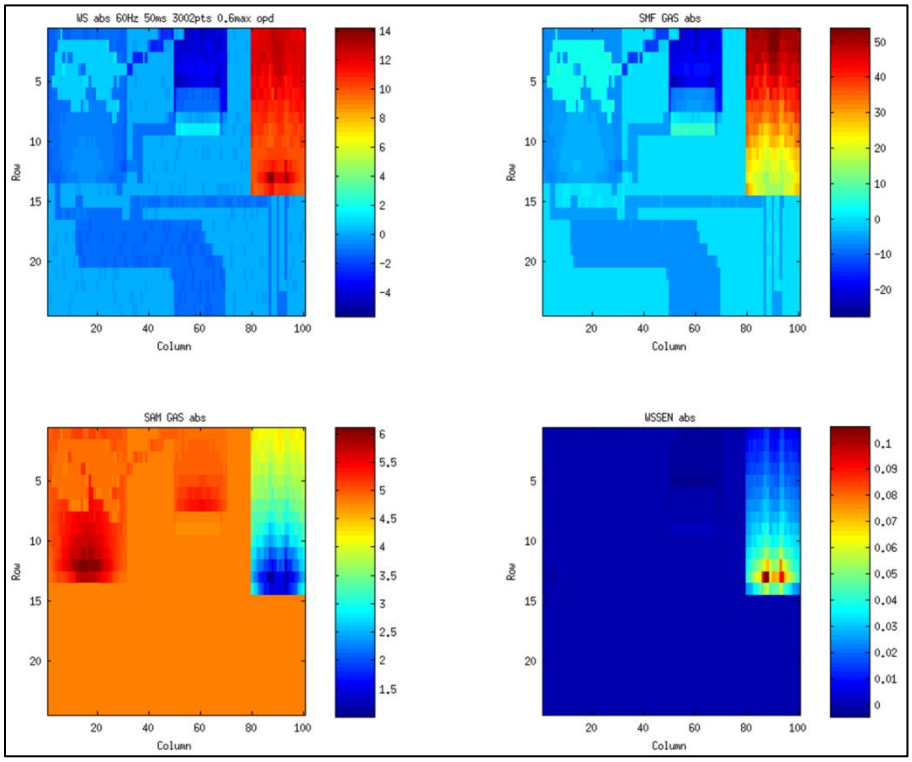

b) 0.6 cm MOPD

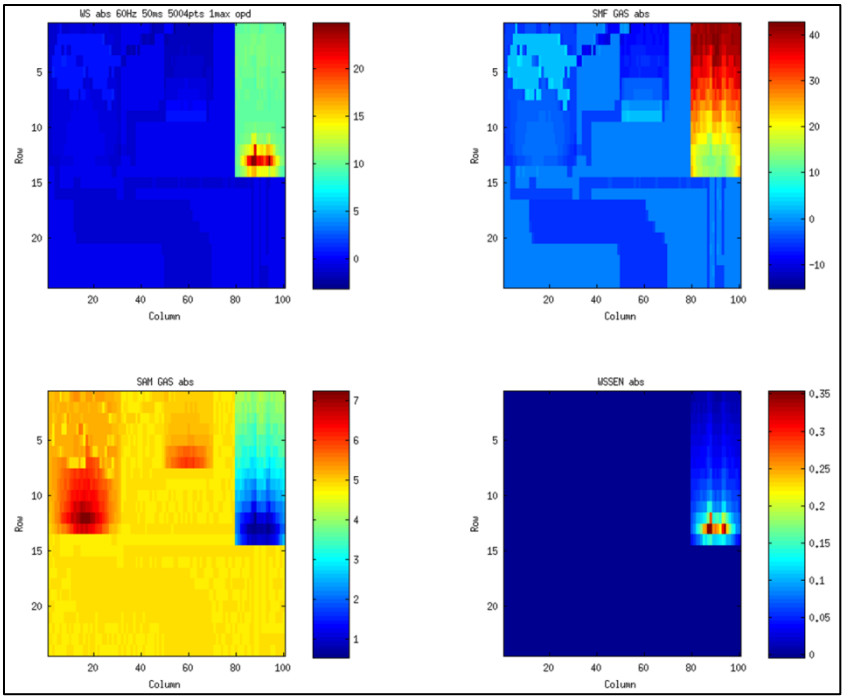

c) 1 cm MOPD

FIGURE 4.32 MATERIALS MAP SCANNING COLLECT WITH BENZENE, AMMONIA, AND METHYL CHLORIDE PLUME, SERIES BY MOPD

The next series shows the ability of the detection metrics to discriminate two different gases when two plumes are situated next to each other. In this trial, a plume consisting of equal spectral strength methyl chloride, with the same temperature profile as benzene, was set next to the original benzene plume so that the left half of the benzene plume and right half of the methyl chloride plume were overlapping in the vertical extent. The other plumes from the previous series remained in the scene.

Figure 4.33 shows the detection plane series for this collect. The scan rate was 60 Hz and the MOPD was set to 1 cm with a platform speed of 50 m/s. It can be seen that the metric correctly rejected the presence of methyl chloride by itself, whereas in the overlapping regions, the overall detection score for benzene was reduced, giving an asymmetric look to the detection score distribution in the plane. The WSSEN score metric seemed to do the best at being able to characterize the full extent of the width of the benzene plume despite the presence of the confuser gas.

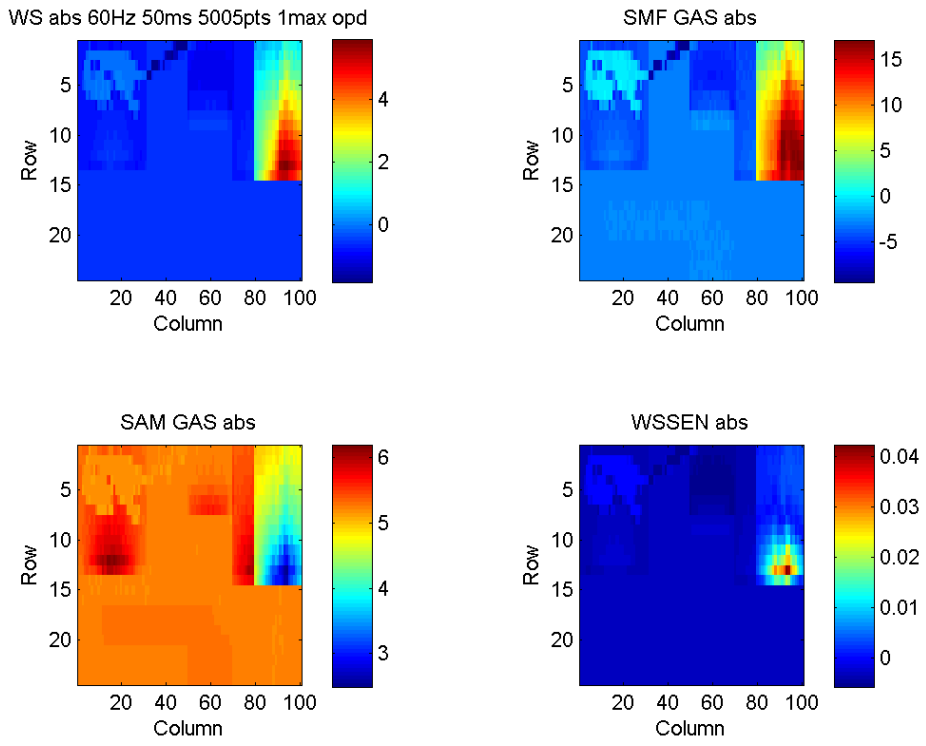

FIGURE 4.33 DETECTION PLANE SERIES FOR OVERLAPPING BENZENE AND METHYL CHLORIDE GAS PLUMES

The next series shows the same trial run under the presence of jitter. Figure 4.34 shows that the jitter had very little effect on the ability to successfully identify the benzene plume under these collection conditions.

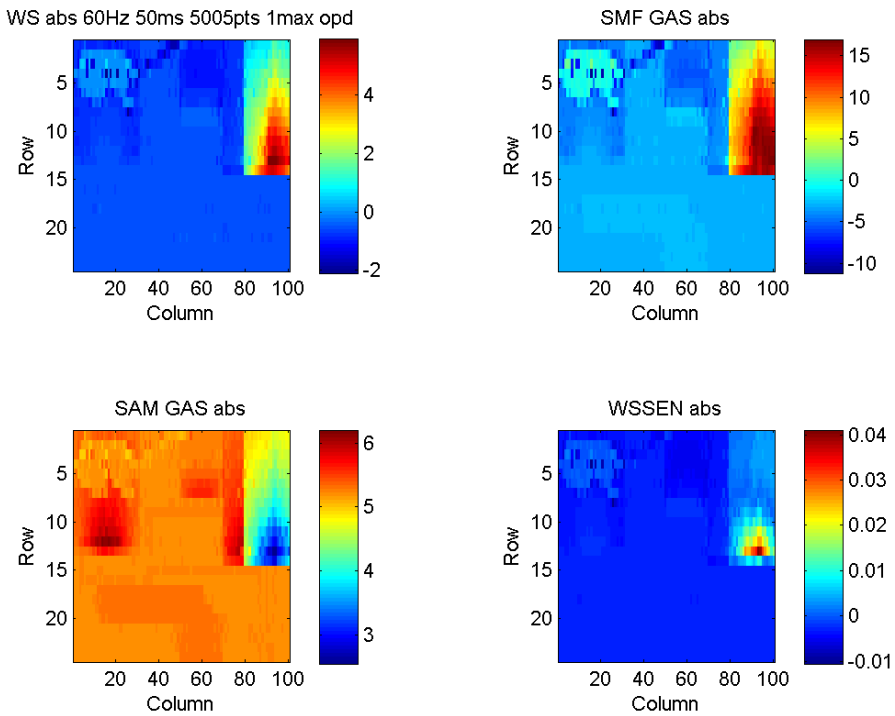

FIGURE 4.34 DETECTION PLANE SERIES FOR OVERLAPPING BENZENE AND METHYL CHLORIDE GAS PLUMES WITH JITTER

4.2.3 DIRSIG Scene Studies

One of the drawbacks of the material maps scene was that it did not have enough variability to produce FA, so beyond the detection planes, it was difficult in characterizing the system performance. The increased realism and variability of the DIRSIG scenes allowed for multiple FA and gave a greater ability to characterize performance.

Figure 4.35 shows detection plane results for the same collection conditions as shown in the previous section. The top panel shows the gas plume being released in an open field, while the bottom panel shows the release point has been moved up in to the tree line. The open field release proved to be a very easy target due to the low spectral and emissivity variability in the field and the high thermal contrast. The detection plane shows that this benzene plume was well characterized.

The second panel shows the results for the plume release in the trees. As can be seen from the detection plane metrics, the plume was a harder target. This was in part due to the obscuration of the ground-level plume by the canopy structure of the taller trees, but also the variability in emissivity and temperature of the ground elements which produced wide ranges of thermal contrast that affected the appearance of the spectral feature. Note, however that the detection score ranges were all largely the same, with the exception of the WSSEN score. This is as they should be because the plume was identical in both scenes and only the release point was changed. Variations in score distributions are due to the aforementioned causes.

