

Subpixel Temperature Estimation from Single-Band Thermal Infrared Imagery

by

Sarah E. Paul

B.S. Rochester Institute of Technology, 2007

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
Rochester Institute of Technology

August 13, 2012

Signature of the 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 Sarah E. Paul
has been examined and approved by the
dissertation committee as satisfactory for the
dissertation required for the
Ph.D. degree in Imaging Science

Dr. Carl Salvaggio, Dissertation Advisor

Dr. John R. Schott

Dr. Michael G. Gartley

Dr. Steven M. LaLonde

Date

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

Title of Dissertation:

**Subpixel Temperature Estimation from Single-Band Thermal Infrared
Imagery**

I, Sarah E. Paul, 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 _____

Subpixel Temperature Estimation from Single-Band Thermal Infrared Imagery

by

Sarah E. Paul

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

Target temperature estimation from thermal infrared (TIR) imagery is a complex task that becomes increasingly more difficult as the target size approaches the size of a projected pixel. At that point the assumption of pixel homogeneity is invalid as the radiance value recorded at the sensor is the result of energy contributions from the target material and any other background material that falls within a pixel boundary. More often than not, thermal infrared pixels are heterogeneous and therefore subpixel temperature extraction becomes an important capability. Typical subpixel estimation approaches make use of multispectral or hyperspectral sensors. These technologies are expensive and multispectral or hyperspectral thermal imagery might not be readily available for a target of interest.

A methodology was developed to retrieve the temperature of an object that is smaller than a projected pixel of a single-band TIR image using physics-based modeling. Physics-based refers to the utilization of the Multi-Service Electro-optic Signature (MuSES) heat transfer model, the MODerate spectral resolution atmospheric TRANsmission (MODTRAN) atmospheric propagation algorithm, and the Digital Imaging and Remote Sensing Image Generation (DIRSIG) synthetic image generation model to reproduce a collected thermal image under a number of user-supplied

conditions. A target space is created and searched to determine the temperature of the subpixel target of interest from a collected TIR image.

The methodology was tested by applying it to single-band thermal imagery collected during an airborne campaign. The emissivity of the targets of interest ranged from 0.02 to 0.91 and the temperature extraction error for the high emissivity targets were similar to the temperature extraction errors found in published papers that employed multi-band techniques.

Acknowledgements

I think I need to begin by acknowledging that I never thought I'd be able to finish this. After almost 5 years of working toward one goal, I can't believe that I'm done. I didn't get to this point on my own, so here's a list of people without whom I would never have finished.

My Committee

Dr. Carl Salvaggio: Carl, you've been my advisor for 8 years. I imagine there have been times where you thought back to the day when I announced in your undergrad programming class that you were my advisor and wished that you would've pawned me off to another faculty member (and frankly, I wouldn't have blamed you), but I'm eternally grateful that you stuck with me. Through the years you've acted as so much more than an academic advisor and have always been there both professionally and personally.

Dr. John Schott: Your reputation precedes you, making you all the more intimidating to your new students (and possibly faculty members). However, over the last couple years, I've found out how deeply you care about your students and I've always appreciated how you were always willing to give me help - even when I just popped into your office. I always appreciated that, even if the help came with a "That was pretty dumb, kid".

Dr. Mike Gartley: Mike, on behalf of all graduate students, I'd like to thank you for always having an answer (and usually a handy piece of code you wrote in 5 minutes but would've taken us an entire day) to our most puzzling problems. You always have really practical advice and knowledge and the thing that I really appreciated was how you still remember how sucky grad school can be sometimes. It's almost like you're one of us, but way, way smarter. Someday I hope to be half as smart as you are.

Dr. Steve LaLonde: The day I walked into your office to ask you to be on my committee I came armed with all sorts of books and facts about my research, expecting you to grill me. I nervously asked you if you wanted me to tell you all about what I was doing and you basically told me that you were there to make sure the process was followed and you'd

leave the grilling to the people from my department. I knew I'd picked the right person. Thank you so much for agreeing to be on my committee and for always making yourself available when I needed to meet with you.

Supporting Cast

Adam Goodenough, Scott Brown, Andy Scott, and Niek Sanders (aka the DIRSIG guys): I literally would not have been able to complete my research without you guys. I'll never forget the time Adam completely re-wrote the sampling mechanism in the adaptive sampling code almost overnight because it kept acting up. That's pretty much the definition of dedication. You guys were always there with a fix or a trick up your sleeve when I needed it the most and though there were times DIRSIG made me want to throw my computer out the window, you guys were always awesome.

Paul Mezzanini and the rest of the Research Computing staff: I spent so many hours in your office while you debugged my terrible scripts (or as you say "got rid of the badness"). There were a few times that you somehow saved the day during what I thought was the 11th hour so I can't thank you enough for all that you did for me. Also, I can't remember which one of you taught me about the "screen" command in Unix, but it's proven to be ridiculously useful and has even made me look smart in front of others when I can pull that out of my arsenal.

Jason Faulring: I'm not sure why you kept letting me in your lab knowing how handy I was with tools. I'll never forget the day you laughed in my face when I spray painted my left hand black. You helped me so much along the way and I'm so grateful for that. No matter if we were in Midland during February, the poo plant at 3am on a summer day, or the RIT campus late on a Friday night you were there with a smile on your face, which made those undesirable situations much more bearable.

Steven Patterson: Thank you so much for all of the MuSES help you provided to me. You were so responsive to my questions and even had a screen sharing session where you debugged one of my major problems. Your help was invaluable to me.

Cindy Schultz and Amanda Zeluff: Thanks to both of you for helping me to keep my grad school life in order. I could always count on you guys to tell it like it is and that's one of the things I enjoyed the most about our conversations.

Nanette Salvaggio: Thank you for being there for me when my real mom couldn't. You were there for me during a time in my life when I needed a friend and I'll never forget your kindness and your great advice.

Alvin: You've been my best friend and also my largest source of frustration, sometimes on the same day. You were always one of my biggest supporters and always helped me whenever I needed it. That meant more to me than you'll ever know.

Justin and Dave: Justin...where do I even begin? You ended up being the brother I never had (and sometimes never wanted, hehe). Thank you for showing me how to loosen up and have fun. I have so many ridiculous memories with you and I cherish each of them. Except for the time you left our apartment unlocked and drunk people came in and stole all of our food. Closest I ever came to punching you. :o) Dave, you were a fantastic officemate. Thank you for not Nilosekking my research or my computer. In all seriousness thanks for all the help you gave me (especially with Matlab) and for being such a good sport...especially the time I made you stand in a swampy hole at 3am at the poo plant where you were subsequently attacked by bats.

May, Cara, and Autie: You three have made everything from getting gas to watching a crappy 'tween vampire movie fun. May, I couldn't have made it through grad school without you. I literally couldn't have asked for three better friends. Thank you for making me laugh and for being your awesome selves.

This has gotten long enough, so I'll finish by thanking all of my family and the rest of the faculty and graduate students. I learned something from each of you and I couldn't have made it all the way to the goal line without you.

To Mom, Nan, Pa, and Cliff.
Thanks for always believing in me, even when I didn't believe in myself.

Contents

1	Introduction	1
1.1	Objectives	2
1.2	Assumptions	2
2	Background	5
2.1	Electromagnetic Radiation	6
2.2	Radiometry	7
2.2.1	Blackbody Radiation	7
2.2.2	Emissivity	9
2.2.3	Conservation of Energy	10
2.2.4	Energy Paths to the Sensor	11
2.2.5	Material Self-Emitted Radiance	12
2.2.6	Reflected Radiance	12
2.2.6.1	Shape Factor	12
2.2.6.2	Reflected Downwelled Radiance	13
2.2.6.3	Background Reflected Radiance	15
2.2.7	Upwelled Radiance	15
2.2.8	Atmospheric Transmission	16
2.2.8.1	Absorption	16
2.2.8.2	Scattering	17
2.2.8.3	Total Transmission	18
2.2.9	The Governing Equation	19

2.3	Mixed Pixels	19
2.3.1	Target position	20
2.4	Heat Transfer	21
2.5	Point Spread Function	22
2.6	Modeling Tools	22
2.6.1	DIRSIG	23
2.6.1.1	Adaptive Subsampling	24
2.6.2	MuSES	25
2.6.3	MODTRAN	27
2.7	Regression	28
2.7.1	Least Squares	30
2.7.2	Model Adequacy	31
2.7.2.1	Residual Analysis	31
2.7.2.2	Testing Individual Coefficients for Significance	32
2.7.3	Sample size selection	34
2.8	Error propagation	34
2.9	Temperature extraction from multiple pixels	35
3	Previous Work	37
3.1	Subpixel Temperature Retrieval with Multispectral Sensors	37
3.2	A Method for Satellite Identification of Surface Temperature Fields of Subpixel Resolution	39
3.2.1	Retrieval of Subpixel Fire Temperature and Fire Area in Mod- erate Resolution Imaging Spectrometer	41
3.2.2	Application of the Dozier Retrieval to Wildfire Characteriza- tion: A Sensitivity Analysis	41
3.3	Estimating Subpixel Surface Temperatures and Energy Fluxes from the Vegetation Index-Radiometric Temperature Relationship	41
3.4	A Novel Method to Estimate Subpixel Temperature by Fusing Solar- Reflective and Thermal-Infrared Remote-Sensing Data with an Arti- ficial Neural Network	43

3.5	Sub-pixel Water Temperature Estimation from Thermal-Infrared Imagery Using Vectorized Lake Features	45
3.6	An Image-Sharpening Method to Recover Stream Temperatures from ASTER Images	46
3.7	Subpixel Temperature Estimation from Low Resolution Thermal Infrared Remote Sensing	48
3.8	Previous Work Summary	50
4	Methodology	52
4.1	Overview	52
4.2	Parameter estimation	53
4.2.1	Target temperature	54
4.2.2	Emissivity data	54
4.2.3	Background temperature	55
4.2.4	Target size	55
4.2.5	Target position	55
4.2.5.1	Comparison metric for target position determination	56
4.2.5.2	Implementation	57
4.2.6	Weather data	62
4.3	Geometry	63
4.4	Physics-based simulations	64
4.4.1	MuSES simulations	64
4.4.1.1	Attribution of material properties	65
4.4.1.2	Weather files	66
4.4.1.3	Running MuSES	69
4.4.2	MODTRAN simulations	70
4.4.3	DIRSIG simulations	71
4.4.3.1	Setting up the DIRSIG scene	72
4.4.3.2	Running the simulations	72
4.5	Building the lookup table	73
4.6	Regression	75

4.7	Testing the method	77
4.7.1	Airborne collection	77
4.7.1.1	Targets	78
4.7.1.2	Location	79
4.7.1.3	Test plan	80
4.7.2	Atmospheric compensation	85
4.7.3	Modeling in MuSES	86
4.8	Error analysis	88
4.9	Synthetic Experiment	93
4.10	MuSES Utility Study	95
5	Results	98
5.1	Airborne collect results	98
5.1.1	Airborne data	98
5.1.2	Ground truth	101
5.1.3	Background temperature estimation	103
5.1.4	Target temperature extraction results	104
5.1.5	Error analysis	113
5.1.6	Lessons learned	116
5.2	Synthetic imagery study results	118
5.3	Heat transfer utility study results	122
6	Conclusion	126
A	Heat Transfer Study	134
B	Novel contributions to the field	136
C	Suggestions for future work	137

Chapter 1

Introduction

Target temperature estimation from thermal infrared (TIR) imagery is a complex task that becomes increasingly more difficult as the target size approaches the size of a projected pixel. At that point the assumption of pixel homogeneity is invalid as the radiance value recorded at the sensor is the result of energy contributions from the target material and any other background material that falls within a pixel boundary. More often than not, thermal infrared pixels are heterogeneous and therefore subpixel temperature extraction becomes an important capability. A typical problem would involve a small, hot target, since most objects of interest that exist at elevated temperatures are small dimensioned. Running engines or motors, exhaust plumes from processes at manufacturing plants, electrical transformers, individual light sources, and hot water discharges from power plants are just a few examples of these types of targets, all of which are small relative to the footprint of the typical remote sensing systems that are used to monitor them.

A methodology has been developed to retrieve the temperature of an object with a uniform temperature that is smaller than a projected pixel of a single-band TIR image using physics-based modeling. The process is broken into two distinct pieces. In the first part, the Multi-Service Electro-optic Signature (MuSES) thermal signature program in conjunction with the Digital Imaging and Remote Sensing Image

Generation (DIRSIG) tool will be used to replicate a collected TIR image based on parameter estimates from the collected image as well as companion high resolution visible imagery of the target. This is done many times to build a multi-dimensional lookup table (LUT). For the second part, a regression model is built from the data in the LUT and is used to perform the temperature retrieval.

1.1 Objectives

The objectives of this research are as follows:

1. Develop a methodology to extract the temperature of an object that is smaller than a projected thermal pixel. This object is assumed to have a uniform temperature across its surface.
2. Provide a methodology that accounts for the heat transfer between objects as well as for the radiometric phenomenology in the physics-based modeling.
3. A complete validation and verification of the developed technique will be carried out using modeled data and data collected from a real airborne sensor acquired during a monitored, ground-truthed experiment.
4. Perform an error analysis of the results to quantify the accuracy of the methodology.

1.2 Assumptions

Temperature extraction from an object that is smaller than a pixel is a complicated undertaking. The following assumptions (in no particular order) are made to limit the scope of this research as well as make attainment of the goals tractable. However, it should be noted that these assumptions are realistic for data obtained from most TIR imaging systems.

High-resolution visible imagery of target is available This approach relies heavily on having accurate estimates of input parameters. The optical and thermal properties of the target material as well as the physical size of the target and the position of the target within the projected thermal pixel can be determined from the having visible imagery available that is of a higher spatial resolution than the thermal image. These parameters would be difficult to accurately estimate using the thermal image alone as the target of interest is smaller than the projected pixel.

Geometry of target is available A three-dimensional model of the target of interest is required for this methodology as it is the input to the physics-based models. It is assumed that the geometry is available (*e.g.* from the visible imagery) by some means and that it is of high enough resolution that no error is introduced into the process from geometrical inadequacies.

Thermal sensor is well characterized It is assumed that sensor characteristics (*e.g.* the point spread function and the spectral response function) are known so that the sensor can be modeled in DIRSIG. It is also assumed that the position of the sensor at the time of image capture is known.

Two pixel constituents The phenomenology for a heterogeneous pixel is complex. To simplify matters, it is assumed that only two materials (known as the target and the background) occupy a pixel. It is also assumed that the background material will fill entire pixels surrounding the mixed pixel(s) of interest so that the background temperature can be estimated.

Target of interest has a uniform temperature It is assumed that a large temperature gradient does not exist on the surface of the target.

Atmospheric data available It is assumed that the surface weather conditions are known for a time period preceding the TIR image collection time.

Atmospheric compensation An atmospheric compensation is assumed to have been performed on the thermal imagery. The atmospheric transmission and

upwelled radiance are not considered, however, the reflected downwelled component will be taken into consideration.

Nadir-looking imagery The thermal infrared image of the target of interest has been captured from nadir or near-nadir angles.

With the objectives and assumptions in mind, the remainder of this document details the theory and methodology for a physics-based target space approach to the subpixel temperature extraction problem.

Chapter 2

Background

As stated previously this research uses a physics-based target space approach to the subpixel temperature retrieval problem. Before diving deeper, it seems prudent to discuss what exactly is meant by a physics-based target space. Physics-based refers to the utilization of the Multi-Service Electro-optic Signature (MuSES) heat transfer model, the MODerate spectral resolution atmospheric TRANsmission (MODTRAN) atmospheric propagation algorithm, and the Digital Imaging and Remote Sensing Image Generation (DIRSIG) synthetic image generation model to reproduce a collected thermal image under a number of user-supplied conditions.

A target space is an N -dimensional space comprised of parameters that have a significant effect on sensor-reaching radiance as well as the sensor-reaching radiance due to the characteristics of an individual pixel. In the context of this subpixel temperature retrieval method, the target space has two dimensions: target temperature and pixel radiance. All other parameters can be estimated from the thermal image or the high resolution visible image. The following sections will describe the phenomenology and the tools required to build and search through the physics-based target space.

2.1 Electromagnetic Radiation

Electromagnetic (EM) radiation is a form of energy that propagates through space as a wave with electric and magnetic field components. EM radiation does not require a medium through which to travel and is classified according to the wavelength of the wave. The electromagnetic spectrum spans the entire range of wavelengths of EM radiation, from gamma rays to radio waves, and is shown in Figure 2.1.




Figure 2.1: The electromagnetic spectrum

Traditionally, in remote sensing, the focus has been on the area of the EM spectrum from the visible portion (0.4 microns) to the thermal infrared (14 microns). That region can be further partitioned into five segments shown in Table 2.1. The partitions occur because of phenomenological changes from region to region as well as coincidence with highly transmissive atmospheric windows.

Name	Abbreviation	Wavelengths [μm]
Visible	VIS	0.4 - 0.7
Near Infrared	NIR	0.7 - 1.0
Shortwave Infrared	SWIR	1.0 - 2.5
Midwave Infrared	MWIR	3.0 - 5.0
Longwave Infrared	LWIR	8.0 - 14.0

Table 2.1: Segments of the EM spectrum traditionally focused upon in remote sensing.

The VIS portion of the spectrum contains the narrow band of wavelengths to which the human visual system is sensitive. The NIR and SWIR portions are still in the solar reflective region and can be used to monitor vegetation health as well as further discriminate between material types (*e.g.* a black roof versus a black car). During the day the energy in the MWIR is a result of both solar photons as well as photons that are emitted directly from an object. However, at night when the sun is down, the energy sensed in the MWIR is a result of self-emission. Finally, the energy in the LWIR is solely a result of self-emission as the solar contribution is deemed negligible. The LWIR is also referred to as the *thermal infrared* (TIR). Using MWIR and LWIR imagery one can potentially determine the temperature of objects within an image. However, this research will rely solely on LWIR imagery.

2.2 Radiometry

In order to understand subpixel temperature extraction techniques it is important to understand the origin of the LWIR energy and how it is propagated from the ground to the sensor.

2.2.1 Blackbody Radiation

A blackbody is an idealized surface that can perfectly reradiate all absorbed incident energy. That is to say that the reflectivity of the surface is zero and the absorptivity

is one [1]. The radiance leaving a blackbody can be modeled by Planck's Law, which is defined in Equation 2.1.

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

where:

h	Planck constant	$6.6256 \cdot 10^{-34} [J \cdot s]$
c	speed of light in a vacuum	$2.99792 \cdot 10^8 [m/s]$
k	Boltzmann constant	$1.3807 \cdot 10^{-23} [J/K]$
λ	wavelength	$[\mu m]$
T	absolute temperature	$[K]$

Equation 2.1 depends only on the temperature of the object and the wavelength, or wavelengths, of interest. If the temperature is held constant, a blackbody curve can be generated for a range of wavelengths as shown in Figure 2.2.




Figure 2.2: Blackbody curves for a 6000 K blackbody (top), 1000 K blackbody (middle), and a 300 K blackbody (bottom).

The magnitude of a blackbody curve varies with temperature, with higher temper-

atures producing curves with larger overall magnitudes. The peak of the blackbody curve also shifts as the temperature changes. As the temperature increases the peak value shifts to shorter wavelengths. Wein's displacement law, defined in Equation 2.2, gives the relationship between the temperature of an object, T , and the wavelength of maximum radiance emission.

$$\lambda_{max} = \frac{2897.768[\mu m K]}{T} \quad [\mu m] \quad (2.2)$$

2.2.2 Emissivity

A blackbody is an idealized surface, therefore real objects will emit less energy than a blackbody with the same temperature. Emissivity is a parameter that takes into account the emission efficiency of an object. It is the ratio of the self-emitted spectral radiance, $L(\lambda, T)$ from an object at temperature T to the spectral radiance from a blackbody at the same temperature $L_{BB}(\lambda, T)$

$$\varepsilon(\lambda) = \frac{L(\lambda, T)}{L_{BB}(\lambda, T)} \quad (2.3)$$

Emissivity is a wavelength-dependent unitless parameter that can vary between 0 and 1. For a Lambertian material that emits radiance equally into the hemisphere above its surface, the emissivity is dependent upon wavelength only. Most surfaces are not exactly Lambertian (although approximately Lambertian surfaces can be found in nature [1]) and will radiate more into one part of the hemisphere than another. In other words the amount of radiance observed from a material will be dependent upon the viewing geometry. In this case the emissivity term needs to take into account where the sensor is in relation to the material of interest. This quantity is known as directional emissivity. Table 2.2 shows a summary of the emissivity terms.

Table 2.2: Summary of emissivity terms

Lambertian?	Emissivity Term	Dependencies
Yes	$\varepsilon(\lambda)$	Wavelength (λ)
No	$\varepsilon(\theta, \phi, \lambda)$	Sensor orientation (θ, ϕ), Wavelength (λ)

2.2.3 Conservation of Energy

Radiance incident upon a material can be reflected, transmitted, or absorbed. Reflectivity, transmissivity, and absorptivity are wavelength-dependent material properties that are described by unitless metrics. Reflectivity describes how well a material can send incident electromagnetic energy back into the medium from which it came. The reflectance, $\rho(\lambda)$, is defined as the proportion of reflected energy to the amount of energy incident on the surface. Transmissivity characterizes how well electromagnetic energy is allowed to pass through a medium. The transmission, $\tau(\lambda)$, is defined as the ratio of the energy at the back of a surface to the energy incident on the front of the surface. Absorptivity is the ability of a material to remove electromagnetic energy from the system by converting it to another form of energy [1]. The absorptance, $\alpha(\lambda)$, is the ratio of the energy taken out of the system to the amount of incident energy onto a surface. Conservation of energy dictates that these three quantities sum to one so that all energy in the system is accounted for.

$$\rho(\lambda) + \tau(\lambda) + \alpha(\lambda) = 1 \quad (2.4)$$

For an opaque material the transmission is zero and the conservation of energy equation becomes

$$\rho(\lambda) + \alpha(\lambda) = 1 \quad (2.5)$$

According to Kirchhoff's Law, the absorptivity of a material at thermal equilibrium is equal to the emissivity $\varepsilon(\lambda)$ of the material [2]. Kirchhoff's Law for an opaque material is

$$\rho(\lambda) + \varepsilon(\lambda) = 1 \quad (2.6)$$

2.2.4 Energy Paths to the Sensor

The radiance recorded by a thermal infrared (TIR) sensor is due to energy originating from multiple sources. Unlike the visible portion of the spectrum, the energy from an object in the TIR is a result of it having a temperature above absolute zero. Figure 2.3 illustrates the four major energy paths to the sensor in the thermal infrared.