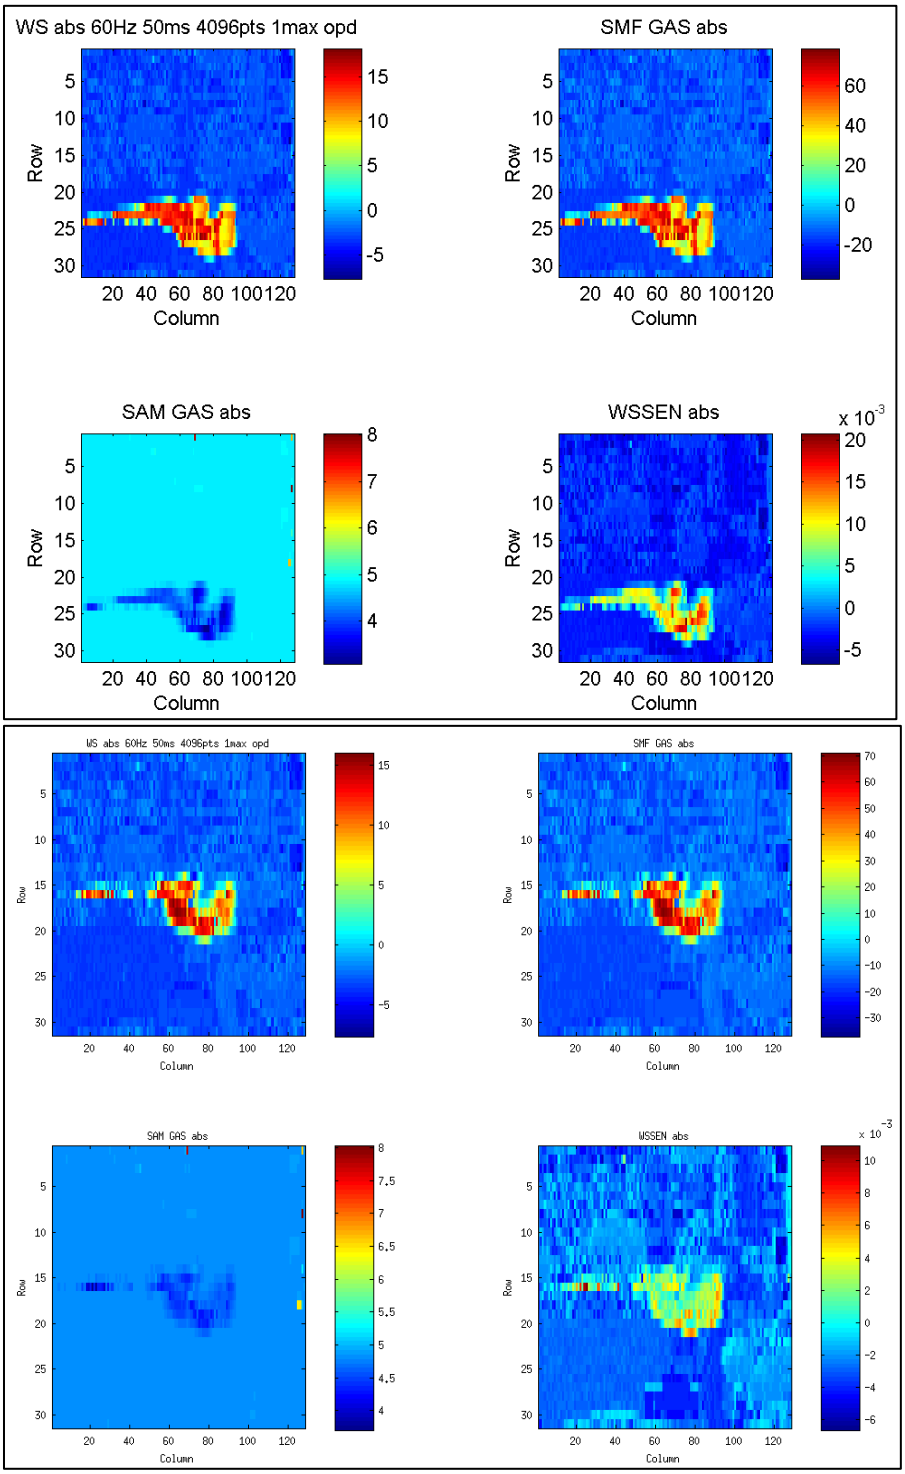

FIGURE 4.35 DIRSIG SCENE DETECTION PLANES WITH GAS RELEASE IN OPEN FIELD (TOP) AND IN TREE LINE (BOTTOM)

Figure 4.36 consists of two images that show a broadband LWIR image of the DIRSIG scene with the benzene plume present in the tree line (the plume is not visible in broadband), while the image on the right is a spectral depth image of the benzene plume of the same scene. Spectral depth was previously addressed in the detection metric section 3.1.4 and is essentially a measure of the distance between the spectral peak of benzene at 1038 cm^{-1} from the baseline of the spectrum in the SRR with adjustments made for the presence of ozone. This image gives a much clearer view of the obscuration of the plume due to the tree line, note however, that the plume does not interact with the 3D objects in the scene. The eddies and propagation of the plume are not affected by the presence of the trees or other objects in the scene.

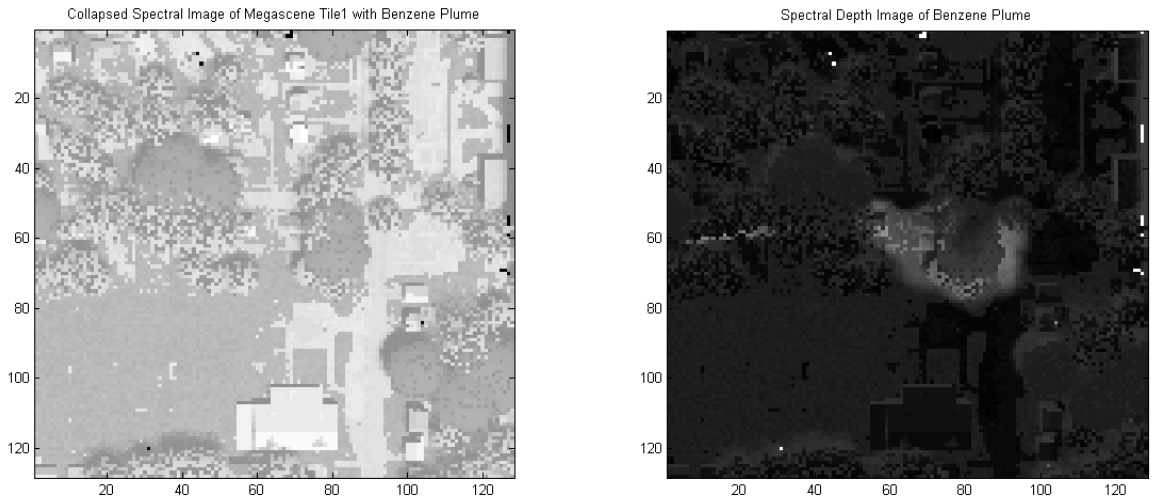

FIGURE 4.36 LWIR IMAGE OF DIRSIG SCENE (LEFT), SPECTRAL DEPTH IMAGE OF BENZENE PLUME IN DIRSIG SCENE (RIGHT)

DIRSIG scenes were produced at each of the atmospheric profiles used in this work. They are rendered with the proper optical effects and radiative transport properties for the specified platform altitude. Thus the previous images, taken at 1 km altitude, present a much larger gas plume in terms of percentage of the scene, whereas those rendered at a 5 km altitude yield a very small target. Figure 4.37 shows the same LWIR broadband image as shown previously rendered from 5 km MSL.

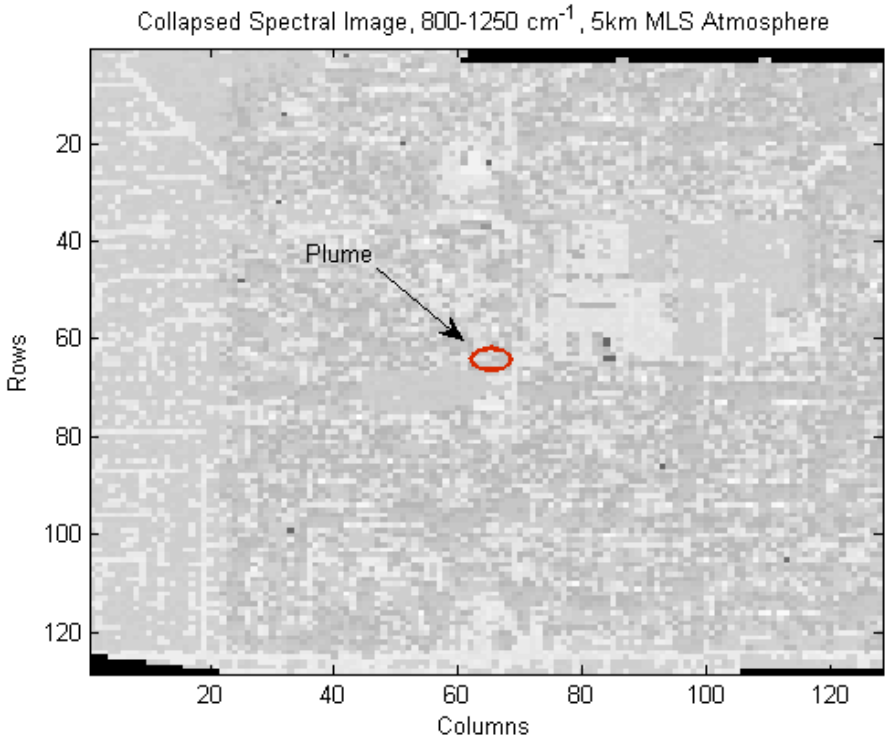

FIGURE 4.37 BROADBAND LWIR DIRSIG SCENE WITH ANNOTATED BENZENE PLUME AT 5 KM MSL PLATFORM ALTITUDE AND 1M GSD

Note the much smaller target the plume represents as a function of spatial area at this altitude. The total extent of the plume is annotated on the image and comprises about 37 pixels, or 0.225% of the image. The dark areas in the scene at the extreme top and bottom of the image are areas where 'the world ends,' in that there is no rendered scene at these locations due to having reached the edge of the modeled scene map.

Surface plots illustrating the tradespace between scan rate and MOPD were created for the DIRSIG scenes as well. These show somewhat similar results to those shown earlier for the material map scene. Figure 4.38 shows the result of the maximum WSSEN score for multiple DIRSIG scene collects at 2, 5, 10, 20, and 40 Hz at 0.2, 0.4, 0.6, 0.8, and 1 cm MOPD. The atmosphere was a 1 km MLS profile for the benzene plume characterized in the experimental section. The target vector was 20,000 ppm-m benzene at 290K over a 300K blackbody. The data shows a general improvement in detection score with increasing scan rate and MOPD with the exception of an anomaly at 5 Hz. This anomaly was due to the fact that at 5 Hz there are three total bins across the image and the first bin happened to capture the entire plume exclusively, yielding abnormally large scores. This higher score trend at 5 Hz was a simple artifact of the plume position coupled with the starting point for the scan at this rate. The data also shows that 20 Hz might be sufficient to characterize the gas plume as there is only a modest increase in score with the 40 Hz scan rate. Figure 4.39 shows the tradespace of the same collection conditions as the previous results but with the inclusion of fine mirror error with a mean displacement of 10% of the sample interval.