Figure 2.3: Energy paths in the thermal infrared. (Adopted from Schott, 2007, pg. 62 [1])

Path A represents the self-emitted radiance from a target of interest. Path B is the reflected downwelled radiance from the atmosphere. The atmospheric self-emission that is propagated directly to the sensor is represented by path C, and the self-emitted radiance from nearby objects that reflects off the target material is path D. These energy paths will be described in more detail in the following sections.

2.2.5 Material Self-Emitted Radiance

The self-emitted radiance from the material of interest (energy path A in Figure 2.3) is the only path to the sensor that contains information about the temperature of the material [1]. To produce the self-emitted radiance, the emissivity modifies the blackbody radiance by serving as a wavelength-dependent multiplicative factor.

$$L(\lambda, T) = \varepsilon(\lambda) L_{BB}(\lambda, T) \left[\frac{W}{m^2 sr \mu m} \right] \quad (2.7)$$

Equation 2.7 gives the radiance from a Lambertian material solely due to its temperature. The self-emission for a non-Lambertian surface is

$$L(\theta, \phi', \lambda, T) = \varepsilon(\theta, \phi, \lambda) L_{BB}(\lambda, T) \left[\frac{W}{m^2 sr \mu m} \right] \quad (2.8)$$

where (θ, ϕ') are the zenith and azimuth angles, respectively, describing the direction of the sensor.

2.2.6 Reflected Radiance

Equation 2.6 can be rearranged to show that when the emissivity of an opaque object is less than one, some fraction of incident energy is reflected from its surface. Energy paths B and D in Figure 2.3 represent radiance from non-target sources being reflected off the target toward the sensor.

2.2.6.1 Shape Factor

It is rarely the case where an object has a totally unobstructed view of the sky and/or is perfectly horizontal. Most of the time the skydome above the object is blocked and/or the object is tilted at some angle (*e.g.* a sloped roof). This reduces the amount of downwelled radiance onto the object as well as introduces energy into the system from the object obscuring the sky. The fraction of the skydome not masked by the background can be quantified by the shape factor, F . To compute the

shape factor, the fraction of the hemisphere that is obstructed by the background is computed and then subtracted from 1. The fraction of the hemisphere blocked by the background is computed by integrating the solid angle element ($d\Omega$) over the solid angle obstructed by the background and dividing that quantity by the solid angle of the hemisphere [1].

$$F = 1 - \frac{\int d\Omega}{2\pi} \quad (2.9)$$

It follows that the fraction of skydome that is masked by the background is given by $(1-F)$. The shape factor is an important concept when talking about the two reflected energy paths to the sensor: reflected downwelled (path B) and reflected background radiance (path D).

2.2.6.2 Reflected Downwelled Radiance

As stated previously the energy in path B is the energy that the atmosphere emits because of its temperature. The question becomes one of how much of that energy is reflected off of the target toward the sensor. To answer this, the energy contribution of every location in the sky must be determined. This is done by approximating the non-homogeneous atmosphere as a series of N homogeneous layers stacked on each other where the first layer is defined to be the one closest to the ground. The blackbody radiance of each layer can be calculated using Planck's Law. However, the atmosphere is not a blackbody so (as in the case of the target self-emitted radiance) the emissivity of the atmosphere will serve as a wavelength-dependent modifier of the blackbody radiance of each layer. To determine the emissivity of each layer recall that conservation of energy requires that the reflectance, transmission, and absorptance of a medium equal one. This can be also applied to each layer of the atmosphere

$$\Delta\rho_i + \Delta\tau_i + \Delta\alpha_i = 1 \quad (2.10)$$

where $\Delta\rho_i$, $\Delta\tau_i$, and $\Delta\alpha_i$ are the reflectance, transmission, and absorptance of

the i th layer of the atmosphere. Since atmospheric scattering is negligible in most portions of the thermal infrared [1], $\Delta\rho$ is zero leaving

$$\Delta\tau_i + \Delta\alpha_i = 1 \quad (2.11)$$

Kirchhoff's Law can be applied by replacing the absorption term with emissivity and the terms can be rearranged to show that the emissivity of the entire i th layer can be expressed as

$$\Delta\varepsilon_i = 1 - \Delta\tau_i \quad (2.12)$$

At this point the self-emitted radiance of a single layer of the atmosphere can be calculated. However, the energy from each layer needs to be propagated down to the target of interest. This is done by multiplying the radiance term of the i th layer by the product of the transmissions of the $(i - 1)$ layers that came before. To obtain the total downwelled radiance contribution from the atmosphere in the direction defined by the angles (σ, ϕ) , the radiance from each of the N atmospheric layers is summed giving

$$L_d(\sigma, \phi) = \sum_{i=1}^N L_{atm.i} (1 - \Delta\tau_i) \prod_{j=1}^{i-1} \Delta\tau_j \quad (2.13)$$

If the spectral dependence is included, the amount of radiance reflected toward the sensor from a horizontal, Lambertian surface with an unobscured view of the sky is

$$L_{rd}(\lambda) = L_d(\lambda) r(\lambda) \quad (2.14)$$

If the surface is non-Lambertian (but still has an unobscured view of the sky), the reflected downwelled radiance has an angular dependence as well as a spectral dependence.

$$L_{rd}(\theta, \phi', \lambda) = \int_{2\pi} L_d(\sigma, \phi) \cos\sigma \frac{r(\sigma, \phi, \theta, \phi', \lambda)}{\pi} d\Omega \quad (2.15)$$

where (θ, ϕ') represent the direction of the sensor.

For surfaces whose view of the sky is obstructed, the shape factor, F , from Section 2.2.6.1 is included in the limit of integration. The equation for the amount of radiance reflected toward the sensor due to the downwelled component of self-emitted sky radiance is

$$L_{rd}(\theta, \phi', \lambda) = \int_F L_d(\sigma, \phi) \cos\sigma \frac{r(\sigma, \phi, \theta, \phi', \lambda)}{\pi} d\Omega \quad (2.16)$$

2.2.6.3 Background Reflected Radiance

The equation for the amount of radiance from background objects that is reflected off of the target toward the sensor, L_{rbk} , is derived in the same way as the reflected downwelled radiance. However, in this case the integration is performed over the amount of sky obstructed by the background (1-F) and the amount of radiance from the background, L_{bk} , is used.

$$L_{rbk}(\theta, \phi', \lambda) = \int_{1-F} L_{bk}(\sigma, \phi) \cos\sigma \frac{r(\sigma, \phi, \theta, \phi', \lambda)}{\pi} d\Omega \quad (2.17)$$

Equation 2.16 and Equation 2.17 can be combined for the sake of simplicity into a reflected radiance term

$$L_r(\theta, \phi', \lambda) = \int_F L_d(\sigma, \phi) \cos\sigma \frac{r(\sigma, \phi, \theta, \phi', \lambda)}{\pi} d\Omega + \int_{1-F} L_{bk}(\sigma, \phi) \cos\sigma \frac{r(\sigma, \phi, \theta, \phi', \lambda)}{\pi} d\Omega \quad (2.18)$$

2.2.7 Upwelled Radiance

The upwelled component of the sensor-reaching radiance (path C in Figure 2.3) is also due to the self-emission of the atmosphere. The energy is emitted by the at-

mosphere directly into the line of sight of the sensor, as opposed to the downwelled component, where the energy is reflected into the sensor line of sight by scene elements. The expression for the upwelled radiance is derived in the same manner as the downwelled radiance except that the layers of the atmosphere are defined with the first layer being at the sensor and the last at the ground. Therefore the transmission term, τ_i , is the transmission of the atmosphere from the sensor to the top of the i th layer. Finally, the upwelled radiance is expressed as

$$L_u(\theta, \phi') = \sum_{i=1}^N L_{BB}(\lambda, T_i)(1 - \Delta\tau_i) \prod_{j=1}^{i-1} \Delta\tau_j \quad (2.19)$$

2.2.8 Atmospheric Transmission

Atmospheric transmission refers to the loss of energy in a beam of radiation as it is propagated from the object of interest to the sensor. The mechanisms by which energy is lost due to the atmosphere occurs are atmospheric absorption and scattering.

2.2.8.1 Absorption

In general, absorption occurs when energy in a beam is converted to another form of electromagnetic energy (typically thermal). The converted energy is therefore considered a loss from the original beam. Atmospheric absorption occurs when molecules in the atmosphere absorb energy from a beam, which causes a molecular vibration, rotation, or electron orbital transition to another energy state.

The amount of energy removed in a given path length depends on the effective size of the absorbing molecules as well as the number of molecules in a volume. The absorption cross section, C_α , is the effective size of a molecule relative to the photon flux at that wavelength [1]. The absorption cross section can be computed as a function of wavelength for an atmospheric constituent. The fractional amount of energy absorbed per unit length of travel is given by

$$\beta_\alpha = mC_\alpha \quad [\text{m}^{-1}] \quad (2.20)$$

where m is the number of molecules per unit volume. For a beam passing through a horizontal layer of gas with an infinitesimally small thickness, dz , the absorption through the layer is

$$d\alpha = \frac{d\Phi}{\Phi} = -\beta_\alpha(z) \sec(\theta) dz \quad (2.21)$$

where θ is the zenith angle, Φ is the incident flux entering the layer, and $d\Phi$ is the flux lost to absorption in the layer. Integrating both sides and taking the antilog gives the transmission,

$$\tau_\alpha = e^{-\int_{z_1}^{z_2} \beta_\alpha(z) \sec \theta dz} \quad (2.22)$$

where the limits of integration are the starting and ending altitudes of the path. If more than one constituent is contained within the layer, the absorption coefficient becomes the sum of the absorption coefficients of the individual constituents.

2.2.8.2 Scattering

Energy can be lost from a beam of radiation via scattering. The transmission loss due to scattering can be written in a similar fashion to the transmission loss due to absorption

$$\tau_s = e^{-\int_{z_1}^{z_2} \beta_s(z) \sec \phi dz} \quad (2.23)$$

where $\beta_s(z)$ is the scattering coefficient. The scattering coefficient depends on the cross-sectional area of the particles as seen by the beam as well as their index of refraction. It can be expressed as a function of a dimensionless size parameter a ,

$$a = \frac{2\pi r}{\lambda} \quad (2.24)$$

The a value is the ratio of the size of the scattering particle to the wavelength of the incident radiation [3]. Rayleigh scattering encompasses the area where the particle size is much smaller than the wavelength ($a \ll 1$). Lord Rayleigh showed that for a given index of refraction the scattering coefficient is inversely proportional to the fourth power of wavelength. Incident radiation is scattered evenly in the forward and backward hemispheres [1].

The Mie scattering region spans the intermediate values of a ($0.1 < a < 50$) [3] and is often referred to as “aerosol scattering”. Mie scattering is weakly dependent on the wavelength of the radiation and, unlike Rayleigh scattering, the radiation is highly forward-scattered [1].

When the particle size is much larger than the wavelength ($a \gg 50$) nonselective scattering occurs. With nonselective scattering the incident radiation is scattered approximately equally in all directions. There is also little dependency on the wavelength of incident radiation. Nonselective scattering is most often associated with raindrops, hail, and large dust particles, for example.

In the thermal infrared, significant scattering is not produced by ordinary gas molecules. However scattering will become important when dust particles and aerosols are present in the atmosphere [3].

2.2.8.3 Total Transmission

Atmospheric absorption and scattering can be taken into account by an overall atmospheric transmission term

$$\tau_{atm} = e^{-\int_{z_1}^{z_2} \beta_{ext}(z) \sec \phi \, dz} \quad (2.25)$$

where β_{ext} is the extinction coefficient that represents the sum of the individual absorption and scattering coefficients.

2.2.9 The Governing Equation

The terms described in Sections 2.2.5 - 2.2.8 can be combined into a single equation that describes how energy is propagated from ground to sensor in the thermal infrared portion of the spectrum. This is given by the *governing equation*.

$$L(\theta, \phi', \lambda) = [L_{self} + L_r] \tau_{atm} + L_u \quad (2.26)$$

The governing equation gives the total spectral radiance reaching the sensor in units of $[W m^{-2} sr^{-1} \mu m^{-1}]$.

2.3 Mixed Pixels

Mixed pixels occur when two or more materials are contained within a pixel boundary. The governing equation discussed in Section 2.2.9 models energy propagation to the sensor for a pixel comprised of a single material. The nature of the subpixel problem requires the consideration of another material within a pixel and therefore another governing equation for heterogeneous pixels. Equation 2.27 shows the governing equation for a pixel with two constituents that are expressed as target and background (note that the angular dependencies are left off for brevity).

$$L_{sensor}(\lambda) = [f_t \varepsilon_t(\lambda) L_{BB}(T_t, \lambda) + (1 - f_t) \varepsilon_b(\lambda) L_{BB}(T_b, \lambda) + f_t L_{rt}(\lambda) + (1 - f_t) L_{rb}(\lambda)] \tau(\lambda) + L_u(\lambda) \quad (2.27)$$

where f_t is the areal fraction of the target. The reflected radiance terms are denoted by L_{rt} and L_{rb} and account for the reflection of downwelled and background radiance from the target and background materials, respectively.

2.3.1 Target position

One important parameter to the small target problem that the governing equation does not take into account is the position of the target relative to the projected thermal pixels. Figure 2.4 illustrates what is meant by the target position parameter.

The target in the leftmost image in Figure 2.4 is located such that each of the four pixels that it is physically projected into will see an equal radiance contribution and therefore they will all have the same radiance. The center image has the target centered in the lower right pixel which will cause the lower right pixel to have a different radiance than the other pixels in the neighborhood. However, due to adjacency effects and the point spread function of the imaging system there will still be some contribution of energy from the target in the surrounding pixels. The same is true for the image on the right.




Figure 2.4: The location of a subpixel target (red circle) changes the radiance values for the pixels in the neighborhood surrounding the target.

In short, pixel radiance values vary with target position. It is important to have a good estimate of the location of the target when modeling a scene in DIRSIG or else there will be error in the pixel radiance values used to derive the target temperature at the end of the process which will in turn cause a larger temperature extraction error.

2.4 Heat Transfer

Heat transfer is thermal energy in transit due to a spatial temperature difference [4]. There are three modes of heat transfer: conduction, convection, and thermal radiation. *Conduction* is the transfer of energy from more energetic atoms or molecules to less energetic atoms or molecules through collisions. The more energy a particle has, the higher the temperature, so energy via conduction must occur in the direction of decreasing temperature. The conductive heat flux, $q''_{conduct}$, is the rate of heat transfer per unit area for conduction and is governed by Fourier's Law. In vector form the heat flux due to conduction is

$$q''_{conduct} = -k \left(\mathbf{i} \frac{dT}{dx} + \mathbf{j} \frac{dT}{dy} + \mathbf{k} \frac{dT}{dz} \right) \quad (2.28)$$

where k is the thermal conductivity in units of $[Wm^{-1}K^{-1}]$ and $\frac{dT}{dn}$ is the change in temperature over a distance in the n direction. The negative sign indicates the heat is being transferred in the direction of decreasing temperature.

Convection occurs between a fluid in motion and a surface when the two surfaces are different temperatures [4]. Convection can be due to force (*e.g.* HVAC systems) or nature (*e.g.* oceanic currents). The convective heat flux is given by Newton's Law of Cooling

$$q''_{convec} = h(T_s - T_\infty) \quad (2.29)$$

where h is the convective heat transfer coefficient, T_s is the temperature of the surface, and T_∞ is the temperature of the free stream. The convective heat transfer coefficient is a measure of how effective a fluid is at carrying heat to and away from a surface. The higher the convection coefficient, the more effective the fluid is at transporting the heat energy. It is dependent on factors such as the fluid density, velocity, and viscosity. In general, fluids that are dense and/or have a higher velocity have a larger value for h . The same is true for rougher surfaces [5].

Thermal radiation is the exchange of heat energy through electromagnetic propagation and requires no matter. The net rate of radiation heat transfer from a surface is

$$q''_{rad} = \varepsilon\sigma (T_s^4 - T_{sur}^4) \quad (2.30)$$

2.5 Point Spread Function

The point spread function (PSF) is a characteristic of all optical systems and is the measure how the energy from a point source is spread spatially across the focal plane. The degraded image, $g(x,y)$, can be modeled as a convolution

$$g(x,y) = f(x,y) * h(x,y) \quad (2.31)$$

where $f(x,y)$ is the incoming energy, and $h(x,y)$ is the point spread function. The $h(x,y)$ term is also called the impulse response of the system. It should be noted that while the optics have an impulse response, other steps in the imaging process (*e.g.* the detector) have their own impulse responses, each of which can be convolved together to arrive at the total impulse response of the system.

2.6 Modeling Tools


Three modeling tools will be used extensively in the completion of this research. The first is the Digital Imaging and Remote Sensing Image Generation (DIRSIG) tool [6]. DIRSIG will be used to model the collected thermal image. In conjunction with DIRSIG, the MuSES (Multi-Service Electro-optic Signature) program [7] will be used to model the heat transfer between objects in a scene. The third tool is the MODerate spectral resolution atmospheric TRANsmission (MODTRAN) atmospheric radiative transfer program. MODTRAN [8] will be used to estimate the downwelled component of the radiance as well as the atmospheric transmission and the upwelled radiance.

2.6.1 DIRSIG

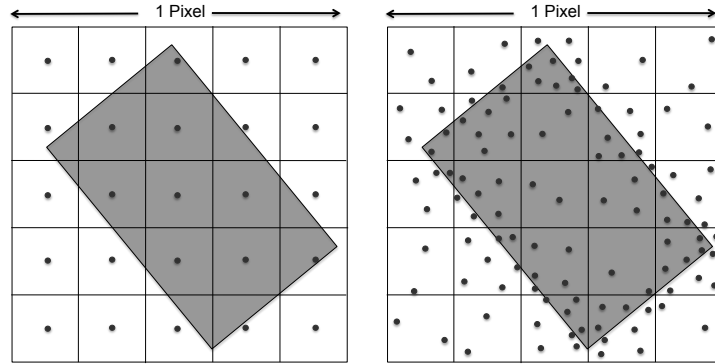
The DIRSIG tool was developed by the Digital Imaging and Remote Sensing (DIRS) Laboratory at the Rochester Institute of Technology (RIT). It is a synthetic image generation application that is capable of modeling imaging systems with sensitivity in the visible through thermal infrared regions of the spectrum as well as polarimetric, RADAR, and LiDAR systems [6]. To accurately reproduce the radiometry of a scene, DIRSIG utilizes ray tracing and first principle physics and mathematical theories. Sub-models based on first principles, including BRDF prediction, facet temperature prediction, sensor models, and atmospheric models are used to generate the synthetic imagery. 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. The DIRSIG model can also model multiple reflections as a photon interacts with a scene. The standard output of DIRSIG is an image file with an associated header. Truth images containing the data used by DIRSIG during simulations (*e.g.* material ID) can also be rendered.

DIRSIG is an appropriate choice for a modeling tool for this research for multiple reasons. One is that it is able to model radiometric interactions between facets of a model as well as pieces of geometry in a scene without making any simplifying assumptions. It can also model (with the help of MODTRAN) the interaction between the atmosphere and the scene. It was also chosen because of the addition of an improved subpixel sampling scheme that would facilitate the subpixel problem present in this research.

2.6.1.1 Adaptive Subsampling



(a) Single ray sampling shown for 4x4 pixels.



(b) Grid oversampling shown for one pixel.

(c) Adaptive subsampling shown for one pixel.

Figure 2.5: Sampling methods in DIRSIG. Each dot represents a single sampling ray.

Typically in DIRSIG, one sampling ray is cast from the center of each detector element and calculates the radiometry for the entire pixel based upon the material the ray intersects with. This single ray sampling method, shown in Figure 2.5(a), will return the correct answer for a material that completely fills a pixel.

Prior to the start of this project a sampling method that accounted for mixed pix-

els was already implemented in DIRSIG. This method, shown in Figure 2.5(b), uses a grid sampling approach where each pixel in the focal plane is divided into sub-elements and a sampling ray is cast from the center of each sub-element. The sampling rays from one pixel are averaged together to produce the pixel's radiance value. The sub-samples in grid sampling are spread evenly over the pixel and therefore spatially sample all sub-elements equally. An improvement to this method would be to send more samples to the important areas of a scene (*e.g.* areas of change) as opposed to areas of less importance (*e.g.* large, uniform areas).

Adaptive subsampling, while not new to the computer graphics community, is a new sampling scheme in DIRSIG introduced for this research that seeks to only send sub-pixel samples to areas of importance. This method, shown in Figure 2.5(c), divides each pixel into a user-specified number of sub-elements and a sampling ray is cast from a random location within the pixel. Adaptive subsampling looks for variation within and between sub-elements and sends more sampling rays where it finds variation. There are thresholds that the user can specify so that the code does not keep casting sampling rays *ad nauseam*. One is the maximum number of sampling rays cast per pixel and the other is a change in radiance threshold, where, if the radiance value computed from one sampling ray to the next changes by less than the threshold, no more rays will be cast. If either of these stopping criteria are met the adaptive sampling code ceases casting rays and moves on to the next pixel.

This improvement in the way DIRSIG handles mixed pixels reduces the amount of computation time needed to arrive at a radiometrically accurate answer by way of sampling based on importance [9]. This reduction in time becomes crucial when multiple DIRSIG simulations are required.

2.6.2 MuSES

One piece of the puzzle that DIRSIG does not compute is the heat transfer via convection, conduction, and thermal radiation between objects in the scene. Appendix

A contains temperature measurements showing the heating of a surface by a transformer. In the span of approximately 7ft the temperature of the surface changed by over 3°C due to heat transfer. In order to accurately model the interactions within a piece of collected imagery the heat transfer needs to be taken into account. This will be done using the Multi-Service Electro-optic Signature (MuSES) thermal signature prediction program from ThermoAnalytics. MuSES employs first-principles physics and empirical data to run simulations [7]. A voxel-based ray-tracer is used to compute view factors for all facets of a model as well as radiosity, and BRDF. The program can simulate low-resolution models for real-time predictions as well as high-resolution models with millions of elements. MuSES can compute a multi-bounce radiosity solution for all model facets as well as background pixels [7]. The output of MuSES is a Thermal Description File, or TDF. This file contains the temperature solutions and the geometry of the simulations.

MuSES can ingest a number of geometric models including the Wavefront Technologies' OBJ file format. MuSES is very sensitive in the type and quality of the surface mesh used in the model. MuSES will only use tri (three vertices) and quad (four vertices) meshes. The thermal solution is calculated at a node placed at the centroid of each mesh element. The size of the mesh elements affects the length and quality of the simulations. A finely meshed surface will produce a more accurate thermal solution but this comes at the cost of increased simulation run time. The mesh elements should be small enough to provide sufficient resolution to avoid large thermal gradients.

The aspect ratio (the ratio of the length to the width) of a mesh element is recommended to be no larger than 4:1, however 10:1 is the absolute largest aspect ratio allowed. In short, long, skinny polygons are not ideal shapes for a mesh element. The quality of the mesh of a model has a large effect on the thermal conductivity within the model. If nodes of adjacent surfaces are joined together so that they form a perfect seam, there will be conduction between the elements. If there is a gap or misalignment of the nodes, there will be no conduction. Other mesh problems

include penetrating elements which occur when one object goes through another without sharing vertices and coincident vertices where two separate vertices occupy the same point in space. In both cases heat transfer via conduction cannot occur. The characteristics of a good quality mesh for import into MuSES [10] are as follows:

1. All adjacent polygons share common vertices
2. All polygons are three or four-sided
3. All polygons are planar
4. All polygons have an aspect ratio near unity
5. Polygons are spread uniformly across the surface
6. No overlapping or repeated facets
7. Surface mesh only (*e.g.* thin plates are represented by their exterior surface only)
8. Mesh is broken into meaningful parts

MuSES is typically run from a graphical user interface (GUI) but can also be run from the command line, which allows for scripting to execute multiple simulations without user interaction.

2.6.3 MODTRAN

The MODerate spectral resolution atmospheric TRANsmission (MODTRAN) algorithm is an atmospheric radiative transfer program developed by Air Force Research Laboratory (AFRL) and Spectral Sciences, Inc (SSI). MODTRAN models the atmosphere as a stack of homogeneous layers [11]. Each layer can be assigned a temperature and pressure either from one of the standard atmospheric models in MODTRAN or radiosonde data input by the user. There are six standard atmospheres in MODTRAN ranging from tropical to mid-latitude summer to sub-arctic

winter. Each of the standard atmospheric profiles can be made to be more representative of a specific location by incorporating user-specified variables such as aerosol and cloud parameters, time of year, time of day, and the latitude and longitude of the location. The standard output of a MODTRAN run is the atmospheric transmission and the upwelled radiance. For this research, DIRSIG will use MODTRAN to compute the total downwelled radiance by placing a sensor at ground level so that it looks at an unobscured sky and running MODTRAN for each combination of 6 zenith angles and 12 azimuthal angles of the sensor. The atmospheric values from these runs are stored in the atmospheric database (ADB) file. After a ray from the sensor strikes a spot on the ground, DIRSIG sends out rays from the spot to sample the radiance from the sky dome by referencing the ADB file. Each ray cast from the sensor will have a specific look-angle and DIRSIG interpolates the value for each ray from the surrounding look-angles in the ADB file.

2.7 Regression

When analyzing a dataset, the goal is often to look for relationships between the variables so that an output can be determined from one or more input parameters. It would be nice to be able to collect an infinite amount of data so that for every possible value of x that there was a corresponding y value and a relationship between the two variables was exactly known. This is impossible in practice so regression analysis is used as a way to interpolate between collected x values by finding a relationship between a sampled set of x and y . In regression analysis, y is typically called the dependent variable and x is called the independent variable, or the regressor variable. It is not uncommon for there to be more than one independent variable.

As stated above, regression analysis is used to functionally relate a dependent variable to a set of independent variables. In order to perform regression analysis a functional form for the relationship must first be specified. This form can and should be checked for adequacy later on in the process. In the case of simple linear

regression the model is

$$y = \beta_0 + \beta_1 x + \epsilon \quad (2.32)$$

Since the model is estimating the relationship between the variables, an error term, ϵ , is introduced. The error is a random model error component that is assumed to have a mean of zero and a variance of σ^2 [12]. The errors are also assumed uncorrelated and normally distributed. The model error term describes the deviation of each dependent variable value from the unknown population regression line and is unknown [13]. An extension of simple linear regression is multiple linear regression where a linear relationship is found between the dependent variable and multiple regressor variables. The functional form for multiple linear regression is

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \cdots + \beta_k x_k + \epsilon \quad (2.33)$$

Models might also include interaction terms which are the product of regressor variables, for example $x_1 x_2$. A third type of regression is polynomial regression where the functional form of the relationship between y and the independent variables is a k th-order polynomial. Polynomial regression is still considered to be a form of linear regression because it is linear in the regression coefficients. For one variable the model is

$$y = \beta_0 + \beta_1 x + \beta_2 x^2 + \cdots + \beta_k x^k + \epsilon \quad (2.34)$$

The model in Equation 2.34 can be extended to account for multiple regressor variables as well as interaction terms. The models listed are not a comprehensive list of functional models in regression analysis, they are simply the ones that are used most often.

Regardless of the functional form of the model chosen for the linear regression analysis, it can be expressed in matrix form as

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon} \quad (2.35)$$

where \mathbf{y} is an $n \times 1$ vector of the observations, \mathbf{X} is an $n \times m$ matrix of the levels of the regressor variables and is also sometimes referred to as the design matrix, $\boldsymbol{\beta}$ is an $m \times 1$ vector of the regression coefficients, and $\boldsymbol{\epsilon}$ is an $n \times 1$ vector of random errors. In this case, n is the number of rows in the \mathbf{X} matrix and m is the number of columns.

Multiple methods exist to calculate the vector of regression coefficients that best fit the model. One of the most widely used is the method of least squares.

2.7.1 Least Squares

Least squares problems can be divided into two main categories: linear and non-linear. This research will only deal with a linear least squares approach. The least squares estimates of the regression coefficients are calculated by

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{y} \quad (2.36)$$

which is the pseudoinverse of \mathbf{X} multiplied by the dependent variable observations. The vector of fitted values \hat{y}_i that correspond to the observed values y_i is calculated by

$$\hat{\mathbf{y}} = \mathbf{X}\hat{\boldsymbol{\beta}} \quad (2.37)$$

The quantity $y_i - \hat{y}_i$ is known as the residual. The residuals are the observed vertical deviation of the dependent variable values from the fitted regression line [13]. The vector of residuals can also be written in matrix notation

$$\mathbf{e} = \mathbf{y} - \hat{\mathbf{y}} \quad (2.38)$$

The method of least squares estimates the regression coefficients such that the sum of the squared residuals is minimized. At this point the least squares estimates of

the regression coefficients can be used to make predictions of a new set of outcomes from a new set of observed regressor values.

2.7.2 Model Adequacy

Once a functional form is chosen for a model it should be tested to determine if it is an adequate representation of the data. This research employs two methods to determine which model is most adequate for the data set. The first method is analysis of the residuals of a model. The second involves comparing models and using a numeric metric to determine the superiority of one model over another.

2.7.2.1 Residual Analysis

Since residuals can be thought of as the deviation between the observed data and the fitted data, they can also be thought of as a measure of the variability in the dependent variable not explained by the regression model [12]. Plotting the residuals versus the fitted values, \hat{y} , is a useful way of detecting common inadequacies. Ideally the residual plot would show the residuals randomly scattered within a horizontal band about zero. This is shown in Figure 2.6.




Figure 2.6: Plot showing the ideal pattern of residuals.

If the residuals are contained within a horizontal band then there are no obvious inadequacies in the model. If, for example, a first order model was being tested and the residuals revealed a polynomial shape, that would indicate that a higher-order model was required to explain all of the variation in the data set.

2.7.2.2 Testing Individual Coefficients for Significance

Another model-to-model comparison involves comparing the number of statistically significant coefficients in different models. One rule of thumb when choosing the order of a polynomial regression model is to check the number of statistically significant coefficients for a given model and compare that to the next higher order model [14]. If the number of significant coefficients increases then the higher order model should be used, the thought being that an increase in the number of significant coefficients means that the model is explaining variability in y that the other models are not able to.

Statistically significant means that a result is unlikely to have happened by chance. When testing the significance of an individual regression coefficient the hypotheses are as follows:

$$H_0 : \beta_j = 0, \quad H_1 : \beta_j \neq 0 \quad (2.39)$$

where H_0 is the null hypothesis and H_1 is the alternative hypothesis. The null hypothesis is that the coefficients are equal to zero. Rejection of the null hypothesis implies that a coefficient is different enough from zero that its regressor variable is of value in explaining the variability in the dependent variable. The test statistic that is used to determine whether the coefficient is different enough from zero to reject the null hypothesis is the two-tail t -test. The t -test quantifies the contribution of a regressor given that all of the other regressors are in the model [12]. The t -statistic for the coefficient β_j is:

$$t_o = \frac{\hat{\beta}_j}{\sqrt{\hat{\sigma}^2 C_{jj}}} \quad (2.40)$$

where C_{jj} is the diagonal element of $(\mathbf{X}'\mathbf{X})^{-1}$ corresponding to the coefficient $\hat{\beta}_j$, and $\hat{\sigma}^2$ is defined as

$$\hat{\sigma}^2 = \frac{\mathbf{y}'\mathbf{y} - \hat{\boldsymbol{\beta}}'\mathbf{X}'\mathbf{y}}{n - p} \quad (2.41)$$

where n is the number of observations and p is the number of regression coefficients. The t -test is performed at a certain significance level, α . The value of α is the threshold for rejecting the null hypothesis and is set prior to the t -test being administered. In other words α sets the standard for how different the coefficient has to be from zero in order for the null hypothesis to be rejected. Common values used for α are 0.05 and 0.01 and, in general, the smaller α is the more the coefficient has to differ from zero in order for it to be deemed statistically significant.

All test statistics have an associated probability called the p-value. The p-value is the probability of obtaining a value of the test statistic that is at least as extreme as the one observed, assuming the null hypothesis cannot be rejected. For the t -test the p-value is computed as

$$p = 2 [1 - F(|t_o|)] \quad (2.42)$$

where $F(|t_o|)$ is the cumulative distribution function of the t -distribution with $n - p$ degrees of freedom. In order to make a determination concerning the null hypothesis, the p-value is compared to α . If the p-value is less than α , the null hypothesis is rejected and the coefficient is statistically significant at that α -level.

One other factor to consider when testing for the significance of the regression coefficients is the multiple testing problem. With an increased number of hypothesis tests comes an increased likelihood of a chance finding [15]. For example if α was

set to 0.05 for one test then the probability of a false positive occurring is 5%. However, if 100 independent tests are conducted then the probability of at least one false positive rises to 99.4% if the null hypothesis is true for all 100 tests. The Bonferroni correction, developed by Italian mathematician Carlo Emilio Bonferroni, provides one solution to the multiple testing problem. The Bonferroni correction says that if one wants a familywise alpha level of α then each of the n individual hypothesis tests should be conducted at the α/n significance level [16].

2.7.3 Sample size selection

One logical question is how many samples are required in order for the regression model to find a significant relationship in the data set. There are many metrics that have been published and the rules-of-thumb give a wide range of sample sizes. Brooks and Barcikowski [17] compiled a listing of nine sample size selection rules-of-thumb. The sample sizes ranged from $N = (10 \cdot p)$, where p is the number of independent variables, to $N > 100$. For this research 50 samples will be used. This number of samples could support a regression model with five independent variables, according to the metric. However, there will only be one independent variable in the equation so the number 50 was chosen because it fell between the minimum and maximum sample sizes as given by the rules-of-thumb.

2.8 Error propagation

A process with one or more inputs will contain some degree of error. If the process can be modeled using a governing equation where the parameter of interest, Y , is a function of one or more inputs, namely

$$Y = f(X_1, X_2, \dots, X_N) \tag{2.43}$$

then the error in Y can be expressed as

$$S_y = \left[\left(\frac{\partial Y}{\partial X_1} S_{X1} \right)^2 + \left(\frac{\partial Y}{\partial X_2} S_{X2} \right)^2 + \dots + \left(\frac{\partial Y}{\partial X_N} S_{XN} \right)^2 \right]^{\frac{1}{2}} \quad (2.44)$$

where s_{X_i} is the error in the individual input variables. The partial derivatives with respect to each input represent the how sensitive the process is to a change in each individual input. Equation 2.44 assumes that the inputs are uncorrelated, meaning a change in one input does not induce a change in another. If some of the input variables are correlated then Equation 2.44 has to be modified to account for the correlation [1].

$$S_y = \left[\left(\frac{\partial Y}{\partial X_1} S_{X1} \right)^2 + \left(\frac{\partial Y}{\partial X_2} S_{X2} \right)^2 + \dots + \left(\frac{\partial Y}{\partial X_N} S_{XN} \right)^2 + \sum 2\rho_{ij} \frac{\partial Y}{\partial X_i} \frac{\partial Y}{\partial X_j} S_{X_i} S_{X_j} \right]^{\frac{1}{2}} \quad (2.45)$$

where ρ_{ij} is the correlation coefficient between X_i and X_j and the sum is over all combinations of correlated variables.

If a governing equation is not available the values of the partial derivatives are computed using simulated data and Monte Carlo methods [1]. In this situation one input variable is varied at a time for multiple values and the relationship between the parameter of interest and the varied parameter is derived from the empirical data.

2.9 Temperature extraction from multiple pixels

If enough terms are known, the governing equation from Section 2.6.5 can be inverted and solved for the background temperature. The imagery is assumed to be atmospherically corrected so the atmospheric transmission and upwelled radiance terms drop out of the equation. This leaves the self-emitted term and the reflected

downwelled radiance term. Section 2.6.3 describes the method for calculating the downwelled radiance term. The downwelled radiance is multiplied by $(1 - \varepsilon_B)$ to arrive at the reflected downwelled radiance. The self-emitted term contains the temperature of the background as well as its emissivity. The background emissivity can be estimated using methods described in the next chapter. This leaves the background temperature as the only unknown in the governing equation. An iterative process can be employed where the radiance is computed for different background temperatures using the governing equation. After each computation, the radiance value is compared to the background radiance value from the collected image and the process continues until the two radiance values are equal. The input temperature that equalized the radiance values is used as the background temperature estimate. This result could also be achieved by inverting the governing equation. This technique was used to estimate the background temperature as the assumption was made that the pixels surrounding the target of interest only contained the background material.

Chapter 3

Previous Work

This section will discuss the previous research into subpixel temperature retrieval methods as well as provide a summary of the methods.

3.1 Subpixel Temperature Retrieval with Multispectral Sensors

Szymanski *et al.* (1999) [18] published a method to extract temperatures from objects smaller than a pixel using multispectral thermal infrared data with emphasis on retrieval accuracy. The focus was on pixels containing two materials, each with a temperature near 300K. They considered the case of materials having substantially different, but known, emissivities as well as a pixel containing a single material at two different temperatures. The bulk of the paper focused on the former case with retrieval results and errors being presented. The analysis was performed on a nighttime synthetic river scene with pixels along the riverbank containing a mixture of land and water. Atmospheric effects were not considered in this study although methods for estimating the atmospheric transmission and the upwelled radiance were described in the paper.

The authors used a spectral mixture analysis approach to solving the problem and

started with the following equation for the top-of-the-atmosphere (TOA) radiance for a mixed pixel in band i

$$L_{mixed}^i = f_1 \epsilon_1^i L_{BB}(T_1) + f_2 \epsilon_2^i L_{BB}(T_2) \tau_{atm}^i + L_u^i + L_d^i(r^i) \tau_{atm}^i \quad (3.1)$$

where f_1 and f_2 are the areal fractions of the pixel that are at temperatures T_1 and T_2 , the ϵ 's are the emissivities of the two materials, τ_{atm}^i is the transmission of the atmosphere, L_u^i is the upwelled radiance, L_d^i is the downwelled radiance, and r^i is the reflectance of the mixed pixel. They noted that the temperatures, emissivities, and transmissions are effective values because of band-averaging. It was assumed that the downwelled radiance would be a negligible effect due to the high emissivities of natural surfaces and therefore was removed from the equation.

One issue that arises from using this approach to the subpixel temperature estimation problem is that there are $2(N+1)$ unknowns (N emissivities for each material and a temperature for each material) and only N observations (bands). In the method presented the assumption is made that the areal fractions can be determined from another source and that the emissivities are known. In this method the emissivity of water is assumed to be well-known in all bands and therefore are assigned no error. This assumption limits the method to calm water surfaces as the emissivity would change with surface roughness introducing error. Atmospheric effects were not included in the simulations so the upwelled radiance was removed from the equation and the TOA radiance becomes

$$L_{reduced}^i = f_1 \epsilon_1^i L_{BB}(T_1) + f_2 \epsilon_2^i L_{BB}(T_2) \quad (3.2)$$

If an atmosphere had been added in their simulation, the quantity $L_{reduced}$ would be the radiance from a mixed pixel after atmospheric correction. The only unknowns in the $L_{reduced}$ equation are the two temperatures and there are N bands, therefore Equation 3.2 can be fit for T_1 and T_2 .

The authors added detector noise, calibration errors, emissivity uncertainties (ex-

cluding water emissivity), and/or areal fraction error into the simulations. Calibration error and detector noise characteristics of the Moderate Resolution Imaging Spectroradiometer (MODIS) and Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) were used. It was assumed that the ground sample distance (GSD) of the visible sensor is four times smaller than that of the thermal sensor which results in the visible pixel area being approximately 6% smaller than the thermal pixel area. Therefore an areal fraction error of $\pm 3\%$ was used in the analysis. The GSD of the ASTER VNIR bands is six times smaller than that of the thermal bands which would translate to a smaller areal fraction error. The land emissivities were determined by performing the temperature retrieval using the average of the emissivities of eight soil types found in the Salisbury database [19]. The pixels analyzed were comprised of 30% - 90% water and 10% - 70% land (soil).

For the simulations of ASTER data with a calibration error of 1.5% and an signal-to-noise ratio (SNR) of 200 the root mean square (RMS) error was 27K for land temperature retrieval and 13 K for water temperature retrieval. These values decreased with lower calibration error and higher SNR. A retrieval was also performed using a hypothetical MWIR band ($3.66\mu\text{m}$ - $4.08\mu\text{m}$) with 1.5% calibration error which resulted in an RMS error of 5K for land temperature and 2.4K for water temperature. For the MODIS simulations a 1.5% calibration error was added and the SNR was 200 for the LWIR bands which resulted in an RMS error of 3.9K for the land temperature retrieval and 1.5K for water temperature retrieval.

3.2 A Method for Satellite Identification of Surface Temperature Fields of Subpixel Resolution

A method for subpixel target temperature and subpixel target area using surface radiance temperatures from two channels of the AVHRR instrument was developed

by Dozier (1981) [20]. It takes advantage of the assumption that a hot target will contribute proportionately more radiance to the MWIR signal compared to the LWIR signal from the same scenario, and can be used with any sensor with a band in each of the spectral regions as long as neither band saturates. The method assumes two temperature fields within a pixel, the target and a uniform background. The target material is at a temperature, T_T and occupies a proportion p of the pixel of interest, where $(0 < p < 1)$. The background material is at a temperature, T_B , and occupies $(1 - p)$. The radiance values observed in AVHRR channels 3 and 4 are denoted by L_3 and L_4 , respectively and are assumed to be blackbodies ($\epsilon = 1$). In the absence of an atmosphere or reflected background radiance, the radiance temperatures sensed by AVHRR are

$$T_3 = L_3^{-1} [pL_3(T_T) + (1 - p)L_3(T_B)] \quad (3.3)$$

$$T_4 = L_4^{-1} [pL_4(T_T) + (1 - p)L_4(T_B)] \quad (3.4)$$

where L_n^{-1} denotes the inversion of the Planck equation. The target temperature and the proportion can be solved for by rearranging Equations 3.3 and 3.4 and solving the set of simultaneous non-linear equations.

$$L_3(T_3) - [pL_3(T_T) + (1 - p)L_3(T_B)] = 0 \quad (3.5)$$

$$L_4(T_4) - [pL_4(T_T) + (1 - p)L_4(T_B)] = 0 \quad (3.6)$$

Knowledge of the background temperature is required, otherwise there are more unknowns than equations. The paper also describes an atmospheric correction factor.

3.2.1 Retrieval of Subpixel Fire Temperature and Fire Area in Moderate Resolution Imaging Spectrometer

Lim *et al.* (2002) [21] applied the Dozier method to fire temperature and area using the MODIS sensor. MODIS was used because its fire sensitive bands have higher saturation values. The background temperature was determined by looking at the cloud-free pixels surrounding the fire in the $10\mu\text{m}$ band and computing the mean temperature. The background emissivity was assumed to be 0.97, which they state is an emissivity typical of vegetation. The Iterative Newton-Raphson Method (INRM) was used to determine fire temperature and area. A sensitivity analysis showed that the fire temperature can be determined to within 5% if the background temperature is known to within $\pm 1\text{K}$ and the emissivity is known to within 0.01 of a unit. The uncertainty of the fire area was as large as 50%.

3.2.2 Application of the Dozier Retrieval to Wildfire Characterization: A Sensitivity Analysis

Giglio and Kendall (2000) performed a sensitivity analysis on the Dozier technique and found that under realistic conditions the random error in retrieved fire temperature to be $\pm 100\text{K}$ and the error in area to be $\pm 50\%$ for a fire that occupies more than 0.05% of the pixel area [22].

3.3 Estimating Subpixel Surface Temperatures and Energy Fluxes from the Vegetation Index-Radiometric Temperature Relationship

Kustas *et al.* [23] developed a procedure, called DisTrad, for subpixel temperature estimation that makes use of the functional relationship between radiometric surface temperatures (T_R) and the Normalized Difference Vegetation Index (NDVI). More often than not, NDVI imagery is available at a higher resolution than the T_R

imagery and therefore can be used to disaggregate the surface temperature pixels to the resolution of the NDVI image. In other words, the DisTrad uses the relationship between NDVI and radiometric temperature to derive T_R at the NDVI pixel resolution [23]. The procedure was developed and tested using data from a single day during the 1997 Southern Great Plains Experiment (SGP97). On this particular day there was significant heterogeneity in T_R from a large spatial variation in moisture and vegetation cover across the scene [23]. The NDVI imagery had a resolution of 24m and the T_R imagery had a resolution of 96m.