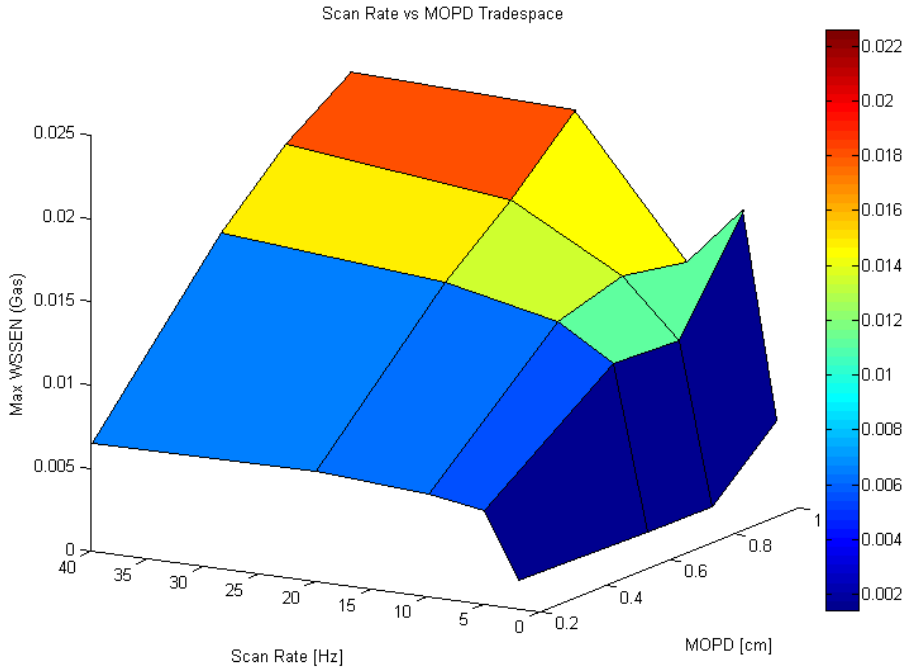

FIGURE 4.38 SCAN RATE VS. MOPD TRADESPACE FOR DIRSIG SCENE WITH BENZENE PLUME IN TREE LINE WITH 1 KM MLS

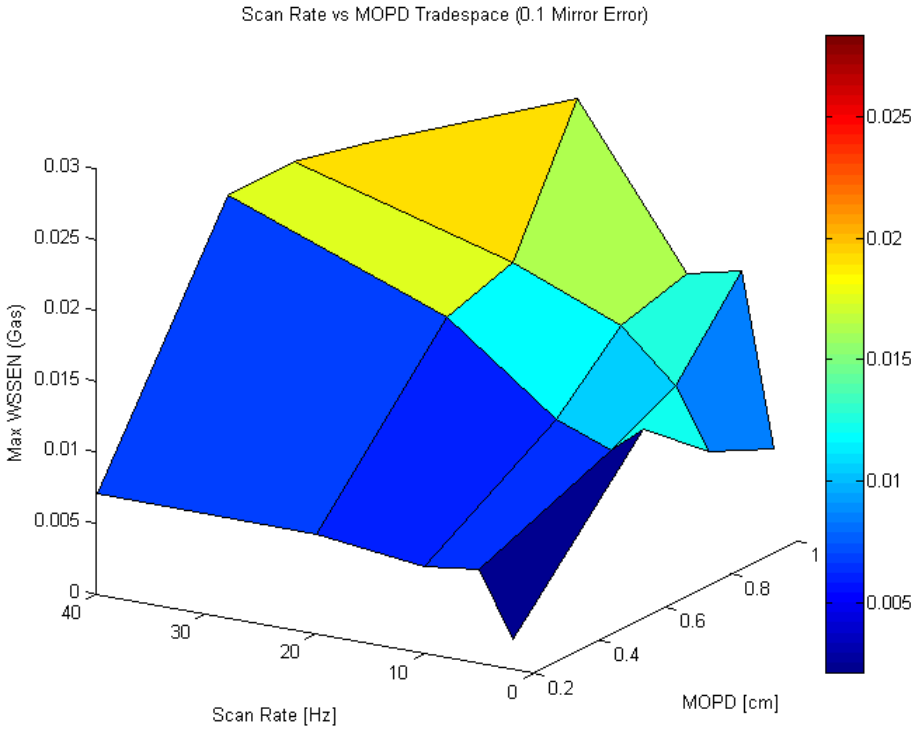

FIGURE 4.39 SCAN RATE VS. MOPD TRADESPACE FOR DIRSIG SCENE WITH BENZENE PLUME IN TREE LINE WITH 1 KM MLS AND FINE MIRROR ERROR

Figure 4.39 shows that the fine mirror error appears to have the greatest impact at high resolution, while the increase in maximum score suggests an increased level of FA and/or distorted spectral features.

In an attempt to capture the extent of the spatial characterization of the plume, which was not captured by the maximum WSSEN score metric just shown, the WSSEN scores within the plume confines were summed for each scan rate and MOPD run and used to construct new surface plots. Figure 4.40 shows the result for the same collection conditions as the previous result. This plot shows a substantial improvement in plume spatial characterization with increasing scan rate as one would expect. Increasing scan rate results in more bins of fewer ground elements in the image meaning that a greater portion of the plume is individually characterized in each interferogram. An unexpected result came from the trends by MOPD which show that performance at 0.6 and 1 cm MOPD yielded similar results. This may be another artifact, but this time due to the number of sample points assigned to each bin at this MOPD.

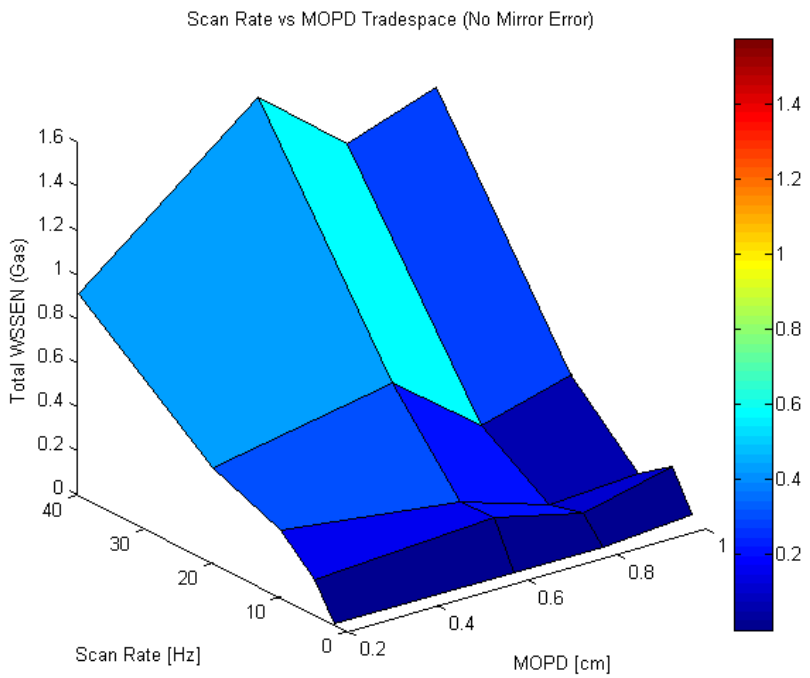

FIGURE 4.40 SCAN RATE VS. MOPD TRADESPACE FOR BENZENE PLUME IN TREE LINE BY TOTAL WSSEN PLUME SCORE

Figure 4.41 shows the tradespace for the same collection conditions as used in the previous results, but with the addition of the same fine mirror error used previously. In this case the sum of the total WSSEN scores were lower than the non-error case. The shape of the plot follows the same general trend as the non-error case with significant emphasis on the 40 Hz, 1 cm MOPD result. It is not clear if this result is due to mirror error corruption of the plume scores, or a true indication of improvement in spatial characterization – prior results would suggest the former.

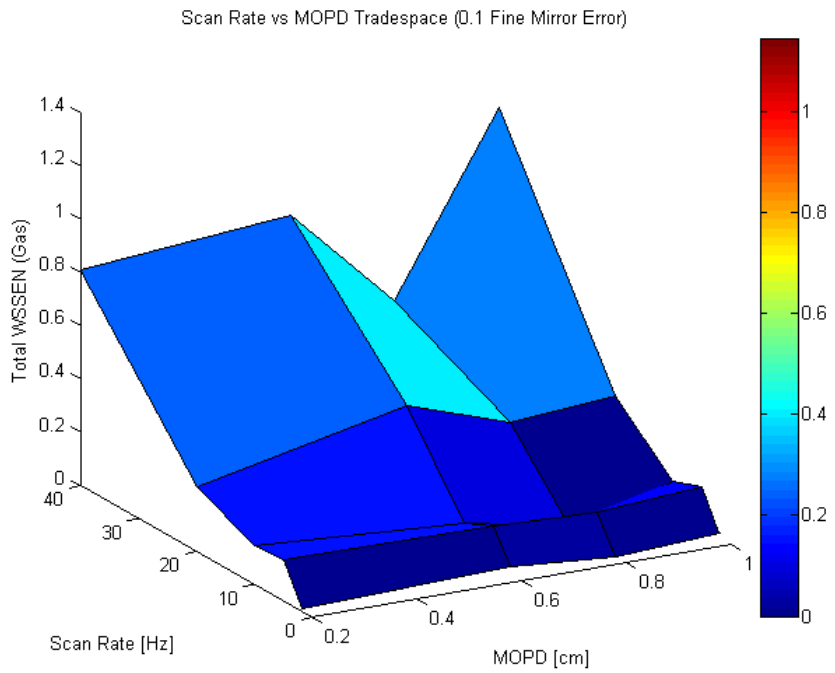

FIGURE 4.41 SCAN RATE VS. MOPD TRADESPACE FOR BENZENE PLUME IN TREE LINE BY TOTAL WSSEN PLUME SCORE WITH FINE MIRROR ERROR

Figure 4.42 shows another way of characterizing the plume scores. Using the same WSSEN metric, the total WSSEN scores from within the confines of the plume were summed, but only above the threshold set by the largest score value outside the plume. This was a way to remove all borderline detects or possible FA from the plume and concentrate on scores that are truly due to the presence of the plume.