The first step in the procedure is to aggregate the 24m NDVI imagery to the 96m resolution of the T_R imagery. This creates $NDVI_{96}$ imagery where 16 of the $NDVI_{24}$ pixels are used to form one $NDVI_{96}$ pixel. The $NDVI_{96}$ pixels are divided into three classes: sparse canopy/bare soil, partial canopy, and high/full canopy. The class assignments are based on the NDVI value. The most homogeneous pixels from each class were selected to be in the analysis. This is done using the $NDVI_{24}$ data and calculating the coefficient of variation (the standard deviation divided by the mean) among the 16 pixels that make up each $NDVI_{96}$ pixel. For each class in the $NDVI_{96}$ image, 25% of the pixels with the lowest coefficient of variation are chosen to be included in the analysis. The thought being that using pixels with uniform cover at the $NDVI_{96}$ scale to derive the relationship between NDVI and T_R at 96m will produce a relationship that will closely resemble the relationship at higher resolutions. To determine the T_R - $NDVI_{96}$ relationship, a least squares fit is performed using a second-order polynomial

$$\hat{T}_{R96}(NDVI_{96}) = \beta_0 + \beta_1 NDVI_{96} + \beta_2 NDVI_{96}^2 \quad (3.7)$$

where \hat{T}_{96} is the estimate of the radiometric temperature at the 96 m resolution. Equation 3.7 does not take into account variability from soil moisture. This is done by calculating a correction factor that accounts for the deviations from the mean value of T_{R96} .

$$\Delta\hat{T}_{R96} = T_{R96} - \hat{T}_{R96} \quad (3.8)$$

The correction factor is applied in the final equation to calculate the radiometric temperature at the 24 m resolution.

$$\hat{T}_{R24} = \hat{T}_{R96}(NDVI_{24}(i)) + \Delta\hat{T}_{R96} \quad (3.9)$$

where $i = 1, 2, 3, \dots, 16$ and $\hat{T}_{R96}(NDVI_{24}(i))$ is calculated using Equation 3.7 with the $NDVI_{24}$ value at the i th pixel.

The procedure was tested by deriving subpixel temperature estimates at 25m and 200m which roughly correspond to the NDVI resolutions of ASTER/Landsat and MODIS, respectively. At the MODIS-equivalent resolution the temperature can be estimated to within $\pm 1.5^\circ\text{C}$ uncertainty. At the ASTER/Landsat resolution it was found that the results from the DisTrad procedure are no more reliable than assuming a uniform subpixel T_R field.

3.4 A Novel Method to Estimate Subpixel Temperature by Fusing Solar-Reflective and Thermal-Infrared Remote-Sensing Data with an Artificial Neural Network

A new approach to estimating subpixel temperature in heterogeneous areas (*e.g.* urban) was developed by Yang *et al.* in 2010. The method makes use of information from high spatial resolution VNIR imagery as well as lower spatial resolution temperature maps to generate temperature estimates at a resolution higher than that of the temperature map [24]. The mechanism used to estimate the subpixel temperatures was the genetic algorithm and self-organizing feature map (GA-SOFM) artificial neural network (ANN) based on the VNIR and TIR data. The authors used

the Kustas method [23] as a benchmark to assess the performance of their algorithm.

Three ASTER VNIR bands (green, red, and near-infrared) were used in the analysis, each with a spatial resolution of 15m. The ASTER surface temperature map, with a resolution of 90m, was used as not only the truth data for the analysis but the basis of the temperature extraction. In other words, a temperature map was used as the input to this process, not a thermal infrared radiance image. Moreover, the temperature map was never converted to a radiance image so all of the analysis was done in temperature space. The VNIR imagery was degraded to a resolution of 90m and the temperature map was degraded to resolutions of 360m and 1080m which gave way to two estimation modes. Mode 1 tried to estimate the surface temperature of the 360m temperature map at the aggregated VNIR resolution of 90m. Mode 2 estimated the surface temperature of the 1080m temperature map at the 90m spatial resolution level. The following example is for Mode 1.

The 90m VNIR image, $VNIR_{90}$, was classified into four land cover types using a supervised classifier. The classes were: vegetation, bare soil, urban, and water. Then the pixels from each land-cover type were separated so that the $VNIR_{90}$ image and the class map derived from the $VNIR_{90}$ image could be degraded to the resolution of the temperature image. This produced a VNIR image with a spatial resolution of 360m ($VNIR_{360}$) and a class map of the same resolution ($CLASS_{360}$). From the VNIR imagery and the class maps at the two resolutions along with the 360m temperature resolution map (T_{360}) estimates of input parameters to the GA-SOFM ANNs could be made. The input parameters included NDVI, temperature vegetation dryness index (TVDI), leaf area index (LAI), vegetation water content (VWC), soil water content (SWC), NIR reflectivity, R waveband reflectivity, degree of vegetation coverage, ratio vegetation index (RVI), soil sand content (SSC), and differenced vegetation index (DVI). Two sets of the input parameters are obtained at each resolution, namely, P_{360} and P_{90} . Then the mapping relationships (MAP_{360}) for each land cover class between P_{360} and T_{360} were built using GA-SOFM ANNs. These mapping relationships were then applied to the P_{90} and the 90m resolution

class map to determine the subpixel temperatures at a resolution of 90m.

For all four land cover types in both modes the GA-SOFM procedure resulted in a lower root mean squared error (RMSE) than the Kustas method. The GA-SOFM procedure produced the smallest RMSE for the bare soil class (1.783K for Mode 1 and 1.573K for Mode 2) and the largest for the urban area for Mode 2 (3.975K).

3.5 Sub-pixel Water Temperature Estimation from Thermal-Infrared Imagery Using Vectorized Lake Features

The work of Sentlinger *et al.* (2008) recognized the need to take land temperature contamination into account when looking at small bodies of water and developed a method to extract subpixel skin temperatures of lakes using *a priori* knowledge of water boundaries [25]. The method was tested using MODIS data and is validated using concurrent ASTER data and in-situ temperature data. MODIS offers temperature and emissivity imagery as well as calibrated at-sensor radiance imagery.

Using the assumption that a mixed pixel is comprised of only two materials, water and land, the at-sensor radiance measured in the k th channel of a sensor is

$$L_{ks} = \tau_k [f_1 \epsilon_{k1} L_{BBk}(T_1) + f_2 \epsilon_{k2} L_{BBk}(T_2)] + L_{kA} \quad (3.10)$$

where τ_k is the atmospheric transmission in band k , f_j is the fractional abundance of the j th material in a pixel, ϵ_{kj} is the emissivity of the j th material in a pixel in band k , and L_{BBk} is the blackbody radiance of a material with temperature T_j . The last term in Equation 3.10 represents a combination of the upwelled and downwelled radiance terms. The reflected downwelled radiance was assumed to be relatively constant across the scene.

In order to solve for the two temperatures, T_1 and T_2 , estimates of the other variables in 3.10 are needed. To determine the fractional abundances, the low resolution radiance image was geo-referenced in a high resolution space using nearest neighbor interpolation. Then the vectorized lake feature data was rasterized into the same high resolution space. The rasterized lake data is cross-referenced into the original low resolution space using the geo-referenced radiance image. The fractional abundance was determined in each low resolution pixel by dividing the number of high resolution pixels of a certain material by the total number of high resolution pixels in a given low resolution pixel [25]. The material emissivity estimates are taken from the MODIS emissivity product using only the pure pixels in the imagery. The atmospheric parameters are estimated from the pure pixels in the radiance image by fitting a line to the measure of at-sensor radiance and estimated ground-leaving radiance, which is the term in square brackets in Equation 3.10. This leaves only two unknowns in Equation 3.10 (the two temperatures). Therefore at least two estimates of the sensor-reaching radiance are required to solve for the unknowns, else the system is underdetermined. The gradient descent method was used to solve for T_1 and T_2 by minimizing the error between the observed at-sensor radiance and the estimated at-sensor radiance from Equation 3.10.

The algorithm was able to estimate subpixel skin temperatures from MODIS data to $\pm 0.32^\circ\text{C}$ of ASTER skin temperature. The absolute error in recovered skin temperature was expected to be $\pm 0.96^\circ\text{C}$.

3.6 An Image-Sharpening Method to Recover Stream Temperatures from ASTER Images

Gustafson *et al.* developed a method to extract stream temperatures from 90 meter spatial resolution ASTER TIR imagery to monitor stream water quality in the Pacific Northwest. They used 15 meter ASTER VNIR data to sharpen the TIR imagery and the method was unique with the use of four endmembers to unmix the

VNIR imagery [26]. First, the VNIR imagery was unmixed into three fraction images corresponding to vegetation, a combined shade/water class, and “other”. The “other” class contained everything that was not vegetation or shade/water.

The NDVI was computed for the VNIR bands and the NDVI values were used to discriminate between vegetation, shade/water, and “other”. The vegetation in the image was expected to have a high NDVI value, the shade/water to have a low value in all bands, and the “other” to have a high NDVI value in Bands 1 and 2. The vegetation endmember spectrum was determined by identifying locations where the NDVI value was in the top 1% of the NDVI histogram and calculating the mean value of the pixels in each band. The shade/water spectrum was calculated in the same way except the bottom 1% of NDVI values in the sum of all three bands was used. The “other” spectrum was computed using the top 1% of pixel NDVI values of the sum of Bands 1 and 2. The endmember spectra were used in the linear mixture model (with the constraint that the fractional abundances add to unity) to unmix the VNIR imagery.

The fraction image for the shade/water endmember was used to separate shade and water into two classes. This was done by finding pure water pixels in the shade/water fraction image and making the assumption that the mixed pixels adjacent to the pure water pixels contained some fraction of water. A 3x3 pixel area surrounding the mixed pixel in question was sampled to determine the “typical” mix of shade, vegetation, and “other”. The fraction of shade was renormalized to the same ratio as the mean of the surrounding non-water pixels, and the remaining shade fraction was used as the water fraction [26].

The four 15m spatial resolution fraction images were aggregated to the resolution of the ASTER TIR imagery (90m) by summing and renormalizing the fractions. For each TIR band, an estimate of the radiance for a “pure” pixel of a specific endmember was made. This was done by selecting a 5x5 pixel area in the fraction image as well as the corresponding 5x5 area in the TIR bands and using the points in a linear

regression model where the dependent variable was radiance and the independent variable was endmember fraction. At the end of this process, there were twenty radiance estimates (one for each of the endmembers for each of the five TIR bands) which were used to estimate the 15m radiance. This was done by multiplying the 90m radiance estimates by the corresponding 15m fraction images for each of the TIR bands, resulting in five 15m TIR radiance images. To extract the temperature the five 15m TIR images were run through the ASTER Temperature Emissivity Separation (TES) algorithm [27].

Since this was a preliminary study of this image sharpening method, an error analysis was not included. However, one observation was that the methodology worked best when the materials to be unmixed had very different kinetic temperatures and/or emissivity spectra [26].

3.7 Subpixel Temperature Estimation from Low Resolution Thermal Infrared Remote Sensing

Ottlé *et al.* (2008) [28] presented a method to estimate subpixel temperatures from TIR imagery using a linear mixing model approach where the model was constrained using estimates of pixel constituent temperatures derived from a land surface model. The derivation of their method begins with the Stefan-Boltzmann law which gives the integrated exitance of an object at a temperature, T , from 3 - 100 μ m.

$$M = \epsilon \sigma T^4 \tag{3.11}$$

where ϵ is the integrated emissivity from 3 - 100 μ m and σ is the Stephan-Boltzmann constant. The total exitance of a homogeneous pixel with n materials is

$$M = \frac{\sum_{i=1}^n \alpha_i M_i}{\sum_i \alpha_i} \quad (3.12)$$

where α_i is the proportion of the pixel area covered by a material. Substituting Equation 3.11 into Equation 3.12 and requiring that the proportions sum to unity gives

$$\epsilon T = \left(\sum_i \epsilon_i \alpha_i T_i^4 \right)^{1/4} \quad (3.13)$$

where T is the temperature of the composite surface and is determined from a thermal sensor, ϵ is the surface emissivity of the entire pixel and ϵ_i is the surface emissivity of a material. The authors make the assumption that for a vegetated area the emissivity is approximately unity so Equation 3.13 becomes

$$T = \sum_i \alpha_i T_i + \xi \quad (3.14)$$

where ξ is an error term that accounts for the error in making the assumption that vegetation has an emissivity of one which they claim to be less than 0.1K in all cases. The authors state if *a priori* estimates of the constituent temperatures, T_i , are known then they can be considered Gaussian random variables because the temperature extraction can be performed in the framework of Bayesian theory and maximum likelihood estimation [28]. Therefore solving for the constituent temperatures is analogous to solving for the minimum of the following cost function

$$J = \frac{1}{\sigma_b^2} \sum_{i=1}^N (T_i - T_{bi})^2 + \frac{1}{\sigma_T^2} \left(\sum_{i=1}^N \alpha_i T_i - T \right)^2 \quad (3.15)$$

where T_{bi} is the *a priori* temperature of a material, σ_b and σ_T are the standard deviations of the uncertainties related to the *a priori* solutions and the observations, respectively. The performance of the method is sensitive to the magnitude of the

observation error relative to the *a priori* error. To beat down the error, multiple pixels can be used in the estimation process if the assumption that the surface temperature for a material is homogeneous within the queried pixels.

A synthetic data set was used to test the method. The results demonstrated the use of *a priori* temperature estimates in reducing the temperature extraction error. They showed the error increased with the number of pixel constituents. The assumption of a homogeneous surface temperature for a material in multiple pixels helped the temperature estimation although the authors concede that this was probably not a valid assumption in real landscapes.

3.8 Previous Work Summary

The methods reviewed in the previous subsections can be broken into two categories: spatial downscaling and spectral mixture analysis. With spatial downscaling information from a high-resolution auxiliary image (*e.g.* VNIR, NDVI) is used to disaggregate a low-resolution thermal image to the resolution of the auxiliary image. The temperature is extracted from the high-resolution TIR image that is created as a result of the process. Spectral mixture analysis starts with the equation for sensor-reaching radiance of a mixed pixel. Variables in the equation are estimated until the only unknowns are the temperatures of the pixel constituents. The number of radiance observations (and therefore bands) required for spectral mixture analysis techniques is equal to the number of unknown temperatures, else the system is underdetermined. The multiple bands requirement is true for spatial downscaling as well because the method requires a thermal image as well as at least one other auxiliary image.

Another similarity between the methods is they are used to determine the temperature of naturally occurring objects that tend to have high emissivities (*e.g.* vegetation, water). Catering the methods to high emissivity targets means that the

reflected downwelled component of sensor-reaching radiance can be assumed negligible. However, it should be noted that Sentlinger *et al.* (2008) did include an estimate of the reflected downwelled radiance in their approach. The targets in the methods were all flat surfaces. Most of the methods that might lend themselves to using a small number of bands require many mixed pixels with the target and background being close to the same temperature. This is a very unlikely scenario in the real world.

The method described in the next chapter will extract the temperature of a uniform small target, where uniform means the lack of a large temperature gradient across the surface of the object. The background will be the only other material in the pixel containing the target. The heat transfer between the target and background will be accounted for as well as the reflected radiance term. Taking the reflected radiance into account means that assumptions about the emissivity and the geometry of the objects do not have to be made, given that the geometry is available by some other means.

Chapter 4

Methodology

Whereas the previous chapter described the phenomenology and the tools that are used to build a physics-based target space, this chapter will detail how the information in the previous chapter is put into practice.

4.1 Overview

An exploitation scenario might go as follows: A thermal image along with a companion higher resolution visible image is captured over a target that is of interest to an analyst. The motivation for extracting the temperature of the target might be to infer an internal process from its surface temperature. As stated in Section 1.2, the assumptions of two pixel constituents and a homogeneous background are made. In order to use the methodology that will be outlined in this chapter, the analyst needs to determine which thermal pixel or pixels contain the target of interest. An iterative algorithm will be used to determine the location of the target and that information will be used to determine in which pixels the target resides. The pixel radiance values from each target-containing pixel are averaged and recorded to use as the input to the final stage of the method.

The analyst can estimate a number of parameter values by utilizing information

from both the visible and thermal image. These parameters will be referred to as the *fixed parameters* as they will not change during the simulations. The fixed parameters are comprised of target emissivity, background emissivity, background temperature, target size, target position, and the weather data. The first four parameters are deemed important because they are included in the governing equation for sensor reaching radiance for a mixed pixel. The target position is another important parameter to the small target problem, even though it is not included in the governing equation, it can influence its prediction. The weather data is used in the estimate of the downwelled radiance as well as the temperature solutions that are computed by the MuSES code. The analyst will also have to decide on a range of target temperatures within which the true target temperature falls. Fifty target temperature values will be generated so that they uniformly span this range. With all other parameter values held constant, a synthetic image will be created for each temperature in the target temperature range resulting in a lookup table (LUT). There will be gaps in the values in the LUT that will be interpolated using linear least squares. The result of the least squares is a set of regression coefficients that are applied to the pixel radiance that was recorded at the beginning of the process to produce a target temperature.

4.2 Parameter estimation

The target space described in the previous section can be thought of as infinitely continuous. In other words, the parameters can take on any value and all possible combinations of parameter values would be represented in the target space. This scenario would be optimum, but it is not physically realizable to have all possible values represented. Therefore the axes need to be bounded according to the problem at hand by making estimates of the range of likely values for the parameters using information from the scene itself. The pixel radiance is physically dependent on the target temperature and comes from the synthetic image created by DIRSIG, so the only parameter range that needs to be estimated is the target temperature.

By bounding the target temperature range, the pixel radiance range will also be bounded. It is important to note that the target temperature range selected must contain the true target temperature because regression analysis performance degrades when extrapolating data. The estimates of the fixed parameters will further bound the range of values that the pixel radiance can take on. The following subsections will discuss methods that are used to estimate the parameters.

4.2.1 Target temperature

Determining the target temperature is the objective of this methodology, however a reasonable guess can be made as to the possible temperature range. For example, if an ice cube were being imaged, the target temperature range would be centered about 0°C. The target temperature range must contain the true target temperature because regression analysis will be used later on in the process to determine the target temperature and it is not always a reliable technique when extrapolating data.

4.2.2 Emissivity data

As stated previously the values of the fixed parameters can be derived from the collected thermal and visible imagery. The analyst might have knowledge of the material the target is made from as well as what type of background material is usually found with the target. The visible imagery can confirm the target and/or background type for the captured imagery as well as reveal surface features that effect the emissivity of a material (e.g. rust). If the visible imager is multispectral or hyperspectral, a reflectance spectrum from both the target and the background could be used for material identification. Using some or all of that information, a spectral database could be queried for an emissivity spectrum for both the background and target.

4.2.3 Background temperature

The background temperature can be estimated from other areas in the scene where the background material fills multiple pixels. Since the pixels are homogeneous, the temperature of the background can be estimated from the method described in Section 2.9, which makes use of the estimate of background emissivity and the assumption of having atmospheric data available to run MODTRAN to compute an estimate of the downwelled radiance.

4.2.4 Target size

An estimate of the physical size of the target can be determined from the higher-resolution visible imagery. The error on this term is ± 1 visible GSD. The analyst's knowledge of the target could be used to reduce that error. For example if the analyst knew that a power plant only has a particular model transformer, its dimensions could be researched from manufacturer specifications.

4.2.5 Target position

An automated algorithm was developed to determine the position of the center of an unresolved target in a thermal image with subpixel accuracy. The idea is to generate synthetic DIRSIG imagery where each image is a simulation of the target center in a different location. The synthetic test images are compared to the collected thermal image to determine which position from the test images is the closest to the actual position. Figure 4.1 pictorially shows the overarching concept of the positioning algorithm.

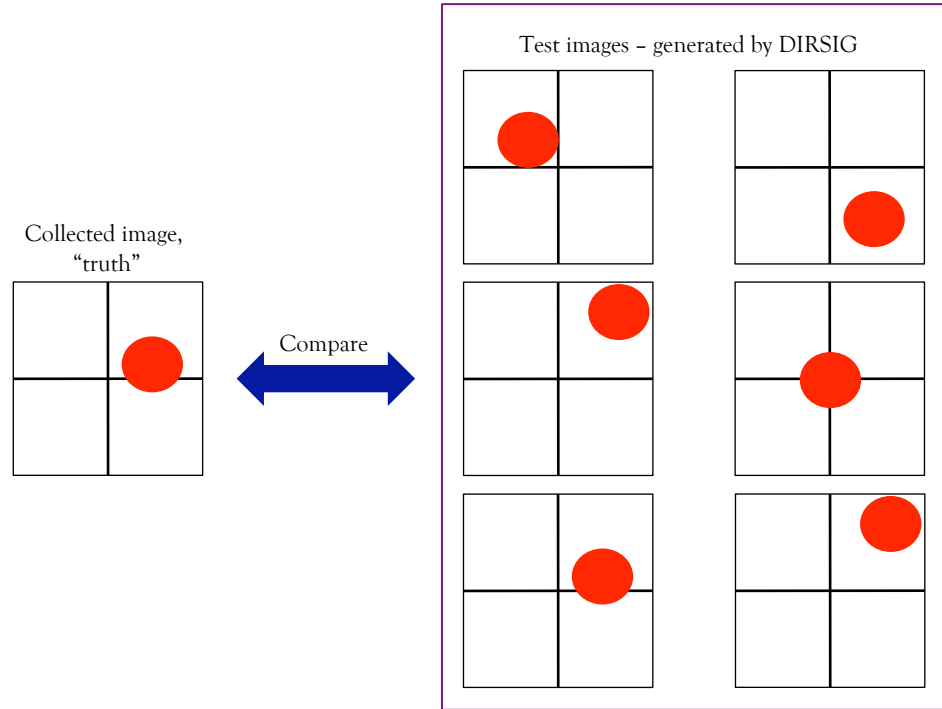


Figure 4.1: Overview of positioning algorithm.

This could be executed in a brute force fashion where hundreds or thousands of synthetic test cases are generated at once and compared to the collected image. It is no surprise that the brute force method is inefficient in terms of the computation time required. This algorithm takes an adaptive, iterative approach where the result from one iteration dictates what happens in the next.

4.2.5.1 Comparison metric for target position determination

A direct pixel radiance comparison cannot be done because the target temperature is unknown. Instead, this algorithm compares normalized pixel radiance values by computing the angle between the vector of radiance values from the collected image and the vector of radiance values from each of the test images. Equation 4.1

computes the angle, θ , between two vectors.

$$\theta = \cos^{-1} \left(\frac{a \cdot b}{||a|| ||b||} \right) \quad (4.1)$$

where a and b are the vectors of the radiance values of the collected and test images, respectively. Each element of the vectors is a radiance value at the individual, observed location. Each vector is divided by its magnitude to create a unit vector, therefore the value of θ is independent of the magnitude of the input vectors. The smaller the angle, the more similar two vectors are.

4.2.5.2 Implementation

The first step in the process is to create a vector of radiance values from the collected thermal image, which will be called the *truth vector*.




Figure 4.2: The radiance values from the 5x5 image chip are reformed into a 25x1 vector. In this example L_1 is the pixel with the most contrast from the background.

It is assumed that the pixel with the most contrast from the background contains the target of interest. A 5x5 pixel region surrounding the pixel with the most contrast from the background is selected and a 25x1 vector is created from the radiance values. This is shown in Figure 4.2. The angle between the truth radiance vector,

L_{truth} , and the synthetic radiance vector, L , can be written as

$$\theta = \cos^{-1} \left(\frac{L \cdot L_{truth}}{\|L\| \|L_{truth}\|} \right) \quad (4.2)$$

The parameter estimates from the previous sections are used to generate the test images. The target temperature is estimated as well. It is important that the choice of target temperature does not reverse the contrast. In other words, if the target pixel appeared brighter in the truth image then it must also be brighter in the test images. Initially the target center is positioned in the center of the simulated scene (position $[0,0,0]$). The camera position information is obtained from the IMU and the GPS data collected during the capture of the thermal image. The camera position information is required so that the GSD of the test images match that of the truth image. The sensor look-angle should also match between the two sets of imagery so that any projected area effects are accounted for. For the first iteration the target is moved one GSD from the center of the scene in every direction. A test image is created for the initial position as well as for every movement giving a total of nine test images. The angle between each test vector and the truth vector is computed and the position from the image that yielded the smallest angle is said to be the closest match.

The position corresponding to the smallest angle is the starting point for the next iteration. In the second iteration the target center is moved by $1/2$ of a GSD plus a small random number in every direction, creating nine more candidate test images. The random number is generated by pulling a value from a random normal distribution with a mean of 0 and a standard deviation of 1 and multiplying it by $1/10$ of $1/2$ the GSD. In the case where the GSD is 1 meter, the value from the normal distribution is multiplied by $1/20$. The random number is added to each position to avoid having half of the target centers always land on the edge of two pixels when they are moved by $1/2$ of a GSD in the second iteration. Again, the angle between the truth and test vectors is computed and the position corresponding to

the smallest angle is used as the starting point for the next iteration.

The third and fourth iterations have the target center moving by $1/4$ and $1/8$ of a GSD, respectively. Figure 4.3 shows the locations of the center of the target for each iteration.




Figure 4.3: The locations of the center of the target are denoted by colored dots. The location of the center of the target for the truth case is marked by the red x. The target center locations in iteration 1 are found in (a). Iteration 2 has the target centers move by $GSD/2 +$ a small random number (b), iteration 3 has them move by $GSD/4 +$ a small random number (c), and iteration 4 has the target centers move by $GSD/8 +$ a small random number (d).

After the angles for the fourth iteration are computed, the position corresponding to the smallest angle is deemed the best estimate of the true target position. This position is recorded and used during the DIRSIG simulations. Using the DIRSIG truth images, the target-containing pixels can be determined.

The algorithm is stopped at the fourth iteration because, for the majority of cases beyond this iteration, the smallest angle does not correspond to the closest position candidate and therefore the starting point for the next iteration is incorrect. The reason for the confusion is illustrated in Figure 4.4 with plots of the normalized radiance vectors for the closest three positions to the truth position in iterations 1 and 4.




Figure 4.4: Plots of the radiance vectors divided by their respective magnitudes versus their array indices for iteration 1 (a) and iteration 4 (b). The red line in (a) and (b) represents the vector from the truth image. The green line represents the vector from the location that was closest to the truth vector. The blue and black lines represent the locations that were the second and third closest, respectively.

In Figure 4.4(a), the maxima and minima of the green line coincide with those of the red truth line, and more importantly the global maximum from each of the lines appear in the same location; which is not true for the black and blue lines. As a result, the dot product between the red and green vectors is larger than the dot product between the red and blue or red and black vectors. It follows that the angle between the red and green vectors is the smallest. This makes sense as large movements of the target center from the truth position will result in vectors that are

very different from the truth vector, in terms of where the maxima and minima lie. However, when the movements about the truth position are small, the maxima and minima of all of the test vectors will coincide with those of the truth vector. This is illustrated in Figure 4.4(b). At this point the comparison metric is “confused” because the only real distinction between any of the vectors is their magnitude. It should be noted that it always selects one of the four closest positions as the final answer when tested with synthetic imagery. In short, the comparison metric employed by this algorithm works well for larger target movements but starts to break down as the movements approach $1/8$ of a GSD and smaller.

4.2.6 Weather data

Section 1.2 stated that it is assumed that upper air atmospheric information was known. This data could come from a radiosonde or a MODTRAN atmospheric model. Another source of atmospheric data as well as surface weather data is the Regional Atmospheric Modeling System (RAMS). RAMS was developed at Colorado State University to simulate and forecast meteorological phenomena and render the results. The major components of RAMS are [29]:

1. An atmospheric model which performs the actual simulations
2. A data analysis package which prepares initial data for the atmospheric model from observed meteorological data
3. A post-processing model visualization and analysis package which interfaces atmospheric model output with a variety of visualization software utilities.

RAMS uses models for solar and terrestrial radiation, effects of latent and sensible heat exchange with surface water, soil, and vegetation, and the formation and interaction of clouds. It can simulate data for areas on the order of hundreds of meters to areas the size of continents [30]. RAMS is used at Savannah River National Laboratory (SRNL), which is part of the Department of Energy.

Weather information could also come from a local weather station or a website like Weather Underground.

Now that the estimates of the parameters have been determined, the geometry of the scene needs to be acquired. This geometry will be attributed with the parameter estimates and used as the input to the physics-based simulations.

4.3 Geometry

One fortunate outcome of the subpixel problem is that even though a thermal image could cover thousands of square meters, only a relatively small area of the scene in the neighborhood of the target has an effect on the pixel radiance. A geometric model of the portion of the scene with the target in it is required for both MuSES and DIRSIG. At a minimum, the MuSES model must contain the target and the background material in order to model the heat transfer between the objects. However, if there is an obstruction (*e.g.* a tree) that shades the target or background, the obstruction needs to be modeled in MuSES as well. It is crucial that the mesh on each piece of geometry follows the meshing guidelines discussed in Section 2.6.2. The geometry for the DIRSIG scene should include any surrounding objects that could contribute to the reflected background radiance.

The geometry can come from a variety of sources. One source could be a three-dimensional model extracted from a point cloud derived from multi-view visible imagery or a LIDAR point cloud [31]. This has advantages in that the model is built using real imagery of the target of interest as well as allowing for the extraction of a large number of objects with varying shapes and sizes. However, the more complex the geometry of the target the less accurate the extracted model will be. Even with the highest resolution imagery available, the model extracted from the point cloud will still contain error due to the model fitting process. If there is *a priori* knowledge of what the target of interest is, a Computer-Aided Design (CAD) program could

be used to produce a detailed model. This requires significant knowledge of the target of interest as well as a knowledgeable CAD user. This method can produce geometry of complicated objects but this comes at the cost of a significant amount of time needed to produce the model.

4.4 Physics-based simulations

This section will discuss the implementation of the modeling tools described in Section 2.6. The total simulation time between MuSES, MODTRAN, and DIRSIG has the potential to become lengthy. The total simulation time depends on factors such as the number of simulations used to build the LUT, the complexity of the scene geometry, and the amount of sampling that is allowed per simulation.

There can only be one instance of MuSES running at a time per license file, so the MuSES simulations are run serially on one core. A Symmetric Multi-Processor large memory machine maintained by Research Computing at RIT is used for the MuSES simulations. It has four dual core AMD Opteron processors, each rated at 2.8 GHz with access to 80 GB of main memory, and is connected via NFS to the high performance file server. Multiple instances of MODTRAN and DIRSIG are allowed to run at one time and the “condominium” cluster (also maintained by Research Computing) is used to process those simulations. The cluster consists of 128 shared cores spread over four nodes, however, the maximum number of cores allowed to one user is limited to 80. Jobs are submitted to the head node which then does each job out to the individual cores. UNIX scripts were written to automate the simulation process.

4.4.1 MuSES simulations

As stated previously, MuSES will be used to incorporate heat transfer due to conduction, convection, and thermal radiation into the synthetic imagery used to build the lookup table. This section discusses how the geometric model from the previous

section will be attributed, how the weather and solar loading will be taken into account in the temperature solutions, and how the MuSES simulations will be run.

4.4.1.1 Attribution of material properties

The meshed model from Section 4.3 can be sectioned into parts by the user. Multiple parts can be defined within a single model and each part contains mesh elements that have common material properties. One such property is the surface emissivity, whose value comes from the parameter estimates. Other properties include the specific heat, density, and conductivity. These values can be assigned using the values for materials in the MuSES library or, for materials not in the library, they can be manually assigned by the user. The type of convection and material thickness are also required. MuSES specifies its own convection types., which are listed in the MuSES User Manual. This research used the “Wind” type, which allows for wind convection to be applied to the part being modeled. The thickness of the material is used to determine the capacitance of the part for conduction and does not change the model geometry. The parts for this research are specified as two-sided, meaning they have a front side and a back side, and each side needs to be attributed with a set of material properties. Each side can have a set of material properties that is independent from the other.

Parts can be either solid, liquid, or gas. It is possible to simulate liquid partially or completely filling a container simply by creating new parts and attributing them appropriately with material properties and convection types. Figure 4.5 shows a drum segmented such that half the volume is filled with water and half is filled with air. To specify a liquid or gas requires knowledge of the convection coefficient, h , which was described in Section 2.4.

At this point the geometry from Section 4.3 will be attributed with the proper material properties and convection types. The next step is to model the effect the environment has on the target and background using the weather data that was




Figure 4.5: A drum broken into two parts. The blue represents the level of the water volume in the container and the green represents the air volume.

collected.

4.4.1.2 Weather files

MuSES has three types of weather files that it can ingest when in the natural environment mode. The PRISM format is the simplest of the three and is used in this research. The fields contained in the PRISM format are found in Table 4.1.

Table 4.1: Weather history parameters

Date	Time	Wind speed
Wind direction	Air temperature	Relative humidity
Air pressure	Cloud cover	Sky radiance

If the sky radiance and solar emission terms have not been measured, MuSES will model it based on geographic location. This requires the user to supply the latitude, longitude, and elevation of the target of interest as well as the day of the year and the time zone where the simulation should occur. The weather file contains cloud cover expressed on a 10 point scale where 0 represents 0% cloud coverage and 10

represents 100% cloud coverage. ThermoAnalytics provides a Perl script that ingests WeatherUnderground weather data and creates a MuSES weather file. When the Perl script reads in the weather data, it correlates “partly cloudy” to 30% cloud coverage, “scattered clouds” to 40% coverage, and “mostly cloudy” to 80% coverage. The coverage is averaged across the sky. In the actual weather on a day with scattered clouds, they occasional drift past the sun and block it for some period some time. MuSES lacks the ability to change cloud coverage with time in this manner (except as it changes in the weather file). Therefore “scattered clouds” is interpreted as 40% of the sun is covered by a cloud and the other 60% is “shining”. However, the value for the amount of cloud coverage does not have to come from Weather Underground inputs so any percentage value can be used in the weather files (*e.g.* 15%, 60%).

Thermal inertia is a measure of a material’s ability of resist a change in temperature and represents the ability of a material to conduct and store heat. The temperature of a material with low thermal inertia changes significantly during the day, while the temperature of a material with high thermal inertia does not change as drastically. Furthermore, the temperature of a material with high thermal inertia can be a result of the environmental conditions more than 24 hours prior to the temperature observation. It is because of this phenomenon, that multiple days worth of weather data are required to get the most accurate temperature solution from the MuSES simulations. The number of days depends upon the properties of the material or materials in question. For example, a slab of asphalt 18” thick requires approximately five days of weather history, whereas a slab 6” thick requires approximately 3 days. Once the amount of weather history data is determined, the weather file is built. The weather history can come from any reliable source, such as the ones mentioned in Section 4.2.6, as long as it contains the parameters found in Table 4.1. The weather data from the 24 hours prior to the collect will be interpolated so that there is a data point every five minutes. The data is interpolated to such a fine resolution so that the weather data can be changed on a smaller time-scale when the cloud cover is refined.

The raw weather data in the MuSES weather file will be inaccurate, to some degree, due to the actual location of the target, especially if the weather data is not pulled from a weather station at the target location. Therefore it is advantageous to refine the weather data so that it more closely represents the conditions at the target location. Cloud cover is a significant driver of temperatures in a heat transfer sense and is also the parameter that will most likely be the least well characterized so it was chosen as the means to refine the weather file. The refinement is done using the imagery-derived estimate of the background temperature. This implies that there is higher confidence in the fidelity of the imagery-derived metric than there is in the cloud cover estimate.

Before the MuSES simulations are started, one simulation is run with the unrefined weather file. When the simulation is complete, the background temperature from the MuSES simulation is compared to the imagery-derived background temperature estimate. If the two temperatures match to within 0.05°C the process is stopped and the original weather file is used during the rest of the process. If the temperatures do not match, the cloud cover is modified as follows:

- If $T_{MuSES} < T_b$, then 0.2 is subtracted from the daytime hours of the previous 24 hours of cloud cover and 0.2 is added to the nighttime cloud cover.
- If $T_{MuSES} > T_b$, then 0.2 is added to the daytime hours of the previous 24 hours of cloud cover and 0.2 is subtracted from the nighttime cloud cover.

The process iterates until the two temperatures match to within the threshold. The result of this process is a weather file with the cloud cover modified such that the modeled background temperature matches its imagery-derived counterpart. A flowchart of the process is shown in Figure 4.6.

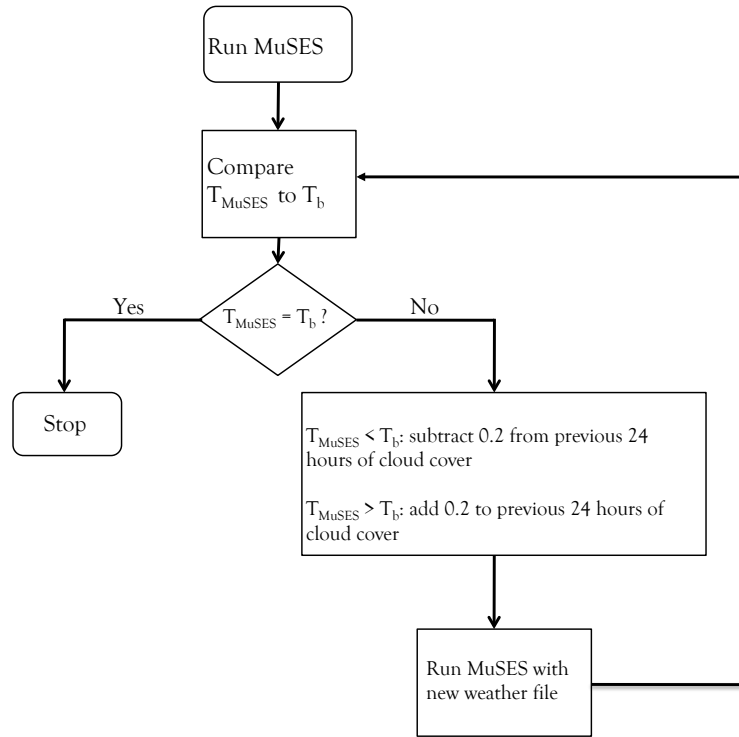


Figure 4.6: FLOWchart showing the process to refine the cloud cover data in the MuSES weather file.

Again, any parameter can be modified. Cloud cover is chosen because that is the parameter

4.4.1.3 Running MuSES

To build the lookup table the target temperature will be varied in MuSES. Since the assumption of a uniform surface temperature has been made, the temperature assigned to the surface of the target will remain fixed during a MuSES simulation. Fifty target temperature values will be generated, each corresponding to a MuSES simulation. The temperature values will be selected by generating 50 numbers from a uniform random distribution and scaling the numbers so that they fall within the pre-defined target temperature range. A uniform distribution was used so that there would be an even distribution of the temperatures within the target temperature

range. It should be noted that there are multiple ways to select the candidate temperature, for example, dividing the temperature range into equally spaced steps. Each temperature will be part of the TDF filename so that it can be used in the LUT later on in the process. The temperature solutions in the TDF files form the basis of the radiance computations in DIRSIG. Before the radiance solution can be computed, MODTRAN runs are set up to compute the downwelling sky radiance.

4.4.2 MODTRAN simulations

MODTRAN is used to create an atmospheric database (ADB) file, which holds all of the atmospheric data that DIRSIG needs to render a scene. Recall the thermal image is assumed to be atmospherically corrected so DIRSIG will only use the section of the ADB file that contains the downwelled radiance information.

The first step in creating the ADB file is to set up the MODTRAN tape5 file so that it matches the collection conditions as closely as possible. Data from a radiosonde can be imported or, if that information is not available, the most appropriate of the six standard atmospheres can be used. To get the temperature of the boundary layer, the weather history file in DIRSIG is invoked. The weather history file contains a time history of meteorological conditions at the scene site. The air temperature data in the weather file is used by MODTRAN as an estimate of the boundary layer temperature. It is important that the air temperature from the weather file, used by MuSES to compute the temperature solution, matches the air temperature that is used by MODTRAN to compute the downwelled radiance so that the boundary layer temperatures are the same. Therefore, before the ADB file is made, the DIRSIG weather file is modified so that the air temperature matches the MuSES simulations.




Figure 4.7: Flowchart depicting the process used to create the downwelled radiance section of the DIRSIG ADB file. One ADB file is created per MuSES simulation.

Once the MODTRAN input is set up properly, the downwelled radiance is calculated by placing the sensor where the target is and sampling the sky at 12 azimuth angles and 6 zenith angles. This creates a hemispherical map of the scattered and emitted radiance in the sky for 72 distinct look-angles. Figure 4.7 shows the process of creating the downwelled radiance section of the ADB file if radiosonde data is present.

At this point, the temperature solutions from MuSES and their corresponding ADB files are ready to be imported into DIRSIG, where the radiance images are created.

4.4.3 DIRSIG simulations

A DIRSIG enhancement that came about as a result of this research was the ability to use a MuSES TDF to drive target temperatures during the radiometric computation. The TDF files containing the temperature solutions are imported into DIRSIG, where one radiance image per TDF file is rendered. The geometry contained in the TDF file can be placed in a pre-existing DIRSIG scene if need be.

4.4.3.1 Setting up the DIRSIG scene

The sensor location (θ , ϕ) and height are known from the collection data and are used in the DIRSIG simulations. The sensor is modeled using knowledge of the spectral response function as well as attributes such as the focal length, pixel pitch, and number of pixels on the focal plane. To sample the scene, adaptive subsampling is invoked using the values for the user-supplied inputs found in Table 4.2.

Table 4.2: Adaptive sampling parameter values.

Parameter	Value
Sub-element grid size	10 x 10 elements
Maximum number of samples per pixel	1000
Radiance difference threshold	0.001

The adaptive sampling code produces imagery with the same number of pixels as the collected image. The oversampling occurs “under the hood” during the radiometric calculation process. Finally, the emissivity curve for each material type must be specified. These curves come from the parameter estimates in Section 4.2. In the DIRSIG gist file, each DIRSIG material ID has to be mapped to their corresponding MuSES material ID. This mapping has to occur because DIRSIG cannot (at the time of this research) automatically correlate the different material IDs.

4.4.3.2 Running the simulations

The scene and sensor parameters from the previous section do not change from one simulation to another in this process. In other words, the only variant between the DIRSIG simulations is the TDF file. Since each TDF contains a target with a different temperature, the only input parameter that is allowed to change in DIRSIG is the target temperature. Each TDF is run through the DIRSIG code to produce a radiance image. The radiance image incorporates the temperature distributions as well as the downwelled radiance and any background reflected radiance from the scene. The flag to only calculate the downwelled radiance component is turned on

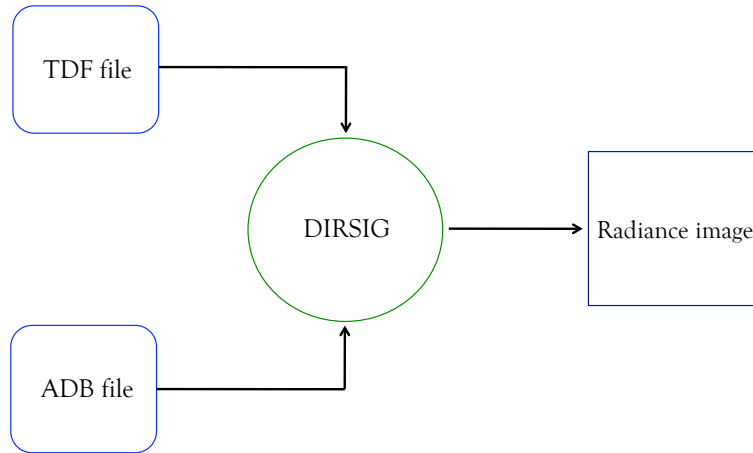


Figure 4.8: Flowchart depicting the output of the DIRSIG simulation process for every TDF.

so the atmospheric transmission and path radiance are not included in the simulations. It also incorporates all of the sensor characteristics, meaning that the GSD of the simulations will match that of the collected thermal image and the radiance reaching the sensor will be modified according the spectral response function of the real thermal imager.

After the DIRSIG simulations are complete, each of the 50 radiance images are read into MATLAB where spatial degradation using the sensor MTF applied. This was done in MATLAB vice DIRSIG because MATLAB serves as a “sanity check” as the user has complete control over computations, whereas DIRSIG computes everything “under the hood”.

4.5 Building the lookup table

The lookup table consists of two values per DIRSIG image. The first is the difference between the target temperature and the background temperature (ΔT). The second is the difference between the radiance averaged over the target-containing pixels and

the radiance from the background. The second is the difference between the radiance averaged over the target-containing pixels and the radiance from the background. This is illustrated in Figure 4.9.




Figure 4.9: Illustration of the delta values computed for the LUT.

The delta values are used to help mitigate any additive errors. For example, the atmospheric compensation might not have been perfect and so there might be an error in the upwelled radiance, which is an additive term in the governing equation. The delta values are computed from each of the 50 DIRSIG simulations and recorded in the lookup table




Figure 4.10: The lookup table is built by computing the delta temperature and delta radiance values from a DIRSIG simulation and recording it as a line in the LUT. This is done for all 50 simulations.