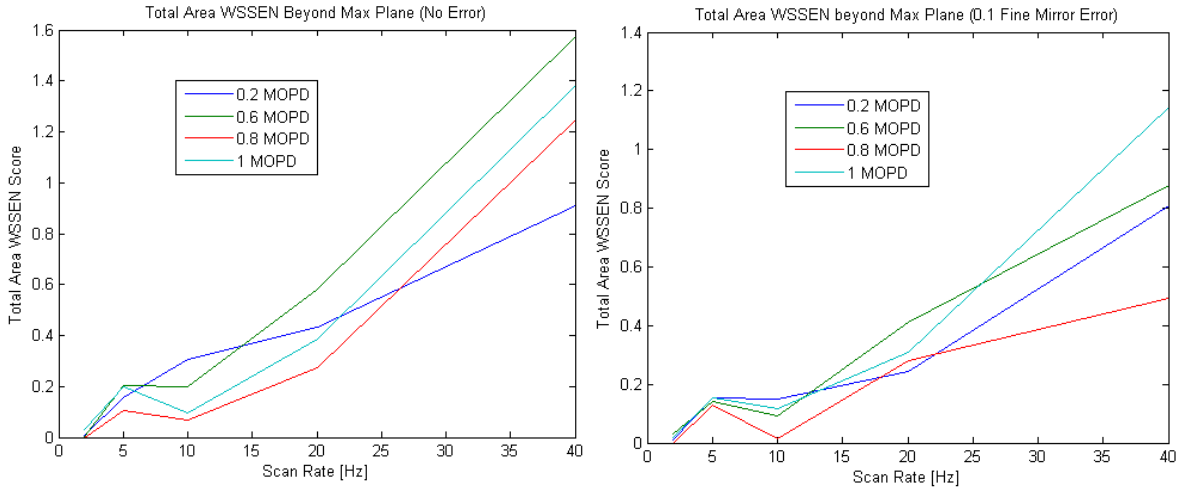

FIGURE 4.42 TOTAL AREA WSSEN SCORE BEYOND MAX PLANE (LEFT), WITH FINE MIRROR ERROR (RIGHT)

Figure 4.42 shows that in the absence of mirror error results in the plot on the left, the 0.6 cm MOPD resolution does indeed outperform the 1 cm MOPD resolution at nearly all scan rates. This may be an artifact of the sampling in the plume. The mirror error version on the right shows that the

0.6 cm and 0.8 cm MOPD resolution in particular, is disrupted. The 1 cm MOPD resolution gives the best characterization, but only at the highest scan rate.

The next series of plots shows another attempt to characterize the detection scores, this time by the difference between the highest in-plume score and the highest score outside of the plume. What this plot is intended to show is that increasing MOPD and scan rate should allow for greater discrimination between scores inside and outside of the plume. This is due to the fact that as resolution is increased, a greater variability in score is seen and especially so at higher scan rates as the bin size decreases. Figure 4.43 shows the results from this metric.

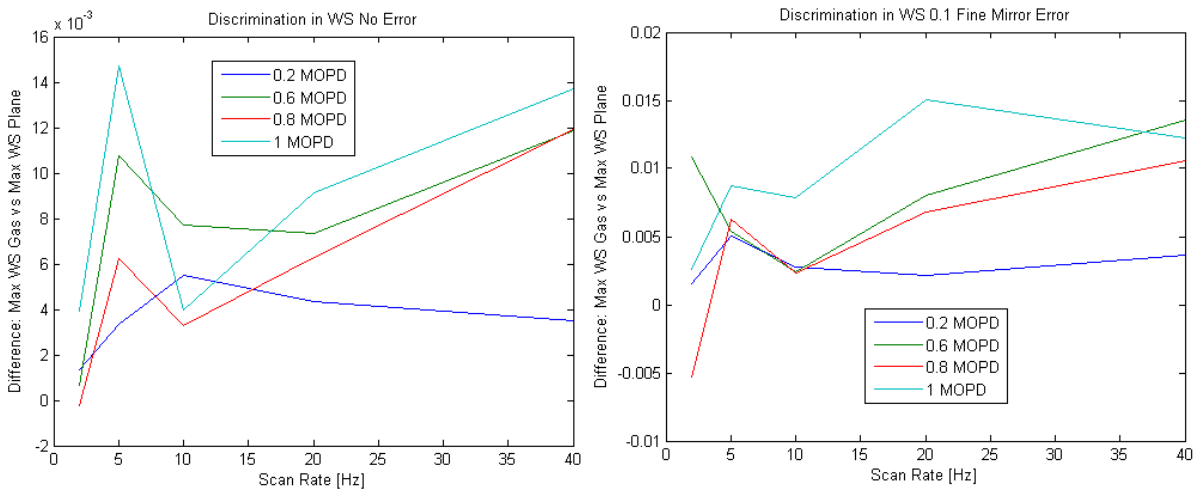

FIGURE 4.43 DIFFERENCE IN THE MAX. WSSEN SCORE IN THE GAS PLUME AND MAX. SCORE OUTSIDE PLUME (LEFT), WITH FINE MIRROR ERROR (RIGHT)

The plot on the left shows great deference to the 5 Hz scan rate, which has the aforementioned anomaly. However, in general the expected trend is found in that the highest MOPD and scan rate settings do result in the greatest discrimination between plume and non-plume regions in most cases. The flat result at 0.2 cm MOPD is due to the fact that the image is sampled in one bin at this platform velocity, yielding very little difference between high scores in and out of the plume. The same trend is generally seen in the plot at right, with deviations attributable to the mirror error.

The next series of plots shows the final attempt to characterize detection performance of the plume in the DIRSIG scene under the presence of FA. Several ROC curves were created as previously described. These curves were limited to a FA rate of 0.01, which stressed the metric to determine which MOPD and scan rate is best at essentially detecting the source of the leak. This point would generally be the highest scoring position, but not always, depending on the amount of scan binning. Figure 4.44 shows the progression through the series. Each plot contains the ROC curves at each MOPD and the series progresses through each scan rate. The 2 Hz scan rate shows that due to a reduced number of scans, the first detection score goes out to 0.025 since each scan makes up a larger proportion of the total number of detection scores.

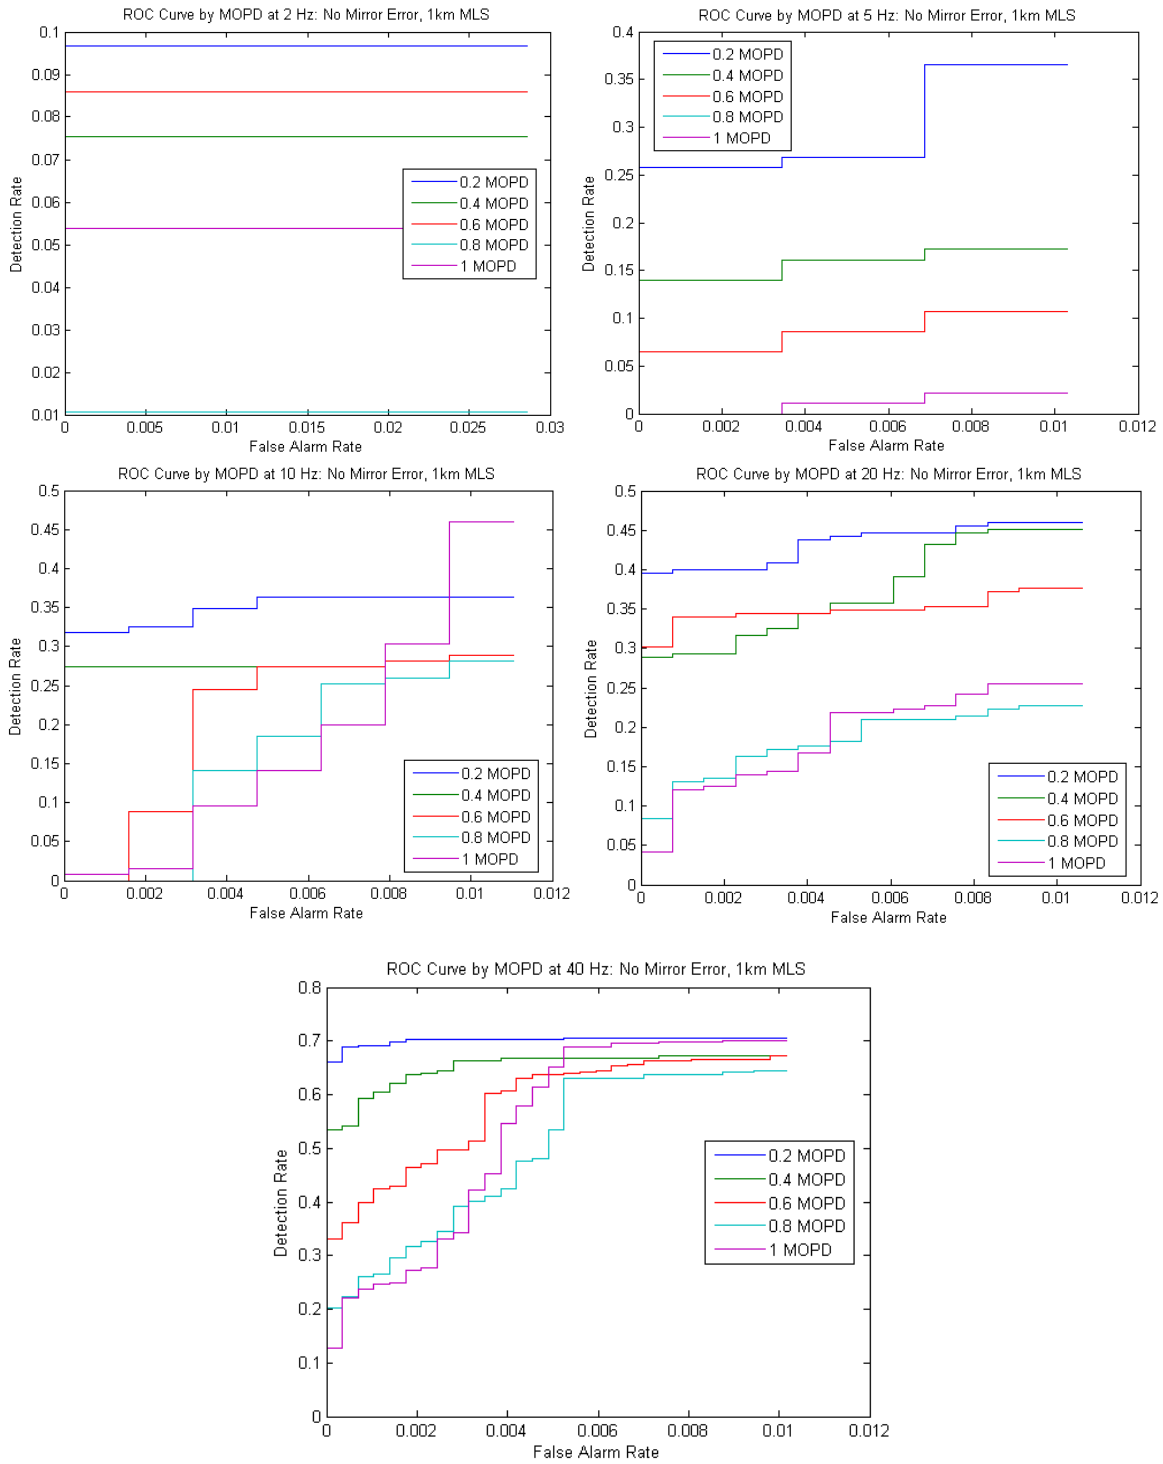

FIGURE 4.44 ROC CURVES AT 0.01 FA RATE, SERIES BY SCAN RATE

Figure 4.44 shows that in terms of collecting data indicating the presence of a gas, the 0.2 cm MOPD setting may yield the best results at virtually all scan rates. However, as the FA rate is allowed to increase beyond what is shown here, the 1 cm MOPD quickly becomes the best detection rate setting as seen in Figure 3.55.