It should be noted that the LUT will only be valid for the geometry and atmospheric conditions at the time of collect. At the end of this process the LUT is complete. The next step is to perform a regression analysis on the data in the LUT to build a statistical relationship between the delta temperature and delta radiance values from the synthetic DIRSIG imagery. The delta radiance value computed from the collected thermal image will be used as the input to the regression equation.

4.6 Regression

Regression analysis is used as a way of modeling the relationship between the data contained in the LUT that was created in the previous section. The dependent variable is ΔT , which will be used to solve for the actual target temperature. The input to the equation is the ΔL value that is computed from the collected thermal image. A single regressor does not allow for interaction terms, so the model is some form of

$$\Delta T = \beta_0 + \beta_1 \Delta L_{pix} + \beta_2 \Delta L_{pix}^2 + \dots + \beta_i \Delta L_{pix}^n \quad (4.3)$$

The number of significant coefficients is computed for a first through fourth-order model at the $\alpha = 0.05/n$ level (where n is the number of regression coefficients) and the model with the greatest number of statistically significant coefficients is selected to be the most appropriate for the data. However, the residuals are also checked to make sure they exhibit the characteristics that indicate a satisfactory model. It should be noted that the order of the model does not imply anything about, or have any relationship to, the physics-based processes used to populate the LUTs. The regression model simply provides a statistical relationship between the dependent and independent variables.

The fifth-order model was not tested because it produced similar results compared to the fourth-order in terms of the magnitude of the residual error (standard deviation of the residuals for a fifth-order model is 0.01°C smaller than the standard deviation of the residuals for a fourth-order model) as well as the extraction error. This conclusion is based on the analysis of LUT data from targets with simple geometry and might not be true for data from more complex targets, like electrical transformers. For the sets of data used in this research, linear interpolation gave the same result as regression analysis. However, regression analysis is extendable to multi-dimensional lookup tables, making it slightly more robust. Once the regression coefficients are calculated, the ΔL value from the collected thermal image is used as the input to the regression equation, which results in a value for ΔT . The background temperature is added to ΔT to arrive at the target temperature

$$T_T = \Delta T + T_b \quad (4.4)$$

4.7 Testing the method

The methodology was tested using imagery collected with a single-band LWIR sensor as well as synthetic imagery.

4.7.1 Airborne collection

The purpose of this study is to test the methodology by collecting thermal infrared imagery with a single-band sensor on an airborne platform. Nine different conditions were captured in each image, ranging from a high temperature/high emissivity target on the asphalt background to an ambient temperature/low emissivity target. The LWIR camera on RIT's Wildfire Airborne Sensor Program (WASP) suite was used as the imager. It is an Indigo Phoenix Quantum Well Infrared Photodetector (QWIP) array whose specifications are found in Table 4.3

Table 4.3: Indigo Phoenix LWIR camera specifications

Specification	Value
Model	Quantum Well Infrared Photodetectors (QWIP)
Spectral range	8.0 - 9.2 μ m
Resolution	640 x 512 pixels
Detector size	25 x 25 μ m
Dynamic range	14 bits
Focal length	25mm
Field of view	35.5° x 28.7°
IFoV	1.00 [mrad]

WASP flew over the targets at an altitude of 1,290ft above ground level, resulting in a GSD of 1.27ft in the thermal infrared. Flying overhead allows for nadir and near-nadir imaging. The experiment took place on the night of 5 August 2011 between 2145 and 2230 to ensure that the sun did not influence the targets. Nighttime imaging also allowed the asphalt to cool down so that a larger temperature contrast between target and background could be achieved. The sensor made four passes in a south-north flight line and an additional two passes in an east-west flight line.

4.7.1.1 Targets

The targets used in the collect were water-tight 11in x 8in x 3in aluminum boxes. When compared to the GSD achieved with the thermal sensor, one box covered approximately 38% of the projected area.




Figure 4.11: Polished box (bottom), sandblasted box (upper left), and box painted with Krylon Ultra Flat Black paint (upper right).

Nine boxes were procured and three were painted with Krylon Ultra Flat Black #1602 spray paint to produce a high emissivity surface, three were polished to produce a low emissivity surface, and three were sandblasted to produce an emissivity in between. Examples of each box are found in Figure 4.11. Spouts were added to the front face of each box so that they could be filled with water and laid out for the collect without water leaking out of the spout. The surface temperature of the boxes was modified by filling the boxes with water at different temperatures. Water was used to modify the surface temperature for a number of reasons. One reason being using hot and cold water was far less complex than designing a heating system that would require a portable power source, electrical wires, etc. Another reason was that water completely filling a box provides a way of modifying the surface

temperature in a uniform fashion. Introducing a heating element into such a small cavity would introduce localized hot spots on the surface.

The coldest water temperature achieved was near the ambient air temperature. If water with a temperature lower than the dewpoint for that night was used, condensation would occur on the surface. This would have made it difficult to get an accurate emissivity measurement as well as an accurate contact temperature measurement during the collect. One box from each of the surfaces (Krylon Ultra Flat Black #1602 spray paint, sandblasted, and polished) was filled with water 12 hours prior to the collect and left outside until the collect took place. Another set of three boxes were filled with the hottest tap water available (46°C) and the final set of three boxes was filled with water at an intermediate temperature (33°C). This experiment produced imagery containing nine combinations of target temperature/background temperature and target emissivity/background emissivity differences.

4.7.1.2 Location

The parking lots on RIT's campus provided a uniform background material large enough such that homogeneous pixels surrounded the target location. The targets were laid out in the parking lot at the location marked with a green "x" on the left side of Figure 4.12. It provided a large, uniform area of asphalt rarely used during the summer months. Two temperature-controlled calibration targets were placed on the roof of the Chester F. Carlson Center for Imaging Science, denoted by the yellow star in Figure 4.12




Figure 4.12: Image from Google Maps showing the experiment locations that have been tested for suitability. Location 1 is marked by red circles, Location 2 by blue circles, and Location 3 is denoted by the green “x”.

4.7.1.3 Test plan

The boxes were laid out in a 3x3 grid depicted in Figure 4.13. The boxes were placed on the ground so that their longest sides faced north and south. With a GSD of 1.27ft, there were approximately 8 pixels in the thermal image separating each box from its neighbors. This left enough space between boxes such that the energy contribution due to the sensor and atmospheric MTF from one box will be negligible in the pixel or pixels containing another box.




Figure 4.13: Test target configuration

The boxes were arranged such that those in the leftmost column were the ones painted with Krylon Ultra Flat Black #1602 spray paint. The boxes in the rightmost column were polished aluminum and the middle column contained the sandblasted boxes. If the boxes had not been modified in any other way (*i.e.* the surface temperatures were the same), a thermal imager would see a difference in the radiance coming from the three columns of boxes simply due to their emissivity differences. Ideally there would be an identical 3x3 array of aluminum boxes laid out on a surface other than asphalt so that multiple emissivity differences between the boxes and the background material were observed. After a complete survey, it was determined that no other location on RIT's campus that would allow for that. The possibility of placing layers of another material (*e.g.* plywood) was explored and approximately

1600 square feet (corresponding to a square with dimensions of 40 feet x 40 feet) of material would have to be laid down. The amount of material required was too large to be laid out and secured. Therefore this experiment only had the parking lot asphalt as a background material.

The aluminum boxes were all the same size, so physical target size differences were not present in the experiment. The boxes remained stationary over the course of the experiment. However, by shifting the camera between image captures, different positions were simulated. Shifting the camera also leads to each target occupying different proportions of pixel area. So, even though the target size did not physically change, the amount of pixel area it occupied with each shift of the camera did change due to its position across pixel boundaries.

The boxes were filled with water and placed on the asphalt 30 minutes prior to the plane flying overhead. This gave enough time for the temperature of the box to come to an equilibrium as well as for heat transfer effects to occur. Five minutes prior to the sensor flying overhead a contact temperature measurement of the top and one side of each box was made using the Exergen Infrared Thermometer as well as a measurement of the asphalt at the base of the box, 1in from the base, and 2in from the base. During the collect, a temperature measurement of the top of each box was made between each pass. The boxes were measured in multiple locations (3) on their top surface and an average of the temperature measurements was computed and recorded. The other locations were not measured due to the limited amount of time available between aircraft passes. Immediately after the final pass of WASP, the more extensive suite of measurements of the boxes and asphalt were taken. The temperature of the background asphalt was measured with 10 samples in a location approximately 12 feet from the Eastern edge of the grid of boxes in order to compute an average background temperature.

The reflectance of each box surface type was measured using the Surface Optics Corporation SOC-400T infrared reflectometer and the reflectance was converted to

emissivity via Kirchhoff's Law. Figure 4.14 contains the emissivity curves for each box.




Figure 4.14: Emissivity curves for (a) aluminum box painted with Krylon Ultra Flat Black #1602 spray paint, (b) aluminum box whose surface has been sandblasted, and (c) polished aluminum box.

The emissivity of the asphalt was not measured coincident with the collect due to atmospheric conditions, since scattered clouds covered approximately 80% of the sky. An asphalt emissivity curve (shown in Figure 4.15) was taken from another project and used in the simulations. The weather station on the roof of the Chester F. Carlson Center for Imaging Science was used as the source of surface weather data. Weather data was recorded starting three days prior to the start of the collect

and continuing until the end of the collect. Three days was used because the asphalt was estimated to be approximately 6in thick, if the asphalt had been thicker, a greater period of preceding weather observations would have been required.




Figure 4.15: Asphalt emissivity curve.

Radiosonde data from Buffalo, NY was gathered to use in the MODTRAN/DIRSIG downwelled radiance calculation. The radiosonde was launched approximately 60 miles away from the site and approximately 2.5 hours prior to the collect. The data below the boundary layer was changed to better reflect the local Rochester weather conditions at the time of collect by inserting the weather data from the rooftop weather station and interpolating. The weather data from the rooftop station was validated using weather data from the Rochester International Airport.

4.7.2 Atmospheric compensation

The dataset was atmospherically corrected using an in-scene technique. Two pools filled with water were placed on the roof of the Chester F. Carlson Center for Imaging Science. One of the pools contained cold water and the other contained hot water. To get the water warm, one of the pools was covered with a solar cover 36 hours prior to the collect. The other pool was filled with cold tap water approximately 30 minutes prior to the collect. The pools were separated by approximately 15 feet and were 6ft long, 4ft wide, and 1ft deep, so there was at least one pure water pixel in each WASP image. A thermometer was placed in the pool to measure the water temperature. However, the thermometer measured the bulk water temperature which can deviate significantly from the surface temperature, which is what the sensor “sees”. To account for this, the water in each pool was stirred right until the moment the sensor was overhead. When the sensor flew over, the person at the pool moved away so as not to influence the radiance values and the temperature recorded with each thermometer was recorded for each pass.

The in-scene atmospheric compensation technique used is similar to the Empirical Line Method (ELM). Recall the governing equation equation for a pure pixel in the thermal infrared

$$L_{pix} = [\varepsilon L_{BB} + (1 - \varepsilon)L_d]\tau + L_u \quad (4.5)$$

Since the emissivity of water is very high (0.986) the reflected downwelled term is negligible leaving

$$L_{pool} = L_{BB}\tau + L_u \quad (4.6)$$

where τ is the slope term and L_u is the intercept. Using the hot and cold pool data as calibration points, the atmospheric transmission and upwelled radiance was determined by computing the slope and intercept of the line going through the two points. The pixel radiance from the WASP imagery for both pools was the

dependent variable and the ground leaving blackbody radiance from each pools was the independent variable. The spectral ground-leaving radiance from the water in each pool from 8 - 9.2 μm was computed by applying the spectral response curve of the WASP sensor, $R'(\lambda)$, and integrating over the bandpass,

$$L_{sensor} = \int_8^{9.2} L_{pool} R'(\lambda) dx \quad (4.7)$$

4.7.3 Modeling in MuSES

The asphalt surface was the most difficult object to model in MuSES. The surface temperature of the asphalt was measured at the time of collect. The rooftop weather station recorded the weather information during the collect as well as the weather information three days prior the collect. RIT Facilities Management confirmed the thickness of the asphalt and the gravel substrate to be approximately 4.5in and 10in, respectively. This information, along with extensive help from the experts at ThermoAnalytics, helped to create and validate a 17-layer model of the parking lot, illustrated in Figure 4.16. The multi-layer approached allowed the heat transfer to be computed at finer intervals than if the parking lot had been modeled as one large layer. When the asphalt was modeled as one solid layer, MuSES computed its temperature to be cooler than what was measured. Dividing up the asphalt into more layers provide put less mass on the surface node and helped it to heat up more in the sun.

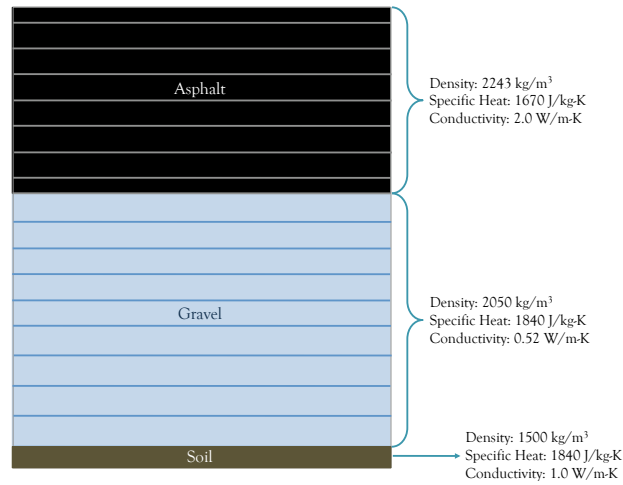


Figure 4.16: Schematic of 17-layer asphalt model in MuSES.

When MuSES was run with all available input data in this configuration, the temperature of the surface of the asphalt was 29.8°C. This was the same value as the average of the 10 temperature samples taken from the parking lot. That is not to say that one should always model an asphalt parking lot as a series of 17 layers. If the material properties of the asphalt (*e.g.* specific heat, density, conduction coefficient) had been measured instead of estimated, the input parameters to the model would have been more accurate. This could have decreased the need for so many layers to increase the modeled surface temperature of the asphalt.

Each box, shown in Figure 4.17, was modeled using one layer that was 3mm thick. The top surface had an emissivity assigned that was indicative of the surface treatment applied (*e.g.* Polished aluminum). Since the temperature of the box remained fixed, the density, specific heat, and conductivity were not needed.




Figure 4.17: MuSES model of aluminum box.

The MuSES weather file contained three days worth of weather history but the box was only out on the asphalt for approximately two hours in total. Therefore two separate TDFs had to be created and linked together. The first file had only the background asphalt object exposed to the environment and the weather file contained data up to the time that the box was placed on the surface. The second TDF had the box placed on top of the asphalt and its weather file contained the weather data for the two hours during the collect period. Because the two TDFs were linked together, the asphalt temperature in the second file was based off of the weather history from the first file.

4.8 Error analysis

The parameter estimates used as inputs to the physics-based modeling processes contain some degree of error. These individual errors propagate through the process and combine to produce an error in the extracted target temperature. There is no governing equation for the entire process, so the error analysis for the physics-based portion is carried out using simulated data and Monte Carlo methods. The MuSES simulations are the first step in the methodology, therefore the description of the error analysis will begin there.

MuSES is not used to model the target temperature as it is fixed during each of the simulations. Therefore the error contained in the MuSES input parameters only

affects the background temperature estimate. The equation of partial differentials that computes the error in the background temperature estimate is

$$s_{Tb} = \left[\left(\frac{\partial T_b}{\partial T_{air}} s_{Tair} \right)^2 + \left(\frac{\partial T_b}{\partial WS} s_{WS} \right)^2 + \left(\frac{\partial T_b}{\partial CC} s_{CC} \right)^2 + \left(\frac{\partial T_b}{\partial D} s_D \right)^2 + \left(\frac{\partial T_b}{\partial SH} s_{SH} \right)^2 + \left(\frac{\partial T_b}{\partial C} s_C \right)^2 \right]^{\frac{1}{2}} \quad (4.8)$$

Table 4.4 describes the error sources corresponding to each abbreviation in Equation 4.8.

Table 4.4: Abbreviations from Equation 4.8

Abbreviation	Corresponding error source
s_{Tair}	Air temperature [°C]
s_{WS}	Wind speed [m/s]
s_{CC}	Cloud cover [% of sky]
s_D	Density of material [kg/m ³]
s_{SH}	Specific heat of material [J/kg - K]
s_C	Conductivity of material [W/m-K]

The parameters in Table 4.4 are varied, one at a time, while the others are held fixed. For each variation, the background temperature is recorded to produce a plot of the background temperature versus the parameter values. The minimum and maximum values of the parameter that are likely to occur are determined and used as the extreme points in the plot. The midpoint between the minimum and maximum is also used. From these three points, one can determine whether or not the relationship between the parameter and the modeled background temperature is linear, or approximately linear ($R^2 \geq 0.9$), over that range. If the relationship is linear then the slope of the best-fit line of the points is the value of the partial differential for that specific parameter in Equation 4.8. If the relationship cannot be described as linear, or approximately linear, then more points between the maximum and minimum need to be computed and the slope between each set

of points is used in a piece-wise fashion for the partial differentials. The values of the s terms comes from the user's best estimate of how much error is contained in each parameter. For example, if the user was confident in their estimate of the air temperature they might set $s_{T_{air}}$ to 0.2°C .

The next step in the error analysis is to use DIRSIG to compute the radiance solution for each of the TDFs from MuSES. The equation of partial differentials that computes the error in the average radiance value of the target-containing pixel, $s_{L_{tcp}}$, is

$$s_{L_{tcp}} = \left[\left(\frac{\partial L_{tcp}}{\partial x} s_x \right)^2 + \left(\frac{\partial L_{tcp}}{\partial y} s_y \right)^2 + \left(\frac{\partial L_{tcp}}{\partial z} s_z \right)^2 + \left(\frac{\partial L_{tcp}}{\partial loc_x} s_{locx} \right)^2 + \left(\frac{\partial L_{tcp}}{\partial loc_y} s_{locy} \right)^2 + \left(\frac{\partial L_{tcp}}{\partial T_b} s_{Tb} \right)^2 + \left(\frac{\partial L_{tcp}}{\partial \varepsilon_b} s_{\varepsilon b} \right)^2 + \left(\frac{\partial L_{tcp}}{\partial \varepsilon_t} s_{\varepsilon t} \right)^2 + \left(\frac{\partial L_{tcp}}{\partial \tau} s_{\tau} \right)^2 + \left(\frac{\partial L_{tcp}}{\partial L_d} s_{Ld} \right)^2 \right]^{\frac{1}{2}} \quad (4.9)$$

Table 4.5 contains the explanation of the error source corresponding to each abbreviation in Equation 4.9.

Table 4.5: Abbreviations from Equation 4.9

Abbreviation	Corresponding error source
s_x, s_y, s_z	Target size in x, y, and z direction [m]
s_{locx}, s_{locy}	Target location in x and y direction [m]
s_{Tb}	Background temperature [$^\circ\text{C}$]
$s_{\varepsilon b}$	Background emissivity
$s_{\varepsilon t}$	Target emissivity
s_{τ}	Atmospheric transmission
s_{Ld}	Downwelled radiance [$\text{Wm}^{-2}\text{sr}^{-1}$]

The same methodology to determine the values of the partials used previously is used in this step as well. The sixth term in the equation represents the error in

the background temperature. That partial is computed by varying the background temperature in DIRSIG and recording how it affects the average radiance value. The value of s_{Tb} is computed from Equation 4.8. The emissivity curves are biased by the same amount over all wavelengths. If the maximum emissivity value likely to occur is 0.05 larger than the emissivity curve used in DIRSIG, 0.05 is added to each emissivity value in the curve. This biasing method is also employed when computing the partials for the atmospheric transmission and the downwelled radiance. The error from the upwelled radiance is not included in the analysis because the difference between two radiance quantities is computed in the next step. This subtraction of the two governing equations gets rid of any error from the upwelled radiance term.

The difference between the average radiance from the target-containing pixels (L_{tcp}) and the radiance from the pure background pixels (L_b) is computed in the next step of the methodology.

$$\Delta L_{avg} = L_{tcp} - L_b \quad (4.10)$$

Equation 4.10 is a simple governing equation, therefore Monte Carlo methods are not employed in this part of the error analysis. The values of the partial derivatives can be computed directly from Equation 4.10. The equation to compute the error in ΔL_{avg} is

$$s_{\Delta L_{avg}} = \left[\left(\frac{\partial \Delta L_{avg}}{\partial L_{tcp}} s_{L_{tcp}} \right)^2 + \left(\frac{\partial \Delta L_{avg}}{\partial L_b} s_{L_b} \right)^2 \right]^{\frac{1}{2}} \quad (4.11)$$

The value of the partial derivative in the first term is 1 and the value of the partial derivative in the second term is -1, leaving the following equation to compute the error in the delta radiance term

$$s_{\Delta L_{avg}} = [s_{L_{tcp}}^2 + (-s_{L_b})^2]^{\frac{1}{2}} \quad (4.12)$$

The value for the error in the background radiance is the calibration error of the

sensor and s_{Ltcp} comes from Equation 4.9.

A regression model (in this example it is a second-order model) is used to determine the difference between the target and background temperatures

$$\Delta T = \beta_0 + \beta_1 \Delta L_{avg} + \beta_2 \Delta L_{avg}^2 \quad (4.13)$$

and the imagery-derived background temperature is added to ΔT to arrive at the target temperature. This can be expressed as

$$T_t = \Delta T + T_b = \beta_0 + \beta_1 \Delta L_{avg} + \beta_2 \Delta L_{avg}^2 + T_b \quad (4.14)$$

Again, since there is an equation that describes the process, Monte Carlo methods are not used in the error analysis for this step. The equation of partial derivatives is

$$s_{Tt} = \left[\left(\frac{\partial T_t}{\partial \beta_0} s_{\beta_0} \right)^2 + \left(\frac{\partial T_t}{\partial \beta_1} s_{\beta_1} \right)^2 + \left(\frac{\partial T_t}{\partial \beta_2} s_{\beta_2} \right)^2 + \left(\frac{\partial T_t}{\partial \Delta L_{avg}} s_{\Delta L_{avg}} \right)^2 + \left(\frac{\partial T_t}{\partial T_b} s_{T_b} \right)^2 \right]^{\frac{1}{2}} \quad (4.15)$$

After computing the partials, Equation 4.15 is reduced to

$$s_{Tt} = \left[s_{\beta_0}^2 + (\Delta L_{avg} s_{\beta_1})^2 + (\Delta L_{avg}^2 s_{\beta_2})^2 + ((\beta_1 + 2\beta_2) s_{\Delta L_{avg}})^2 \right]^{\frac{1}{2}} \quad (4.16)$$

where s_{β_0} , s_{β_1} , and s_{β_2} are the standard error in the regression coefficients.