The next series of figures are bar charts that roll-up the results of each ROC curve by summing the area under the curve between the 0 and 0.01 FA rate. The results are shown as grouped by MOPD for each scan rate. The top chart in each series is for a given atmosphere with no mirror error, while the bottom chart was for the same atmosphere with mirror error. The trend for the results in general was that the 40 Hz scan rate gave the best performance across nearly all MOPD values and this is likely due to the increased spatial characterization of the plume at this scan rate.

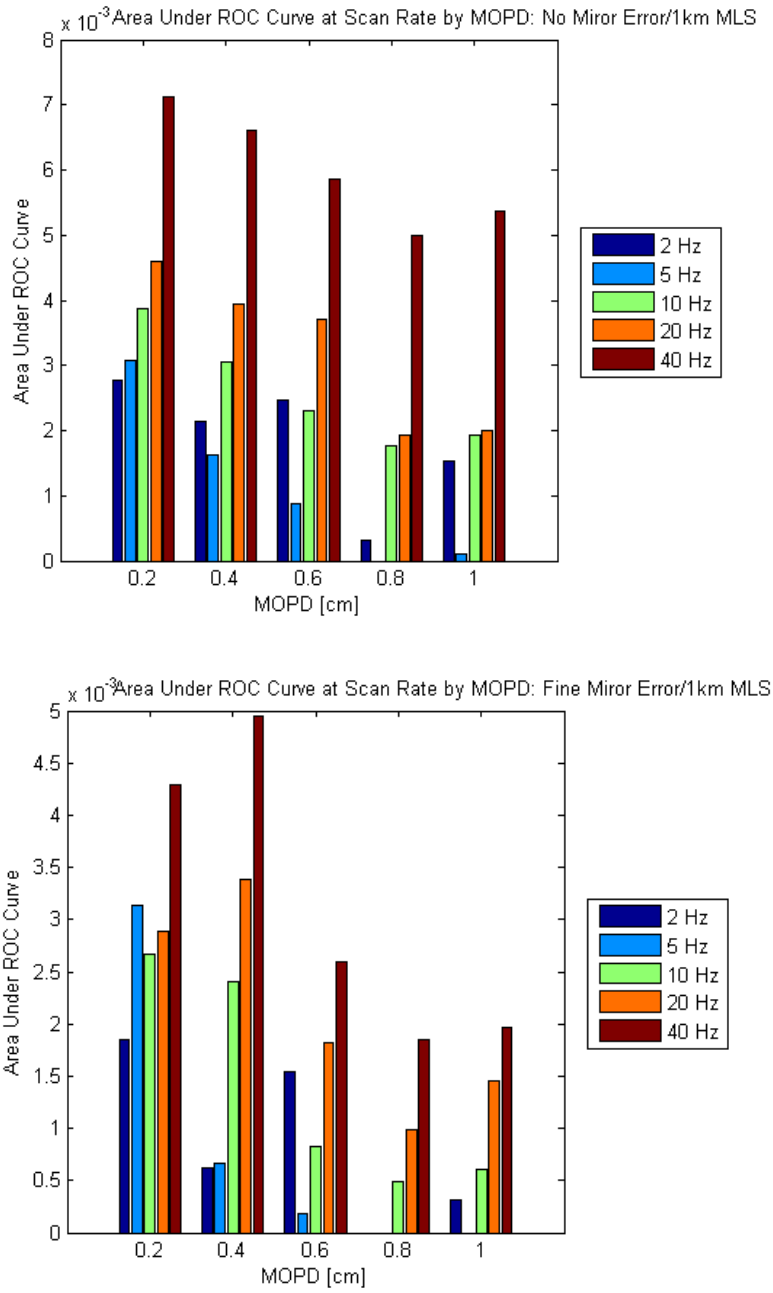

FIGURE 4.45 AREA UNDER ROC CURVE FOR 1 KM MLS WITH NO ERROR (TOP), FINE MIRROR ERROR (BOTTOM)

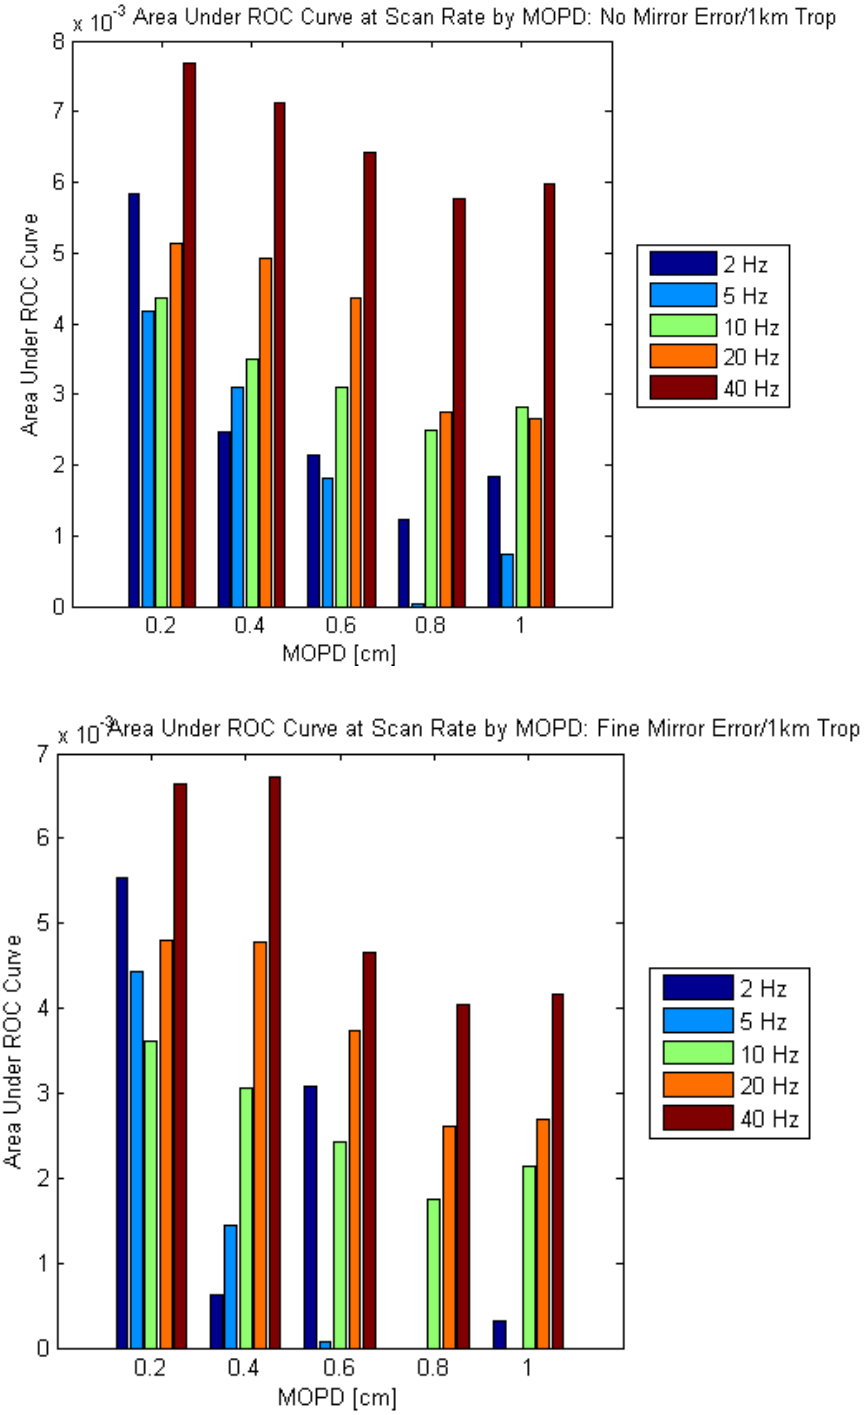

FIGURE 4.46 AREA UNDER ROC CURVE FOR 1 KM TROPICAL WITH NO ERROR (TOP), FINE MIRROR ERROR (BOTTOM)

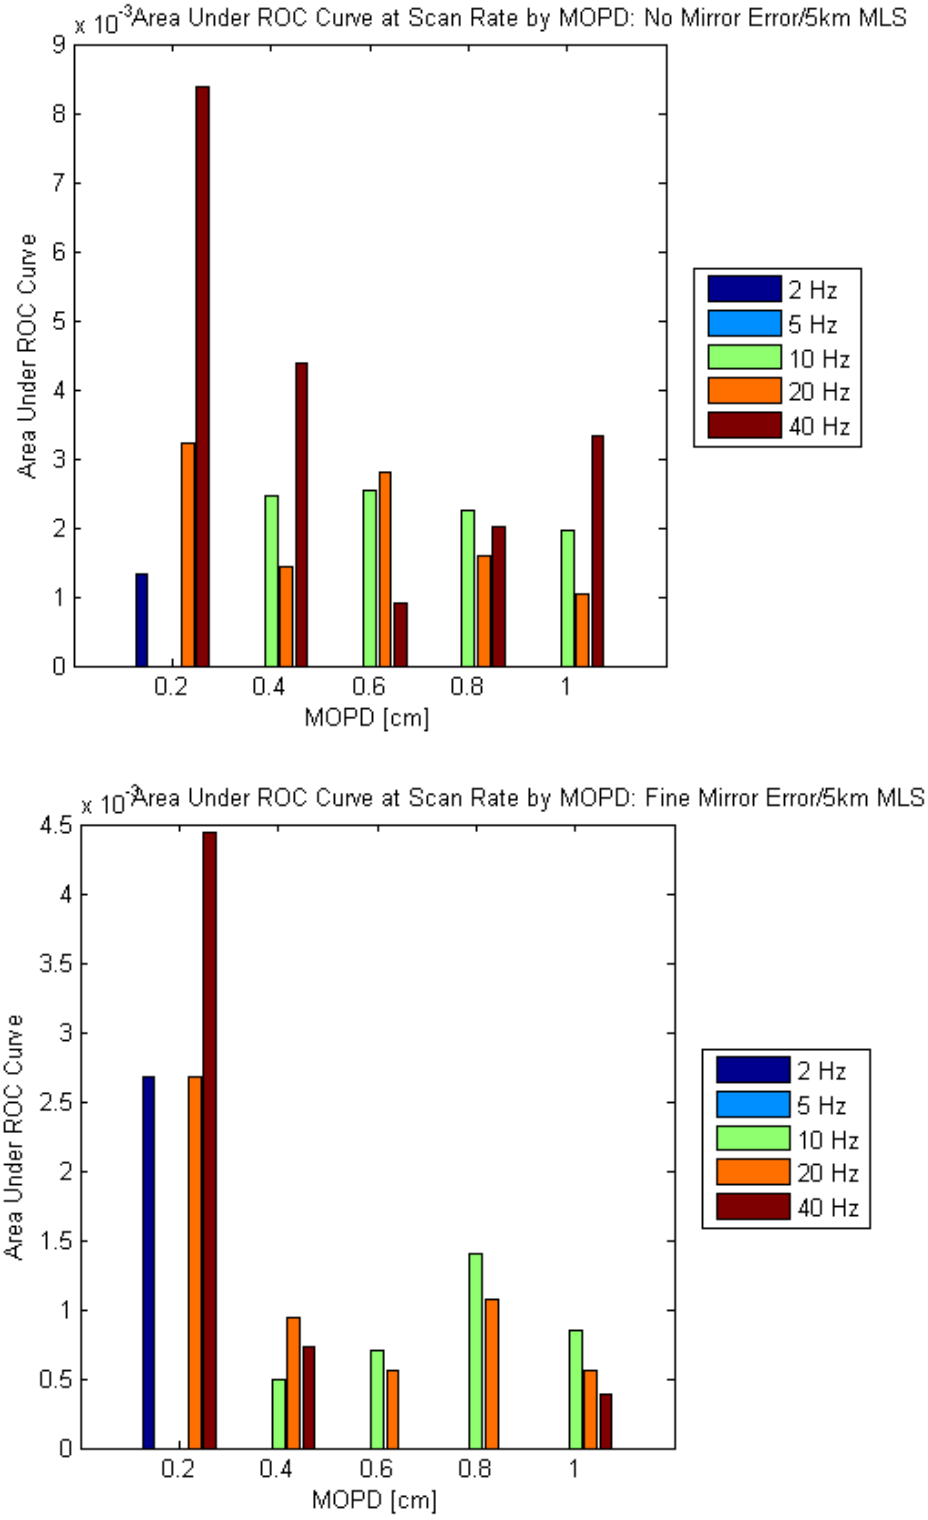

FIGURE 4.47 AREA UNDER ROC CURVE FOR 5 KM MLS WITH NO ERROR (TOP), FINE MIRROR ERROR (BOTTOM)

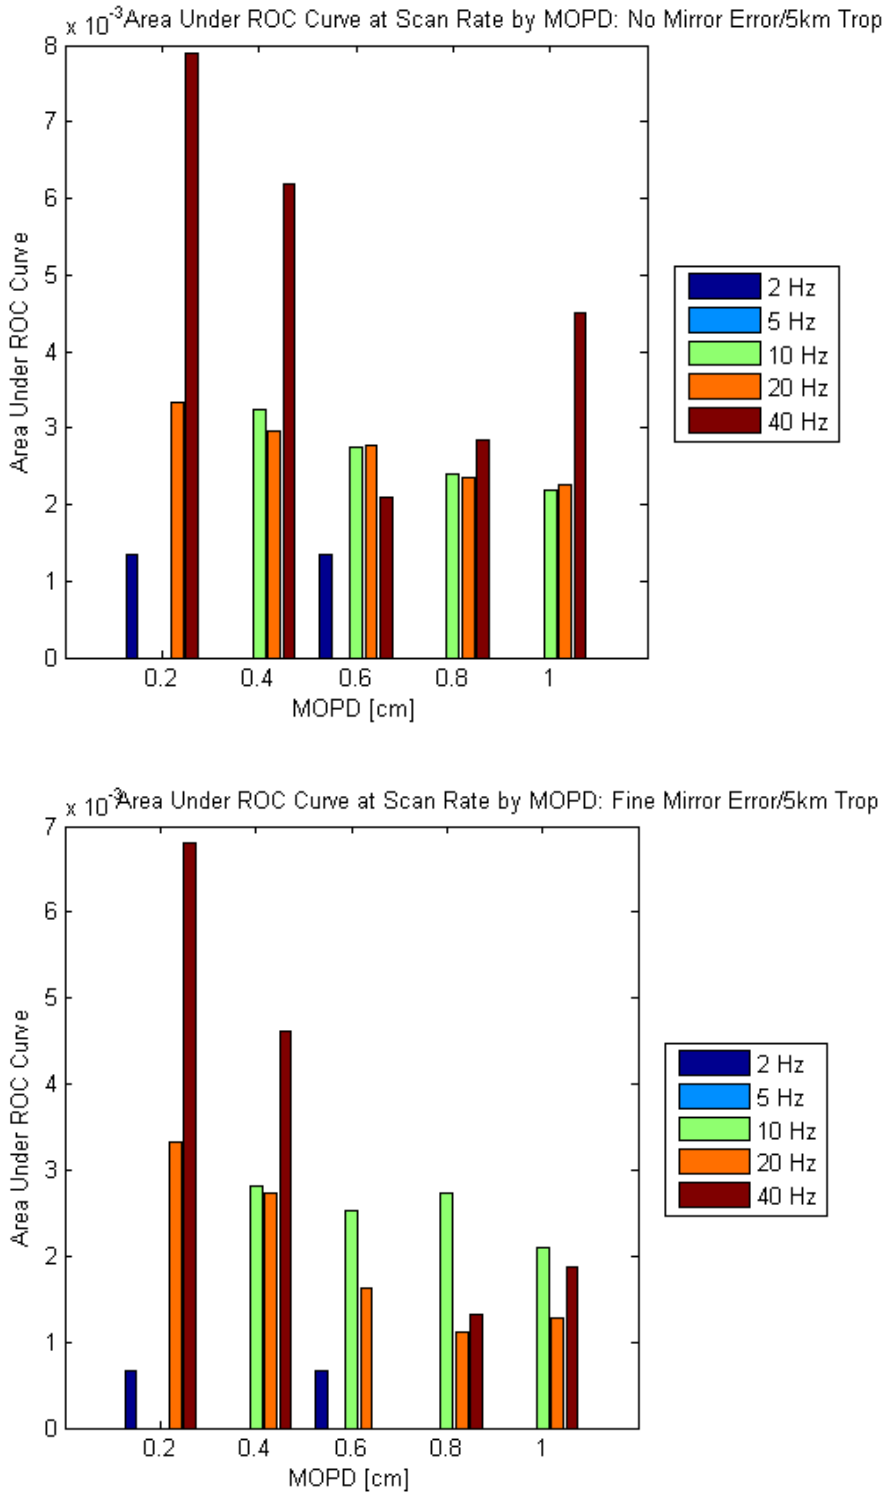

FIGURE 4.48 AREA UNDER ROC CURVE FOR 5KM TROPICAL WITH NO ERROR (TOP), FINE MIRROR ERROR (BOTTOM)

To summarize the results, very little change was seen between the 1 km MLS and tropical atmospheric profiles in terms of detection performance as previously demonstrated, the addition of water vapor has little impact to the detectability of benzene. At the 5 km altitude, the target

became much harder to find and the 2 and 5 Hz scan rates proved to be ineffective at nearly all MOPD values due to the substantial reduction in plume samples at the higher altitude in one and three bins, respectively. At 10 Hz, the target plume was characterized with some confidence with nearly equal performance at 20 Hz. The 40 Hz scan rate was affected by the increased presence of ozone, particularly at low resolution where ozone and the benzene spectral feature become indistinguishable, resulting in an unreliable high score. From a 0.6 cm MOPD resolution and on, it is clear that the discrimination between benzene and ozone has improved and the scores continue to improve with increasing MOPD. Another interesting observation from the 5km data set is that the mirror error had a particularly damaging affect on many of the scores at the 40 Hz scan rate.

4.3 Summary of Results

What follows is a bullet list summary of the major results from all of the trials, experiments, and studies conducted in this work.

- The airborne FTS model demonstrated that when a gas is present in multiple regions in the interferogram it can be detected from a scanning platform, even if only a portion of the interferogram is attributable to samples from the gas spectral feature. Some semblance of the gas spectral feature will be transformed and can be detected with robust detection metrics.
    - This point allows for scanning to be accomplished at low scan rates or higher platform velocities that only allow capture a portion of the gas plume in the scan while still providing a reasonable chance for detection.
- When a gas is present in multiple regions in the interferogram it is more resilient to error sources such as LOS jitter and mirror uncertainty error.
- Atmospheric degradations to high resolution gas detection in the LWIR are primarily due to the increased presence of water vapor and ozone due to the proximity of these features, as well as the similar size and shape to the spectral features of many target gases used in this work.
- The source of the impact of the background in LWIR high resolution gas detection is not due to spectral variation, as most target gas spectral features are finer than those of the background, but rather variation in emissivity, which causes significant changes to the thermal contrast which results in changes to the shape of the spectral feature. These changes can result in a muted absorption for a given CPL and thermal contrast, reduction of thermal contrast to zero making the gas undetectable, or a transition from absorption to emission. These effects are most pronounced in low emissivity, man-made materials. The

variation in the background primarily used in this work, up to 15%, resulted in only minor impact to the detection of target gases.

- Further, it was posited that rapidly shifting the radiance values between OPD samples due to changes in the atmosphere or background material or presence of the gas plume would cause anomalies in the transformed spectrum. The majority of these sources only varied by about 10% in total SRR and were well tolerated by the FTS model.

- Detection of a target gas in the presence of jitter, and to a lesser extent, mirror error, is dependent on where in the interferogram the error occurs. If the error removes or corrupts the one region in the interferogram that was sensitive to the target gas spectral feature, the gas will not be detected. Fortunately many gases, particularly at high values of MOPD will appear in multiple regions in the interferogram, increasing their resiliency to error sources. If the error occurs late in the interferogram, the high resolution spectral details will be corrupted or removed, leaving the bias where the spectral feature would have been. This may still result in a positive detection, particularly for SMF. The opposite case was also demonstrated in this work.

- Despite significant disruption to the interferogram due to mirror error, including critical sampling errors, while all bias will be lost resulting in a loss of radiometric accuracy, in many cases recognizable spectral features from the gas can still be found among the noise and spectral artifacts.

- The FTS model demonstrated that the nature of interferogram formation under ideal optical conditions on a scanning, moving platform does not preclude the ability to detect the presence of a gas plume in a scene and even supports identification at high levels of MOPD. The fact that these scan rates are not achievable today may explain the lack of success in high resolution airborne FTS systems to date. Consideration of additional factors suggested in the future work section may place further restrictions on the performance of the system that have not been previously considered.

5 Conclusion and Future Work

The objectives of this work were multifaceted. It began with a thrust to understand how the unique way an interferometer forms an interferogram and resulting transformed spectrum is impacted by environmental conditions, instrumental settings, and major error sources. The results of this effort were compelling and needed to be investigated from an operational approach. The data obtained from the early work provided an operational envelope from which performance trades could be made for a simulated collection platform providing spectra from a realistic gas plume in a highly complex scene. The ability of the model to be run against any SRR data cube allows the ability for specific customization of a scene to determine the efficacy of an airborne FTS for a given detection scenario. Further, the modeling of anticipated errors and scan rates can provide design constraints in the engineering phase of a system that would allow for realistic expectations of performance within an error tolerance interval.

5.1 Conclusion

Multiple objectives were achieved in this work. First, the major performance parameters that drive fugitive gas detection in an FTS instrument were mapped and their interrelationships determined. These parameters were screened in initial studies to determine those with maximal impact and carried forward for implementation into an airborne Michelson-based FTS computer model, capable of both single point and scanning spectral image formation, for characterization and tradespace evaluation.