4.9 Synthetic Experiment

The synthetic experiment simulated the airborne collect using imagery created entirely in MuSES and DIRSIG. The primary reason for this study was to show how temperature extraction errors improve with better atmospheric knowledge. Recall that the downwelled radiance was computed using atmospheric information from a radiosonde that was neither launched in the same place as the collect nor at the same time. Also, imprecise knowledge of the cloud cover distribution was available. This almost certainly leads to errors in the downwelled radiance computation which will have a large effect on the sandblasted and polished boxes because of their low emissivities. A secondary function of this experiment was to look at target-background combinations with larger temperature differences than were achieved during the airborne collect. The largest temperature difference between target and background in the airborne experiment was 5°C , whereas objects such as operating electrical transformers produce a much larger temperature difference. Even though the hottest water used had a temperature of around 46°C , the largest temperature difference was 5°C because the boxes were placed on the asphalt for 30 minutes prior to the airplane flying overhead, which allowed for heat transfer to occur. The heat transfer from the box to the asphalt and the environment resulted in the decreased temperature difference between the target and background. Finally, in the airborne collect the box emissivities fell at the extreme ends of the spectrum. This experiment allowed for a more moderate emissivity value of 0.8 to be tested.

Forty truth cases were created using all combinations of the parameters in Table 4.6. All of the truth cases used the asphalt emissivity curve from Figure 4.15 as the emissivity of the background.

Table 4.6: Synthetic truth image inputs. The first three target emissivity descriptions match the emissivity curves found in Figures 4.14(a), 4.14(b), 4.14(c), respectively.

Parameter	Value
<i>Target temperature</i> [°C]	82.6
	62.6
	32.6
<i>Target emissivity</i>	Krylon Ultra Flat Black
	Sandblasted Aluminum
	Polished Aluminum
	0.8 Graybody
<i>Position</i>	Center of center pixel
	Across boundary of four adjacent pixels

The position of the target when it was located across four pixels was set to be (1.87, 1.77) in DIRSIG. This means that the center of the target was translated 1.87m in the x-direction and 1.77m in the y-direction. Heat transfer was modeled using MuSES as described in Section 4.7.3. The downwelling radiance was computed using the Buffalo radiosonde data. Each TDF was imported into DIRSIG, where the downwelling radiance was applied and the radiance images were produced. The radiance from the target-containing pixels was recorded to be used later in the process as an input to the regression equations derived from their respective LUTs.

One LUT was created per truth case. The input parameters to the LUT images were the input parameters to the truth cases plus randomly assigned error. The errors are summarized in Table 4.7.

Table 4.7: Error sources for the synthetic image study.

Parameter	Error amount
Background temperature	4%
Background emissivity	2%
Target emissivity	3%
Position	1/4 pixel
Atmosphere	none

The truth values for the background temperature, background emissivity, and target emissivity were biased by the error amount in Table 4.7. The true asphalt temperature was 29.9°C and the average imagery-derived asphalt temperature was 28.6°C. This corresponded to the imager-derived asphalt temperature being 4% lower than the true asphalt temperature. Therefore a 4% error for the background temperature was used in the synthetic experiment. The derivations of the true asphalt temperature and the image-derived temperatures are described in more detail in the Results chapter. The emissivity curve for the background was biased upward by 2% and the emissivity curves for the target emissivities were biased by 3%. To add error to the position, the center of the target was moved 1/4 of a GSD diagonally. Since the main function of this study was determine the improvement that would occur with more accurate knowledge of the atmosphere, no error was added to these values. This represents the best case scenario where the atmosphere is characterized perfectly.

4.10 MuSES Utility Study

Heat transfer is present between all objects in the real world. The methodology described in the beginning of this chapter incorporates heat transfer into the modeling process. One could ask whether or not modeling the heat transfer actually results in a more accurate temperature extraction in the end. If not, one entire step of the process can be omitted. This study used synthetic imagery to determine the importance of modeling heat transfer when using this methodology to determine the temperature of a subpixel object.

Six surrogate truth cases were created and the parameters used are summarized in Table 4.8. Each case used the same simple box geometry with the emissivity curve of Krylon Ultra Flat Black. The same asphalt model and emissivity curve used in previous studies was used again here.. The asphalt was modeled in MuSES the same way as before, where the weather file and material properties dictated its temperature. The boxes were fixed at their respective temperatures.

Table 4.8: Input parameters for the six truth cases of the heat transfer utility study.

Case	<i>Background Temperature</i>	<i>Target Temperature</i>	<i>Target Emissivity</i>	<i>Background Emissivity</i>	<i>Position</i>
1	29.8 °C	82.6 °C	0.93	0.95	Center
2	29.8 °C	62.6 °C	0.93	0.95	Center
3	29.8 °C	32.6 °C	0.93	0.95	Center
4	29.8 °C	82.6 °C	0.93	0.95	Across 4 pixels
5	29.8 °C	62.6 °C	0.93	0.95	Across 4 pixels
6	29.8 °C	32.6 °C	0.93	0.95	Across 4 pixels

The target temperatures for cases 1 and 4 were 60°C above the ambient air temperature, cases 2 and 5 were 40°C above ambient air temperature, and cases 3 and 6 were 10°C above ambient air temperature. The background temperature was 29.8°C, so the target temperature in cases 1 and 4 was 52.8°C above the asphalt temperature. The target temperature for cases 2 and 5 was 32.8°C above the asphalt temperature and cases 3 and 6 were 2.8°C above the asphalt temperature.

The truth cases were modeled in MuSES to produce a TDF. The TDF for each case was then run through DIRSIG to produce a radiance image. The position labels from Table 4.8 refer to the position of the target in DIRSIG. The center position refers to the target being located in the geometric middle of the center pixel in the image. When the target was across four pixels, it was shifted relative to the center pixel such that a portion of the target was physically located within four adjacent pixels. The atmosphere file containing clouds was used to compute the downwelling radiance.

The input parameters to the LUTs used to extract the target temperature were biased from the truth values (found in Table 4.8) by the amount shown in Table 4.9.

Table 4.9: Error sources for heat transfer study.

Parameter	Error amount
Background temperature	4%
Background emissivity	2%
Target emissivity	3%
Position	1/4 pixel
Atmosphere	none

The data collected was used to test the methodology outlined in this chapter. A study to determine the how large the temperature extraction error would be if the heat transfer were not modeled was also conducted using this data. Two lookup tables were created, one that had MuSES data built in and one that did not, and a regression equation was built from each. The average pixel radiance value across the target-containing pixels was used as the input to both equations and the temperature extraction error was compared. If the temperature extraction errors produced by the heat transfer LUT were the same as those produced by not modeling heat transfer, that would indicate that modeling heat transfer could be removed from the temperature extraction process.

Chapter 5

Results

5.1 Airborne collect results

5.1.1 Airborne data

The methodology described in Chapter 4 was applied to three of the images collected during the airborne campaign. Two images were collected during the North-South passes and were acquired at the beginning of the experiment. The third image was captured during the East-West pass at the end of the mission. The images were identified by their WASP image numbers and Table 5.1 shows the image ID with the collection information.


Table 5.1: Image ID numbers with corresponding collection information from 5 August 2011.

ID #	Pass direction	Time collected (EDT)
71	North-South	2230
148	North-South	2251
185	East-West	2301

The images are shown in Figures 5.1, 5.2, and 5.3 with a companion image chip that highlights the boxes.



(a) Image 71




(b) Chip of image 71

Figure 5.1: Image 71 with corresponding image chip highlighting the experiment boxes.



(a) Image 148



(b) Chip of image '48

Figure 5.2: Image 148 with corresponding image chip highlighting the experiment boxes.

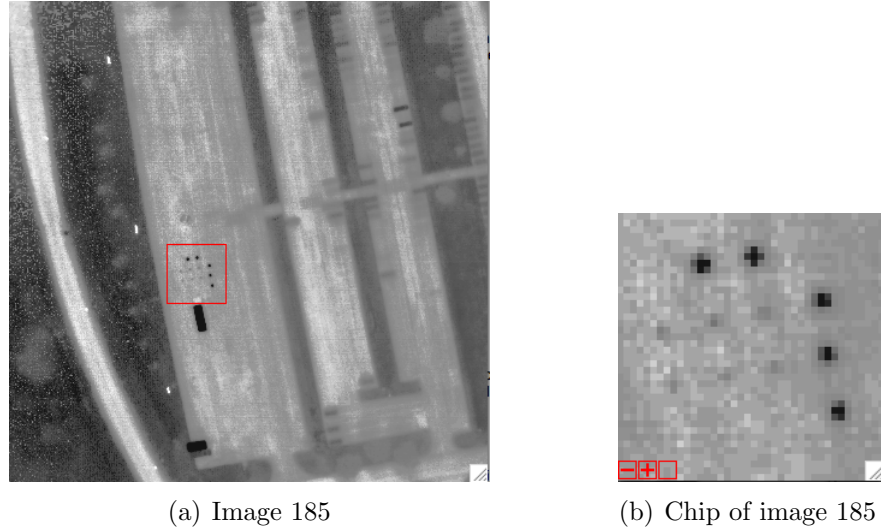


Figure 5.3: Image 185 with corresponding image chip highlighting the experiment boxes.

The image chips from Figures 5.1(b), 5.2(b), and 5.3(b) illustrate the difficulty of the task at hand, as it is difficult to even locate the boxes in the image. The Krylon flat black painted boxes containing the warm water and the ambient air temperature water were not found in any of the images. This was likely due to the temperature and/or the emissivity of those particular boxes being similar to the asphalt. Figure 5.4 shows the image chips with orange squares around each of the boxes that were found in the scene.





Figure 5.4: Image chips with the boxes from the experiment highlighted by orange boxes. (Note: not all of the boxes were accounted for)

Images of the calibration pools were also collected. One of these images is shown Figure 5.5. The FOV of WASP was not large enough to contain both the pools and the boxes in one image, so the images used for atmospheric compensation were taken prior to the images of the boxes. For the North-South passes, the time difference was approximately two minutes and for the East-West passes the box images were taken immediately after the pool images.



(a) Pool Image



(b) Image chip highlighting the pools

Figure 5.5: Image of the calibration pools on the roof of the Chester F. Carlson Center for Imaging Science and its companion image chip.

5.1.2 Ground truth

After each pass of the plane, the temperature of the top of each box was measured with an Exergen Infrared Thermometer and recorded. Table 5.2 contains the temperatures of the boxes in images 71, 148, and 185.

Table 5.2: Box surface temperatures measured after each pass for three images.

Box description	Image 71 temperatures [°C]	Image 148 temperatures [°C]	Image 185 temperatures [°C]
Flat black painted (hot)	31.3	30.3	29.9
Sandblasted (hot)	30.6	30.1	29.6
Polished (hot)	27.3	28.2	28.5
Flat black painted (warm)	28.0	26.9	26.7
Sandblasted (warm)	27.4	28	27.6
Polished (warm)	27.0	26.8	27.3
Flat black painted (ambient)	27.6	27.4	26.6
Sandblasted (ambient)	27.1	26.6	26.9
Polished (ambient)	25.6	26.4	26.2

At the end of the mission, the temperature of the asphalt was sampled in ten different locations. These measurements were taken approximately 30ft away from the site of the collect to ensure that neither the experiment nor the experiment team had influenced the surface temperatures. They were also taken in the thoroughfare of the parking lot (as opposed to the parking spaces) to ensure that shading of the surface by a car for an extended period of time or surface contaminants (*e.g.* oil) were not present. The ten measurements were averaged to obtain the background asphalt temperature. Table 5.3 contains the temperature measurements for the asphalt.

Table 5.3: Asphalt temperature measurements.

Measurement #	Temperature [$^{\circ}\text{C}$]
1	30.0
2	30.1
3	29.8
4	29.8
5	30.0
6	29.8
7	30.2
8	29.9
9	29.8
10	29.9
Average temperature	29.9

The final piece of ground truth information collected was the calibration pool temperatures. A temperature measurement from each of the pools was collected immediately after the plane passed overhead and a listing can be found in Table 5.4. These measurements, and the associated blackbody radiance values for each, were used to remove the atmospheric transmission and upwelled radiance terms from the pixel radiance values in the images.

Table 5.4: Calibration pool temperatures.

Time	Hot pool temperature [$^{\circ}\text{C}$]	Cold pool temperature [$^{\circ}\text{C}$]
2228	33.5	22.5
2249	33.3	22.2
2301	32.9	22.1

5.1.3 Background temperature estimation

The method described in Section 2.12 was used to estimate the background temperature. A region of interest, with approximately 200 pixels, was selected in the parking lot away from the experimental site and the average asphalt radiance was computed for each image. Using knowledge of the downwelled radiance and the background emissivity, the ground-leaving radiance for the background was com-

puted. An iterative process was used to back out the asphalt temperature from the Planck equation. The background temperature for each image is given in Table 5.5.

Table 5.5: Extracted background temperature and the temperature error for each of the images.

Image #	Background temperature [°C]	Temperature error [°C]
71	28.8	1.76
148	28.7	1.55
185	28.6	1.33

5.1.4 Target temperature extraction results

The weather data for Rochester from Weather Underground [32] indicated scattered clouds covering approximately 80% of the sky at the time of the collect. The data from the Rochester International Airport reported the presence of broken clouds, meaning over 50% of the sky was covered. When MODTRAN incorporates clouds into a model, it assumes a solid layer of cloud cover. It cannot directly model broken cloud cover. Since both weather sources indicated the presence of broken clouds covering more than half of the sky area, it was decided that the atmosphere would be modeled with and without cloud cover. The cloudy atmosphere was modeled as an altostratus cloud layer ($ICLD = 2$ in MODTRAN “tape5” file). Altostratus was chosen because the altitude of the cloud base was above 2.4km. In both of the atmospheric models, the corrected radiosonde data was imported into the “tape5” file. The temperature extraction results from the cloudy atmosphere for the three images are shown in Figures 5.6, 5.7, and 5.8.




Figure 5.6: Temperature extraction results for the boxes in image 71 (captured at 2230).

The sandblasted and polished boxes were easy to locate within the images because their emissivities were so much lower than that of the asphalt. The painted boxes were challenging to find, so much so that the painted boxes filled with warm water and water at the ambient air temperature could not be located in any of the images. Those particular boxes are labeled with an N/A in the Figures 5.6, 5.7, and 5.8.

The smallest temperature error in image 71 came from the painted box, which had the highest emissivity. The polished boxes produced very large errors (greater than 200°C). It is believed that such large temperature errors resulted from two issues. The primary issue was the lack of quality atmospheric knowledge. The high emissivity of the painted surface prevents the reflected radiance from the atmosphere from making a large contribution to the ground-leaving radiance of the box, which is why the temperature errors for the Krylon flat black painted boxes were small,

even though the errors for the other boxes were large. Consider a scenario where the target has a temperature of 29°C and the sky has a temperature of -13°C . For a box painted with Krylon Ultra Flat Black paint, 96% of the ground-leaving radiance is due to self-emission. This means that the atmosphere only contributes 4% of the energy to the ground-leaving radiance. One can see how even a large error in atmospheric knowledge would not translate into a large temperature extraction error due to the reflected component's small contribution to the ground-leaving radiance. Using the previous example, a 10% error in the downwelled radiance knowledge translates to an error in pixel radiance of $0.06 \text{ W}/(\text{m}^2\text{sr})$. This radiance error corresponds to a blackbody temperature error of 0.35°C . If the case of the polished box is considered under the same assumptions, the reflected downwelled component contributes 95% of the ground-leaving radiance from the box whereas the self-emitted component is only a 5% contributor. With such a large effect on the ground-leaving radiance, even a small error in the knowledge of the atmosphere can manifest as a large error in temperature.

The secondary contributor to the large extraction errors of the sandblasted and polished boxes was the poor knowledge of the atmosphere resulting in delta radiance values that were outside the range modeled by the regression equation. The target temperature range modeled by MuSES and DIRSIG was 20°C to 40°C so the true box temperatures fell in the middle of the modeled range. The background temperature of the asphalt used in the modeling process was approximately 29°C , therefore the largest temperature difference between the target and background that was modeled was 11°C . Any temperature error greater than 11°C generated by the regression model was a result of extrapolation and therefore the magnitude of the error might be far less accurate.




Figure 5.7: Temperature extraction results for the boxes in image 148 (captured at 2251).

The results from Images 148 and 185 follow the same pattern as those from Image 71. With the exception of the first data point in each image (the Krylon flat black painted box with the warmest water), all of the errors are very large. This means that the extracted temperatures were a result of imprecise atmospheric knowledge and spatial errors due to broken cloud formations or extrapolation and therefore the magnitudes are not accurate.




Figure 5.8: Temperature extraction results for the boxes in image 185 (captured at 2301).

The only temperature errors that were valid were the ones belonging to the high emissivity box. The largest temperature error was 7°C, which is two degrees larger than the land temperature errors predicted from Szymanski *et al.* Szymanski looked at mixed pixels including water, and did not include an emissivity error on the background in their analysis. The results from this research did have error in the background emissivity.

The results were recomputed with an atmosphere that did not have clouds. The results are found in Figures 5.9, 5.10, and 5.11.




Figure 5.9: Temperature extraction results for the boxes in image 71 (captured at 2230).

The most obvious feature to note is the dramatic decrease in the temperature extraction errors in image 71 when the cloud-free atmosphere is used. This indicates that the clear atmosphere may be a more appropriate model for the conditions that the boxes were exposed to as compared to an atmosphere with a solid cloud layer. Since the cloud cover was broken the night of the experiment, and the portion of the sky the boxes reflected was so small, it is possible that the reflected radiance from the boxes originated in a cloud-free portion of the sky.

With the exception of the polished box containing warm water, which had a temperature error twice as large as the next largest error, the temperature errors of the boxes were less than 13°C. The polished box containing warm water was probably an anomalous result (perhaps due to sensor artifacts) because the average pixel radiance from that particular box was lower than the polished box filled with water

at a lower temperature. Since the boxes had the same emissivity, the box with the warmer water should have produced a higher radiance value. This could have been a result from a number of factors including; an error in the ground truth temperatures, an error in the atmospheric correction that significantly affected this target radiance value. A more likely source of error is the positioning algorithm selecting the wrong target-containing pixels for this box.




Figure 5.10: Temperature extraction results for the boxes in image 148 (captured at 2251).

The temperature errors from image 148 were generally larger than those from image 71, however, they were still much smaller than the errors from image 148 when clouds were modeled in the atmosphere. The reflected radiance from the boxes was a combination of sky and clouds, and the ratio of this combination almost certainly changed as the experiment progressed. In image 148 the boxes might have reflected a larger portion of cloud radiance than in image 71, so the clear atmosphere was

not as appropriate of a model. The temperature errors from image 185 followed the same trend as the other cloud-free atmosphere results. The temperature errors in Figure 5.11 decreased significantly compared to those in Figure 5.8.




Figure 5.11: Temperature extraction results for the boxes in image 185 (captured at 2301).

The temperature extraction errors from the cloud-free atmospheres were smaller than those from the atmosphere where a layer of altocirrus clouds were modeled. Modifying the atmospheric model such that the altocirrus clouds were removed had the most significant effect on the temperature extraction for the sandblasted and polished boxes. These two surfaces had very low emissivities (0.2 and 0.02, respectively) and reflected almost all of the downwelled radiance incident upon them, thereby making them sensitive to errors in atmospheric knowledge during the modeling process. This is further illustrated by the polished boxes consistently producing the largest temperature extraction errors and, with a few exceptions, the sandblasted

boxes produced the second largest temperature extraction errors.

Individual extraction results are dominated by the error in the system. So, while it is easy to discuss the general trends that are common to all of the analyzed images, it is difficult to talk about the errors for individual boxes because of all of the uncertainty within the experiment. The atmosphere played a significant role in the errors, especially with the low emissivity surfaces. The boxes painted with Krylon Ultra Flat Black showed only a few degree improvement (which is significant but small compared to the improvement from the sandblasted and polished boxes) in temperature error between the temperatures extracted using a cloudy atmosphere and the temperatures from a clear atmosphere. This makes sense because the high emissivity of the painted surface prevents the reflected radiance from the atmosphere from making a large contribution to the ground-leaving radiance of the box. Given the polished and sandblasted boxes sensitivity to the atmosphere, it is easy to see how there could be large extracted temperature fluctuations between images as the atmospheric conditions changed from scene to scene.

Since the boxes only took up 38% of the GSD, the background radiance had a significant effect on the total pixel radiance. Therefore, errors in the background temperature estimation as well as the background emissivity would have a significant effect on the pixel radiances that were modeled. The error in the estimated background temperature and emissivity parameters that was used in the modeling process certainly contributed to the temperature extraction errors and also could be a reason why the extraction errors changed with the scenes.

The errors for the painted boxes were in line with the temperature errors found in the literature review. However, all of the methods found in the literature review used multiple thermal bands to determine the temperature of a subpixel object. The other methods did not include low emissivity targets so a comparison cannot be made in that realm, but for the most part the temperature errors from the sandblasted box were less than 10°C. This result appears to be good given the lack

of atmospheric knowledge and the low emissivity of the surface.

5.1.5 Error analysis

The error analysis was carried out as described in Section 3.8. First, the partial derivatives for the MuSES input parameters were computed. Recall the terms in Equation 3.6 were air temperature, wind speed, cloud coverage, density, specific heat, and conductivity. Table 5.6 shows the parameter, the range of values used in the analysis, and the partial derivative values.

Table 5.6: Range of values used and the error term value computed for each MuSES input parameter. The error term values represent the partial derivatives multiplied by their respective parameter error estimate.

Parameter	Value range	Error term value
Air temperature	$22.5^{\circ}\text{C} \pm 3^{\circ}\text{C}$	0.17°C
Wind speed	$2.1\text{m/s} \pm 5\text{m/s}$	-0.54°C
Cloud coverage	$84\% \pm 10\%$	-0.004°C
Density	$2250 \pm 150\text{kg/m}^3$	0.09°C
Specific heat	$1500\text{J/kg-K} \pm 500\text{J/kg-K}$	0.3°C
Conductivity	$1.5 \pm 0.75\text{W/m-K}$	1.17°C

The conductivity had the largest error term value and therefore it has the largest effect on the background temperature modeled in MuSES. The value used for the conductivity of the asphalt was found in the CRC Handbook of Tables for Applied Engineering Science (Second Edition) [33]. The range of values was found by looking through Google search results. Without directly measuring the conductivity of the asphalt, the value used was only a best estimate and therefore subject to error.