Multiple detection metrics including the SAM, SMF, Spectral Ratio, Spectral Energy/Spectral Depth, and Spectral Energy Normalized Spectral Ratio were implemented to quantitatively assess the impact of each parameter as independent from SNR. These detection metrics were used against Monte Carlo random trials to measure the impact and likely operational error envelopes of employing uncertainty in the OPD positional measurement and LOS jitter in the primary mirror. Multiple performance trade studies were conducted to determine detection performance under variable atmospheric, surface emissivity, surface temperature, gas temperature and CPL, instrument resolution, and instrument error scenarios.

Finally, scene-based scenarios were developed to maximize the variability in the aforementioned parameters to determine FTS instrument performance and operational parameters in a scanning sense across a scene containing a target gas. The results from these scenarios indicated the ability of a scanning FTS system to detect a fugitive gas leak from 1 to 5 km altitude under various atmospheric conditions, even when only a portion of the interferogram was attributable to spectral information from the target gas. The study illustrated the ability to rapidly identify the source of the leak at a low MOPD while being able to fully characterize the extent of the plume at higher MOPD.

5.2 Future Work

In terms of future effort, much can be done to upgrade the fidelity of the FTS model. While duplication of effort of aspects already accomplished by Allain's DIRSIG model was avoided in part, much improvement can be made to both models. The first step would consist of adding in all of the optical and detector effects ignored in this work. This would include a remodeling of SRR to calculate the correct atmospheric slant path distance from each ground element in the IFOV to the aperture at each OPD sample point (or some interpolated value between them). If this could be done on the fly, a forward staring FTS model would become possible, allowing for a rock point to be established. It was argued early in the proposal for this work that the geometry of a forward staring FTS would allow for greater success in that the average CPL would be higher, the fluctuations in background and atmosphere would be less dramatic, and the time allowed for multiple scans to increase SNR would be afforded. This represents another set of performance trades to evaluate. Of course, in an operational sense, the gas location would already have been identified via a nadir detection scheme, then approached again in a forward staring mode, similar to the ACIS implementation of a vibration stabilized, steerable pod.

Next, all optical effects including stray light, off-axis rays, and optical losses would need to be accounted for. One study suggests that as the maximum OPD is extended, the assumption of collimated light in the interferometer becomes less valid resulting in a narrowing of the beam. This narrowing reduces spatial accuracy and induces off axis effects (Smith, 1996). This angular divergence in the beam causes longer path lengths to be traveled by the outer edges of the beam as compared to the axial rays, which causes spatial errors in the interference fringe pattern when the beams recombine across the FPA. Apertures can be placed along the moving mirror arm to shape the beam again, but result in a reduction of overall optical throughput, reducing the SNR. Further, Smith points out that as the OPD becomes greater, the amplitude of the measured irradiance is less while the same magnitude of [dark] noise is present at every interferogram sample. Thus the SNR at each OPD sample point at longer OPD necessarily declines, possibly causing the MOPD to become self-limiting at a given SNR. Realistic OPD sample point SNRs would need to be determined and included in the model to understand the most likely MOPD achievable under a given set of collection conditions.

The addition of a modeled detector response, detector noise, F number, detector area, and instrument self-emission would add realism to simulation and help determine the impact to SNR of the various simulated errors and conditions. This would allow calculation of additional SNR-based metrics such as MDQ or the Noise Equivalent Concentration Pathlength (NECL) for comparison of system performance across sensor types. Further, the impact of quantization and associated clipping (particularly at ZPD) to radiometric accuracy could be studied and the impact to resulting spectra determined. The large dynamic range presented by the first several OPD points presents a challenge to quantization efforts, particularly if a dynamic scale model or extremely deep quantum wells are not used (Smith, 1996). If the ZPD value is clipped due to saturation, the radiometric accuracy of the entire spectrum is lost, unless the points can be recovered by other methods, such as estimation or

curve fitting/modeling. Finally, the optical frequency due to the moving mirror velocity can be used to drive the data rate and detector sampling rate to estimate system date rates (throughput) at various MOPD or scan rate settings, to determine SNR impacts with increasing detector sampling rates, and SNR impacts from increased mirror velocities. This relationship between a detector sampling-limited mirror velocity, which limits scan rate, and ultimately MOPD, would be interesting trade space to explore.

The scanning module can be improved to include integration time effects at each OPD sample to determine impact on the interferogram, particularly with increasing mirror velocity and/or sample points. Pitch, roll, and yaw effects can be added to the uncertainty in the sensor LOS and the impact to the resulting spectra can be characterized.

Finally, to truly characterize the performance of an imaging FTS on a given scene, the spatial registration of the interference pattern projected on the detection plane would need to be modeled, with allowance made for differences in SRR due to off-nadir image points and keystone effects from optical issues. If this were in place, the model could be converted from a pushbroom that performs the task of a framing array, to a true framing array model.

The combination of all of the current model capabilities with the aforementioned 'realism' suite of capabilities would yield an unrivaled FTS model, suitable for the simulation and trade space evaluation of competing designs, or even real systems.

References

Alain, F. (1999). *Simulation of Imaging Fourier Transform Spectrometers Using DIRSIG.* Rochester: Thesis, Rochester Institute of Technology.

Allard, J.-P., Chamberland, M., Farley, V., & Villemaire, A. (2008). Airborne Measurements in the Longwave Infrared Using an Imaging Hyperspectral Sensor. *Imaging Spectrometry XIII*, *Proceedings of SPIE 7086*, pp. 1-12.

Beer, R. (1992). *Remote Sensing by Fourier Transform Spectroscopy* (Vol. 120). New York: John Wiley & Sons, Inc.

Bennett, C. L. (1994). FTIR Measurements of Thermal Infrared Sky Radiance Transmission. *SPIE , 2266*, 25-35.

Bennett, C. L. (1997). The Effect of Jitter on an Imaging FTIR Spectrometer. *SPIE , 3063*, 174-184.

Bennett, C. L., Carter, M. R., & Fields, D. J. (1995). Hyperspectral Imaging in the Infrared Using LIFTIRS. *SPIE , 2552*, 274-283.

Bennett, C. L., Carter, M. R., Fields, D. J., & Lee, F. D. (1995). Infrared Hyperspectral Imaging Results from Vapor Plume Experiments. *SPIE , 2480*, 435-444.

Berk, A., Bernstein, L., & Robertson, D. (1989). *MODTRAN: A Moderate Resolution Model for LOWTRAN 7.* Burlington, MA: Specral Sciences, Inc.

Berk, A., Cooley, T. W., Anderson, G., & Acharya, P. K. (2004). MODTRAN5: A Reformulated Atmospheric Band Model with Auxiliary Species and Practical Multiple Scattering Options. *Remote Sensing of Clouds and the Atmosphere* (pp. 78-85). Proceedings of SPIE.

Blackadar, A. K. (1997). *Turbulence and Diffusion in the Atmosphere.* New York: Springer.

Boonmee, M. (2007). *PhD Dissertation: Land Surface and Emissivity Retrieval from Thermal Infrared Hyperspectral Imagery.* Rochester: RIT.

Brown, S. (2010, May 7). DIRSIG Developer, DIRS Group Staff Scientist. (A. Weiner, Interviewer)

Carlson, R. C., Hayden, A. F., & Telfair, W. B. (1 December 1988). Remote Observations of Effluents from Small Building Smokestacks Using FTIR Spectroscopy. *Applied Optics , 27* (23), 4952-4959.

Carter, M. R., Bennett, C. L., Fields, D. J., & Hernandez, J. (1993). Gaseous Effluent Monitoring and Identification Using an Imaging Fourier Transform Spectrometer. *SPIE , 2092 Substance Detection Systems*, 16-26.

Carter, M. R., Bennett, C. L., Fields, D. J., & Hernandez, J. (1993). Imaging Fourier Transform Spectrometer. *SPIE , 1937*, 191-200.

Carter, M. R., Bennett, C. L., Fields, D. J., & Lee, F. D. (1995). Livermore Imaging Fourier Transform Infrared Spectrometer (LIFTIRS). *SPIE , 2480*, 380-386.

Chaffin, C., Marshall, T., & Kroutil, R. (1995). Passive Fourier Transform Infrared (FT-IR) Monitoring of SO2 in Smokestack Plumes: A Comparison of Remote Passive Spectra of an Actual Hot Plume with Emission Spectra Collected with a Heatable Cell. In O. A. Simpson (Ed.), *Optical Sensing for Environmental and Process Monitoring*, (pp. 302-313).

Chamberland, M., Villemaire, A., & Tremblay, P. (2004). Performance Model of Imaging FTS as a Standoff Chemical Agent Detection Tool. *Chemical and Biological Standoff Detection*, *Proceedings of SPIE 5268*, pp. 240-251.

Chemical Rubber Company. (1992). *Handbook of Chemistry and Physics* (73 ed.). (D. R. Lyde, Ed.) CRC Publishing Co.

Clark, R., Swayze, K., & Livo, K. (2003). Imaging Spectroscopy: Earth and planetary remote sensing with the USGS Tetracorder and expert systems. *J. Geophys. Res.*

Cohen, D. L. (1997). Performance degradation of a Michelson interferometer when its misalignment angle is rapidly varying, random time series. *Applied Optics , 36* (18), 4034-4042.

Cohen, D. L. (1999). Performance degredation of a Michelson interferometer due to random sampling errors. *Applied Optics , 38* (1), 139-151.

D'Amico, F. M., Emge, D. K., & Marinelli, W. J. (2004). The Airborne Chemical Imaging System (ACIS). *Chemical and Biological Standoff Detection*, *Proceedings of SPIE 5268*, pp. 194-201.

Digital Imaging and Remote Sensing Laboratory. (2006). *The DIRSIG User's Manual.* Rochester, NY: Rochester Institute of Technology.

Elachi, C., & van Zyl, J. (2006). *Introduction To The Physics and Techniques of Remote Sensing* (2 ed.). Hoboken, New Jersey: John Wiley & Sons, Inc.

Farley, V., Chamberland, M., & Vallieres, A. (2006). Radiometric Calibration Stability of the FIRST: a Longwave Infrared Hyperspectral Imaging Sensor. *Infrared Technology Applications XXXII*, *Proceedings of SPIE 6206*, pp. 1-11.

Farley, V., Chamberland, M., & Villemaire, A. (2007). Chemical Agent Detection and Identification with a Hyperspectral Imaging Infrared Sensor. *Imaging Spectrometry XII*, *Proceedings of SPIE 6661*, pp. 1-9.

Flanigan, D. F., & Quinn, T. G. (1996). Model for the Prediction of Sensitivity of Passive Fourier Transform Infrared (FTIR) Sensors. *SPIE , 2763*, 96-116.

Foy, B. R., Petrin, R. R., Quick, C. R., Shimada, T., & Tiee, J. J. (2002). Comparisons Between Hyperspectral Passive and Multispectral Active Sensor Measurements. *Chemical and Biological Sensing III*, *Proceedings of SPIE 4722*, pp. 98-109.

Fry, H., McVey, B., & Schmitt, M. (November 1998). Assessment of Infrared Remote Chemical Sensing Systems with Numerical Simulations. *Air Monitoring and Detection of Chemical and Biological Agents*, *SPIE 3533*, pp. 2-9. Boston.

Funk, C. C., Theiler, J., & Roberts, D. A. (July 2001). Clustering to Improve Matched Filter Detection of Weak Gas Plumes in Hyperspectral Thermal Imagery. *IEEE Transactions on Geoscience and Remote Sensing , 39* (7), 1410-1420.

Gallagher, N. B., Wise, B. M., & Sheen, D. M. (2003). Error Analysis for Estimation of Trace Vapor Concentration Pathlength in Stack Plumes. *Applied Spectroscopy , 57* (6), 614-621.

GATS, Inc. (n.d.). *High-Resolution Spectral Modeling*. Retrieved various various, 2009-2010, from Spectral Calc.com: http://www.spectralcalc.com/info/about.php

Gittins, C. M., & Marinelli, W. J. (1998). AIRIS Multispectral Imaging Chemical Sensor. *Electro-Optical Technology for Remote Chemical Detection and Identification III*, *SPIE 3383*, pp. 65-74. Orlando, Florida.

Griffin, M. K., Kerekes, J. P., Farrar, K. E., & Burke, H. K. (2001). Characterization of Gaseous Effluents from Modeling of LWIR Hyperspectral Measurements. In S. S. Shen, & M. R. Descour (Ed.), *Algorithms for Multispectral, Hyperspectral, and Ultraspectral Imagery VII*, *Proceedings of SPIE 4381*, pp. 360-369.

Griffiths, P. R., & de Haseth, J. A. (2007). *Fourier Transform Infrared Spectrometry* (2nd ed.). Hoboken, New Jersey: John Wiley & Sons, Inc.

Hackwell, J. A., Warren, D. W., & Bongiovi, R. P. (1996). LWIR/MWIR Imaging Hyperspectral Sensor for Airborne and Ground-Based Remote Sensing. *SPIE , 2819*, 102-107.

Hairg, R., Matz, G., & Rusch, P. (2004). Remote Detection of Methane by Infrared Spectrometry for Airborne Pipeline Surveillance: First Reults of Ground-Based Measurements. *Remote Sensing of Clouds and the Atmosphere VIII*, *Proceedings of SPIE 5235*, pp. 435-446.

Harig, R., Martz, G., & Rusch, P. (2002). Scanning Infrared Remote Sensing System for Identification, Visualization, and Quantification of Airborne Pollutants. *Instrumentation for Air Pollution and Global Atmospheric Monitoring*, *Proceedings of SPIE 4574*, pp. 83-94.

Hearn, D. R. (1999). *Fourier Transform Interferometry.* Technical Report 1053, Massachusetts Institute of Technology, Lincoln Laboratory, Lexington, MA.

Hirsch, E., & Agassi, E. (2007). Detection of Gaseous Plumes in IR Hyperspectral Images Using Heirarchical Clustering. *Applied Optics , 46* (25), 6368-6374.

Huheey, J. E., Keiter, E. A., & Keiter, R. L. (1993). *Inorganic Chemistry: Principles of Structure and Reactivity* (4 ed.). New York: HarperCollins College Publishers.

Jellison, G. P., Mitchell, H. J., & Miller, D. P. (2003). Theory, Modeling and Measurements of Gas Plumes. *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery IX*, *Proceedings of SPIE 5093*, pp. 172-183.

Jensen, J. O., Ifarraguerri, A. I., & Loerop, W. R. (2002). Chemical Imaging System. *Instrumentation for Air Pollution and Global Atmospheric Monitoring*, *Proceedings of SPIE 4574*, pp. 52-62.

Kerekes, J. P., Griffin, M. K., Baum, J. E., & Farrar, K. E. (2001). Modeling of LWIR Hyperspectral System Performance for Surface Object and Effluent Detection Applications. In S. S. Shen, & M. R. Descour (Ed.), *Algorithms for Multispectral, Hyperspectral, and Ultraspectral Imagery VII*, *Proceedings of SPIE 4381*, pp. 348-359.

Kroutil, R. J., Knapp, R. B., & Small, G. W. (2002). Airborne Passive FT-IR Spectrometry. *Vibrational Spectroscopy-based Sensor Systems*, *Proceedings of SPIE 4577*, pp. 213-225.

Kroutil, R., Lewis, P. E., Miller, D. P., Shen, S. S., & Small, G. (2008). Implementation of an Automated Signal Processing Approach for the Analysis of Chemical Spectral Signatures Collected from FTIR Mounted in an Aircraft. In S. S. Shen, & P. E. Lewis (Ed.), *Imaging Spectrometry VIII*, *Proceedings of SPIE 7086*, pp. 1-12.

Laidler, K. J., Meiser, J. H., & Santuary, B. C. (2003). *Physical Chemistry* (4 ed.). Boston: Houghton Mifflin Co.

Lewis, P. E., Thomas, M. J., & Miller, D. P. (2004). Encapsulation of the US EPA Standoff Chemical Detection Requirements for Commercial Sensor System Developers. *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery X*, *Proceedings of SPIE 5425*, pp. 218-231.

Lucey, P. G., & Wilcox, B. B. (2003). Mini-SMIFTS: An Uncooled LWIR Hyperspectral Sensor. *Imaging Spectrometry IX*, *Proceedings of SPIE 5159*, pp. 275-282.

Lucey, P. G., Williams, T., & Mignard, M. (1998). AHI: An Airborne Long Wave Infrared Hyperspectral Imager. *Airborne Reconnaissance XXII*, *Proceedings of SPIE*, pp. 36-43. San Diego, California.

Malik, H. (22 July 1997). *Advanced FTIR Singal Processing for Airborne and Spaceborne Remote Sensing of Chemical Clouds.* Phillips Laboratory, Directorate of Geophysics, AIL Systems, Inc. Hanscom AFB, MA: Air Force Materiel Command.

Manning, C., Bowie, B., & Griffiths, P. (2000, March 16). *Very Rapid-Scan Time-Resolved FT-IR Spectrometry*. Retrieved March 22, 2009, from Manning Applied Technologies, Inc.: http://www.appl-tech.com/Pittcon00HTML/Pittcon00HTML.htm

Mitchell, H. J., Hemmer, T. H., Lewis, P. E., & Salvaggio, C. (2001). The Effects of Temporally Changing Sources on Fourier Transform Spectrometers. *Algorithms for Multispectral, Hyperspectral, and Ultraspectral Imagery VII*, *Proceedings of SPIE 4381*, pp. 516-526.

Mitchell, H. J., Jellison, G. P., Miller, D. P., & Salvaggio, C. (2003). The Importance of Background in the Detection and Identification of Gas Plumes Using Emissive Infrared Hyperspectral Sensing. *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery IX*, *Proceedings of SPIE 5093*, pp. 206-217.

Patterson, J. D., & Bailey, B. C. (2007). *Solid-State Physics: Introduction to the Theory.* Berlin: Springer.

Polak, M. L., Hall, J. L., & Herr, K. C. (20 August 1995). Passive Fourier-Transform Infrared Spectroscopy of Chemical Plumes: An Algorithm for Quantitative Interpretation and Real-Time Background Removal. *Applied Optics , 34* (24), 5406-5412.

RCA. (1974). *Electro-Optics Handbook* (Vol. 11). Harrison, NJ: RCA Commercial Engineering.

Schaum, A. (2001). Spectral Subspace Matched Filtering. *Proceedings of SPIE.*

Schott, J. R. (2007). *Remote Sensing: The Image Chain Approach.* New York: Oxford University Press.

Sharpe, S. W., & Chu, P. M. (2001). Creation of 0.10 cm-1 Resolution, Quantitative, Infrared Spectral Libraries for Gas Samples. *Proceedings of SPIE , 4577*.

Sharpe, S. W., Sams, R. L., & Johnson, T. (2002). The PNNL Quantitative IR Database for Infrared Remote Sensing and Hyperspectral Imaging. *Proceedings of the 31st Applied Imagery Pattern Recognition Workshop.* IEEE.

Sheen, D. M., Gallagher, N. B., Sharpe, S. W., Anderson, K. K., & Schultz, J. F. (2003). Impact of Background and Atmospheric Variability on Infrared Hyperspectral Chemcial Detection Sensitivity. *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery IX*, *Proceedings of SPIE 5093*, pp. 218-229.

Shepherd, G. G. (2002). *Spectral Imaging of the Atmopshere.* London: Elsevier Science Ltd.

Silverstein, R. M., & Webster, F. X. (1998). *Spectrometric Identification of Organic Compounds* (6 ed.). New York: John Wiley & Sons, Inc.

Skoog, D., Holler, F., & Nieman, T. (1998). *Principles of Instrumental Analysis* (5 ed.). Philidalphia: Harcourt Brace & Co.

Smith, B. C. (1996). *Fundamentals of Fourier Transform Infrared Spectroscopy.* Boca Raton, FL: CRC Press.

Thomas, M. J., Lewis, P. E., Kroutil, R. T., & Combs, R. J. (2002). Infrared Detection and Analysis of Vapor Plumes Using Airborne Sensor. *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery VIII*, *Proceedings of SPIE 4725*, pp. 47-64.

Tonooka, H. (2001). An Atmospheric Correction Algorithm for Thermal Infrared Multispectral Data Over Land -- A Water Vapor Scaling Method. *IEEE Transactions on Geoscience and Remote Sensing , 39* (3), 682-692.

Vallieres, A., Chamberland, M., & Farley, V. (2005). High-Performance Field Portable Imaging Radiometric Spectrometer Technology for Chemical Agent Detection. *Optically Based Biological and Chemical Sensing, and Optically Based Materials for Defense*, *Proceedings of SPIE 5990*, pp. 1-9.

Vallieres, A., Villemaire, A., Chamberland, M., & Belhumeur, L. (2005). Algorithms for Chemical Detection, Identification and Quantification for Thermal Hyperspectral Imagers. *Chemical and Biological Standoff Detection III*, *Proceedings of SPIE 5995*, pp. 1-11.

Villemaire, A., Farley, V., Giroux, J., & Chamberland, M. (2004). Development and Testing of a Hyperspectral Imaging Instrument for Standoff Chemical Detection. *Chemical and Biological Standoff Detection II*, *Proceedings of SPIE 5584*, pp. 135-143.

Wadsworth, W., Dybwad, J. P., & Stone, D. (2005). Mosaic Imaging Fourier Transform Spectrometer Applications. *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XI*, *Proceedings of SPIE 5806*, pp. 255-266.

Williams, D. J., Feldman, B. L., & Williams, T. J. (2005). Detection and Identification of Toxic Air Pollutants Using Airborne LWIR Hyperspectral Imaging. *Multispectral and Hyperspectral Remote Sensing Instruments and Applications II*, *Proceedings of SPIE 5655*, pp. 134-141.

Wozniak, B., & Dera, J. (2007). *Light Absorption and Absorbents in Sea Waters.* Poland: Springer.

Ynghua, F., Yanli, Q., & Ye, H. (2005). Airborne Detection of Urban Environmental Pollutants. *Optical Technologies for Atmospheric, Ocean, and Environmental Studies*, *Proceedings of SPIE 5832*, pp. 279-282.

Zachor, A. S. (1977). Drive nonlinearities: their effects in Fourier spectroscopy. *Applied Optics* , 1412-1424.