The best estimates of the error terms (the “s” terms in Equations 3.6 and 3.7) for the material properties were 100kg/m^3 for the density, 0.5W/m-K for conductivity, and 200J/kg-K for specific heat. The best estimates for errors in environmental condition were 0.2°C for the air temperature, 1m/s for the wind speed, and 10% for

the cloud coverage.

The next step in performing the error analysis was to look at how the input DIRSIG parameters affected the modeled pixel radiance. The DIRSIG input parameters, their values, and the partial derivatives computed are found in Table 5.7.

Table 5.7: Range of values used and the error term values computed for each DIRSIG input parameter. The error term values represent the partial derivatives multiplied by their respective parameter error estimate.

Parameter	Value range	Error term value
Background temperature	$28.7^{\circ}\text{C} \pm 2^{\circ}\text{C}$	$0.14 \text{ [W/(m}^2\text{sr)]}$
Background emissivity	emissivity curve ± 0.02	$0.05 \text{ [W/(m}^2\text{sr)]}$
Target emissivity	emissivity curve ± 0.03	$0.02 \text{ [W/(m}^2\text{sr)]}$
Size (X dimension)	$0.2\text{m} \pm 0.06\text{m}$	$0.04 \text{ [W/(m}^2\text{sr)]}$
Size (Y dimension)	$0.28\text{m} \pm 0.06\text{m}$	$0.09 \text{ [W/(m}^2\text{sr)]}$
Size (Z dimension)	$0.08\text{m} \pm 0.06\text{m}$	$0.03 \text{ [W/(m}^2\text{sr)]}$
X location	starting point + 1/4, 1/2 pixel	$0.098 \text{ [W/(m}^2\text{sr)]}$
Y location	starting point + 1/4, 1/2 pixel	$0.05 \text{ [W/(m}^2\text{sr)]}$
Atmospheric transmission	MODTRAN output $\pm 10\%$	$-0.004 \text{ [W/(m}^2\text{sr)]}$
Downwelled radiance	make adb output $\pm 20\%$	$0.018 \text{ [W/(m}^2\text{sr)]}$ for painted $0.13 \text{ [W/(m}^2\text{sr)]}$ for polished

The largest error term value from the background temperature. This means that a change in the background emissivity in DIRSIG had the largest effect on the pixel radiance value. Since over 60% of the pixel area was the background material, it makes sense that an error in the background emissivity would make a larger contribution to the error compared to the target emissivity. The error term for the downwelled radiance was computed for the case with a box painted with Krylon and a polished box, yielding a very different partial derivative for both. Since the box with Krylon Ultra Flat black had a high emissivity, errors in the downwelled radiance term have a smaller effect when compared to the box with the low emissivity polished surface. The partial derivative value for the position terms was large because the pixel radiance value changes drastically as the target leaves the projected area of the pixel.

Ground truth information shows that the extracted background temperature error was between 1.3°C and 1.76°C, so the values in Table 5.5 were used as the best estimates of the error. The background and target emissivities were estimated to be 0.02 units off from the true value. The targets' physical dimensions were measured beforehand, so no error was included for the target size. The target position was estimated to 0.1m, which was about a quarter of a GSD in each direction. The atmospheric transmission error was estimated to contain a 5% error from the truth and the downwelled radiance was estimated to contain a 15% error.

Finally, Equation 3.13 was used to compute the overall error in the target temperature extraction. The standard error of each regression coefficient for every box was computed. The input delta radiance values to Equation 3.13 came from the real image. The errors from each of the image are found in Table 5.8.

Table 5.8: Temperatures and their errors from the airborne experiment for the atmosphere modeled without clouds.

Box	Image 71 [°C]	Image 148 [°C]	Image 185 [°C]
Krylon hot	35.9 ± 0.12	29.9 ± 0.45	35.59 ± 0.26
Sandblasted hot	30.3 ± 2.03	23.58 ± 2.55	29.32 ± 2.23
Polished hot	19.71 ± 13.2	20.44 ± 4.03	5.28 ± 12.2
Sandblasted warm	21.38 ± 2.35	13.2 ± 9.98	19.7 ± 1.99
Polished warm	-8.41 ± 5.03	1.49 ± 8.5	13.7 ± 10.92
Sandblasted ambient	14.13 ± 6.2	11 ± 3.25	16.82 ± 3.05
Polished ambient	19.57 ± 13.35	9.57 ± 8.48	3.51 ± 11.56

The values of the errors in Table 5.8 are small compared to the difference between the actual and extracted temperatures for the sandblasted and polished boxes. This indicates that some of the estimates of the error terms are too optimistic. The main driver of error in the low emissivity box cases is the downwelled radiance term. The estimate of 15% might be too low although it is difficult to say exactly how much error there actually is in that particular term. It should also be noted that instrument error was not included in this analysis. The thermal infrared imager

on WASP surely adds noise to the imagery it collects and therefore the radiance values contain some amount of error due to the sensor itself. However, the noise characteristic of the WASP thermal imager were not characterized for this research.

5.1.6 Lessons learned

As with any collection/ground truth effort, there were a number of lessons learned from this airborne collect. If this experiment could be done over, the following changes would be made:

Calibration targets

One aspect of the experimental design that made the data analysis difficult was the calibration method. Imagery of the calibration pools used showed that there were only 1-2 pure water pixels. This made it difficult to accurately determine the radiance of the water pixels due to sensor degradations due to optical performance. Larger calibration pools would solve that problem. Another issue had to do with mixing the water. Two volunteers stirred the water up until the moment they judged the plane to be nearby. Then, as the plane flew over and imaged, the volunteers stepped away from the pool. This allowed the water to settle and form a skin whose temperature was not the same as the bulk temperature that was measured. A better way would be to install a mixing system in the pools so that the water was constantly churning, even as the plane flew overhead. Finally, a solution that addresses both of the issues with the pools would be to purchase some large, flat, temperature controlled panels.

Time of year

The experiment was conducted during the warmest part of the year in Rochester. Therefore the temperature of the asphalt was high, even at night, resulting in a reduction in the temperature contrast between the targets and the background that could be achieved with the method employed. This not only created a harder temperature extraction problem, it did not allow for the experiment to

emulate a situation where there is a large temperature contrast between target and background (*e.g.* an operating transformer on a concrete background).

Boxes

Perhaps the largest disappointment of the experiment was the inability to heat the surface of the boxes to more than 34°C. The water from the tap was very hot, but the amount of water that could fit in the boxes was small. Once the boxes were placed on the asphalt, heat transfer quickly reduced the temperature of the surface of each box. The experiment would have benefitted from designing an internal heating element for the boxes so that they could not only achieve higher temperatures than the tap water would allow, but would remain at those temperatures throughout the experiment. Another change would be to incorporate a box with an intermediate emissivity. The majority of natural and manmade objects have emissivities that fall between 0.8 and 1.0, but two thirds of the boxes had emissivities less than 0.2. Although an intermediate emissivity was modeled in the synthetic experiment, it would have been nice to have had that in the real experiment as well.

Atmospheric data

The atmospheric data from the Buffalo radiosonde was a large contributor to the temperature extraction error, especially in the case of the low emissivity boxes. It would be advantageous to launch a radiosonde from the collection area at or near the time of the collect. At the very least, performing the experiment closer to the time the Buffalo radiosonde is launched might help mitigate atmospheric errors.

Along with a list of things to change in the future, it is important to note that there were a few things that made the experimental analysis easier. One was the use of reflective duct insulation as fiducial markers to mark the top of each row and column in the grid of boxes. In some of the images it was difficult to determine the location of a box. Using the fiducials and knowledge of the spacing between the boxes, the missing box was located. Performing

the experiment at night ended up being a good idea because the campus was quiet and the experiment was not disturbed while in progress due to pedestrian traffic.

5.2 Synthetic imagery study results

This main purpose of this study was to look at the possible reduction of temperature extraction errors with better knowledge of the atmosphere. This study used a best case scenario, where the atmosphere was modeled without the inclusion of error. It also looked at larger target-background temperature differences than those achieved in the airborne collect, as well as a more moderate target emissivity.

The set of boxes were modeled under an atmosphere with and without a cloud layer included. The study positioned the boxes within the center of the pixel in the middle of the study area, as well as across the boundary of four adjacent pixels. Every combination of those parameters led to four analysis sets whose results are graphically summarized in Figures 5.12, 5.13, 5.14, and 5.15. Figures 5.12 and 5.13 show the temperature extraction errors for the cloud-free atmosphere.




Figure 5.12: Synthetic experiment temperature extraction results for the clear atmosphere and boxes in the center of the center pixel.

The temperature extraction errors for the target representing the high emissivity painted box were very similar to those found in the airborne experiment. This makes sense because the temperature extraction errors from the high emissivity surface can tolerate errors in atmospheric knowledge. The temperature errors from the box with the 0.8 graybody emissivity were similar in magnitude to those from the painted box, however, they did not have a counterpart in the airborne collect so comparisons could not be made. Not surprisingly, the boxes with the most improvement over the airborne collect were those with the sandblasted and polished surfaces. The errors from the polished surfaces were again the largest overall, but in all cases are smaller than those from the airborne collect. This shows that with precise atmospheric knowledge, one could expect errors around 12°C or less for a box with an emissivity of approximately 0.02. Recall that while the atmosphere was perfectly modeled, errors for the target and background emissivity, background

temperature and target position were included. The errors would also decrease with better knowledge of these input parameters, until the only error left would come from the regression model itself.




Figure 5.13: Synthetic experiment temperature extraction results for the clear atmosphere and boxes offset from the center pixel.

The magnitude of the temperature errors were similar for the two positions. This indicates that as long as the target-containing pixels are correctly identified, the temperature extraction errors are not sensitive to positional errors. The temperature extraction errors also generally had the same magnitude regardless of the temperature difference between the target and the background. This indicates that this method may work just as well for large delta temperatures as it does for small temperature contrast situations.




Figure 5.14: Synthetic experiment temperature extraction results for the cloudy atmosphere and boxes in the center of the center pixel.

The results from the atmosphere with clouds yielded very similar results to the previous case. The magnitude of the temperature extraction errors were consistent with what was seen in the synthetic case with no clouds. The analysis did not include atmospheric error so the magnitude of the temperature extraction errors were not expected to fluctuate between the two atmospheric cases. The cloudy atmosphere simply contributed more downwelled radiance than the cloud-free atmosphere. However, the temperature errors for the cloudy cases might be optimistic because the downwelled term, which contained no error in the synthetic experiment, is an important term in cloudy weather. The larger downwelled contribution will reduce the temperature contrast between materials as well, therefore instrument errors will also become important. Instrument error was not modeled in the synthetic case nor was it treated in the airborne experiment. The temperature extraction errors were also consistent between the two positions (centered and offset).




Figure 5.15: Synthetic experiment temperature extraction results for the cloudy atmosphere and boxes offset from the center pixel.

In short, this exercise showed that improved atmospheric knowledge also improves the temperature error for low emissivity surfaces. It also showed that this method is not sensitive to large temperature difference between the target and the background and if the target-containing pixels are correctly identified, it will produce consistent results regardless of the target location relative to a pixel or a group of pixels.

5.3 Heat transfer utility study results

Section 4.10 described a study to determine whether or not modeling heat transfer decreased temperature extraction error. This section provides the analysis and conclusion of that study.

Figure 5.16 depicts the temperature extraction error as a function of the difference

between target and background temperature. The blue bars represent the error when heat transfer was not modeled in the process and the green bars represent the temperature error when heat transfer was taken into account via MuSES. All of the other input parameters were the same between the two data sets. For this case, the target was positioned in the center of the pixel in the middle of the study area.




Figure 5.16: Temperature extraction error as a function of the difference between target and background temperature for a target placed in the center of central image pixel.

Incorporating heat transfer into the modeling process decreased the temperature error across the board. The improvement was most significant where the difference between the target temperature and the background temperature was the largest. This makes sense because a larger temperature difference between the two surfaces causes a larger amount of heat transfer to occur. More heat transfer causes there to be a larger temperature gradient across the surface of the background material in

the vicinity of the target. Therefore, it is not valid to assume that the temperature of the background that was derived using pixels away from the target is the same as the target temperature in the vicinity of the target where heat transfer occurs. A MuSES visualization of a Krylon-painted box Figure 5.17 is provided to illustrate the difference temperature gradient across the background resulting from a box whose temperature is 53.6°C larger than the asphalt and a box whose temperature is 3.6°C larger than the asphalt temperature.




Figure 5.17: MuSES visualizations of Krylon-painted boxes with surface temperatures of 82.6°C (a) and 32.6°C (b) on asphalt with a temperature of 29°C . The images use the same color map, so the colors map to equivalent temperatures.

The temperature of the background in the area highlighted by the white outline in Figure 5.17(a) is 38.6°C , while the temperature of the same area in Figure 5.17(b) is 30.8°C . Modeling heat transfer becomes less important as the temperature difference between target and background decreases. Figure 5.16 shows that there is an improvement in the temperature extraction error of approximately 0.5°C when the difference between the background and target temperature is small.

This study was repeated for the case where the target was located within four

adjacent pixels. Figure 5.18 shows the temperature extraction error as a function of the difference between the target and background temperature.




Figure 5.18: Temperature extraction error as a function of the difference between target and background temperature for a target located within four adjacent pixels.

The results from this case show the same trend as when the target was placed in the center of the pixel in the middle of the study area. The magnitude of the temperature extraction errors are larger than they were in Figure 5.16, however, the temperature extraction errors for the centered and offset cases are similar to the extraction errors for the centered and offset Krylon Ultra Flat Black painted boxes from the synthetic experiment.

Chapter 6

Conclusion

This research investigated a method of using physics-based modeling as a means to estimate the temperature of a *subpixel* object from a *single-band* thermal infrared image. Boxes with different emissivities ranging from 0.02 - 0.91 were arranged on an asphalt background and single-band thermal imagery was collected to test the methodology. The GSD of the sensor was approximately 0.4m and the targets filled 38% of the area of the projected pixels. The average temperature extraction error for the Krylon Ultra Flat Black painted boxes was 3.3°C. The average temperature extraction errors for the sandblasted and polished boxes (excluding the anomalous data point from image 71) were 8.3°C and 15.35°C, respectively. The error analysis showed that errors in the knowledge of conductivity and background temperature had the largest effect on the magnitude of the temperature error. Finally, for the low emissivity cases, errors in the knowledge of the downwelled radiance had a large effect on the temperature extraction errors. The airborne collect also produced many lessons learned.

To show how the temperature extraction errors could be improved with perfect atmospheric knowledge, a follow-on synthetic experiment was done. The temperature extraction errors decreased across the board with the Krylon Ultra flat boxes having an average temperature error of 1.56°C, the sandblasted boxes having an average

temperature error of 5.1°C , and the polished boxes having an average temperature error of 9.93°C . The synthetic experiment also indicated that the methodology was insensitive to positional errors of the target as long as the target-containing pixels are correctly identified.

Finally, a study to determine the utility of modeling heat transfer in the process was performed. It was determined that modeling heat transfer reduces temperature extraction errors and the importance of modeling heat transfer increases as the temperature difference between the target and the background increases.

Table 6.1 provides a summary of this research and four previous works on the sub-pixel temperature extraction problem discussed in Chapter 4. The four methods were selected for comparison because they presented extraction errors and involved the use of solid targets. The extraction errors listed for this research are from the airborne experiment.

Table 6.1: Summary of subpixel temperature extraction methods.

Author	# bands used	Sensor GSD	Targets	Heat transfer	Extraction error
This research	1	0.4m	Single object ϵ_T : 0.2 - 1.0	Yes	3.3°C ($\epsilon_T = 0.91$) 8.3°C ($\epsilon_T = 0.2$)
Szymanski <i>et al.</i>	5 (ASTER) 15 (MODIS)	Not stated	Water/soil shore lines	No	13K (water), 27K (land) [ASTER] 1.5K (water), 3.9K (land) [MODIS]
Kustas <i>et al.</i>	1 thermal NDVI image	25m	Land	No	$\pm 1.5^\circ\text{C}$ uncertainty
Yang <i>et al</i>	1 temperature 3 VNIR	90m	Vegetation, soil, urban, water	No	1.8K (soil) 4K (urban)
Sentlinger <i>et al.</i>	2	100m	Water	No	0.96°C (expected)

None of the other papers tested low emissivity targets so there are no other points of comparison for the $\varepsilon_T = 0.2$ or 0.02 box from this research. The 3.3°C error from the Krylon Ultra Flat Black painted box is less than the temperature extraction errors of the non-water targets from the other methods. The emissivities from these targets are more in line with the emissivity of the painted box, so the comparison is fair. However, the error from the painted box in the synthetic experiment are similar to the temperature errors of the water found in the other papers and the average error from the sandblasted boxes is only a few degrees larger than the extraction errors from the other papers. It should also be noted that Szymanski's analysis was performed solely using synthetic data.

In conclusion, the methodology used in this research is a viable method for determining the absolute temperature of a high emissivity object that is unresolved by a single-band thermal imager. With improved atmospheric knowledge, this method could also be applied to low emissivity objects. The limiting factor in generating low temperature errors is the quality of knowledge of the input parameters to the modeling process. With improved parameter knowledge comes improved temperature extraction errors.

Bibliography

- [1] John R. Schott. *Remote Sensing: The Image Chain Approach*. Oxford University Press, 2 edition, 2007.
- [2] Pierre-Marie Robitaille. Kirchhoff's law of thermal emission: 150 years. *Progress in Physics*, 4, 2009.
- [3] J.M. Wallace and P.V. Hobbs. *Atmospheric Science An Introductory Survey*. Academic Press, New York, 1977.
- [4] Frank P Incropera, David P Dewitt, Theodore L Bergman, and Adrienne S Lavine. *Fundamentals of Heat and Mass Transfer*. John Wiley and Sons, 6 edition, 2007.
- [5] Tufts School of Engineering. Lecture notes.
- [6] Digital Imaging and Remote Sensing Laboratory. *The DIRSIG User's Manual*, 2006.
- [7] ThermoAnalytics. Muses 9.3 brochure. Calumet, MI.
- [8] Alexander Berk, Lawrence S. Bernstein, and David C Robertson. Modtran: A moderate resolution model for lowtran 7. Technical report, AFGL, Hanscomb AFB, 1989.
- [9] Sarah Paul, Adam A. Goodenough, Scott D. Brown, and Carl Salvaggio. Sub-pixel radiometry: a three-part study in generating synthetic imagery that incorporates sub-pixel variation. volume 7695, page 76950N. SPIE, 2010.

-
- [10] ThermoAnalytics. *MuSES IR Signature Prediction Software Training Manual*.
 - [11] Prabhat K. Acharya, Alexander Berk, Gail P. Anderson, George P. Anderson, North F. Larsen, Si Chee Tsay, and Knut H. Stamnes. Modtran4: multiple scattering and bidirectional reflectance distribution function (BRDF) upgrades to modtran. volume 3756, pages 354–362. SPIE, 1999.
 - [12] Douglas C. Montgomery, Elizabeth A. Peck, and G. Geoffrey Vinning. *Introduction to Linear Regression Analysis*. John Wiley, 2006.
 - [13] John Neter and William Wasserman. *Applied Linear Statistical Models: Regression, Analysis of Variables, and Experimental Designs*. Richard D. Irwin, 1974.
 - [14] Conversation with Dr. Peter Bajorski, June 2010.
 - [15] Kristin L. Sainani. The problem of multiple testing. *Physical Medicine and Rehabilitation*, 1(12), 2009.
 - [16] Mikel Aickin and Helen Gensler. Adjusting for multiple testing when reporting research results: The bonferroni vs holm methods. *American Journal of Public Health*, 86(5):726 – 728, 1996.
 - [17] Annual Meeting of the American Education Research Association. *A new sample size formula for regression*, 1994.
 - [18] John J. Szymanski, Christoph C. Borel, Quincy O. Harberger, Pawel Smolarkiewicz, and James P. Theiler. Subpixel temperature retrieval with multispectral sensors. In *Algorithms for Multispectral and Hyperspectral Imagery V*, volume 3868, pages 478–486, 1999.
 - [19] John W Salisbury and Dana M D’Aria. Emissivity of terrestrial materials in the 8–14 m atmospheric window. *Remote Sensing of Environment*, 42(2):83 – 106, 1992.

- [20] Jeff Dozier. A method for satellite identification of surface temperature fields of subpixel resolution. *Remote Sensing of Environment*, pages 221 – 229, 1981.
- [21] A. Lim, Soo Chin Liew, Kim Hwa Lim, and Leong Keong Kwoh. Retrieval of subpixel fire temperature and fire area in moderate resolution imaging spectrometer. In *Geoscience and Remote Sensing Symposium, 2002. IGARSS '02. 2002 IEEE International*, volume 6, pages 3205 – 3207, 2002.
- [22] Louis Giglio and Jacqueline D. Kendall. Application of the dozier retrieval to wildfire characterization: A sensitivity analysis. *Remote Sensing of Environment*, pages 34 – 49, 2001.
- [23] William P. Kustas, John M. Norman, Martha C. Anderson, and Andrew N. French. Estimating subpixel surface temperatures and energy fluxes from the vegetation index-radiometric temperature relationship. *Remote Sensing of Environment*, 2003.
- [24] Guijun Yang, Ruiliang Pu, Wenjiang Huang, Jihua Wang, and Chunjiang Zhao. A novel method to estimate subpixel temperature by fusing solar-reflective and thermal-infrared remote-sensing data with an artificial neural network. *Geoscience and Remote Sensing, IEEE Transactions on*, 48(4):2170 – 2178, 2010.
- [25] Gabriel I. Sentlinger, Simon J. Hook, and Bernard Laval. Sub-pixel water temperature estimation from thermal-infrared imagery using vectorized lake features. *Remote Sensing of Environment*, 112:1678 – 1688, 2008.
- [26] William T. Gustafson, Rebecca N. Handcock, and Alan R. Gillespie. An image-sharpening method to recover stream temperatures from ASTER images. *Remote Sensing for Environmental Monitoring, GIS Applications, and Geology II*, 4886:72–83, 2003.
- [27] A. Gillespie, S. Rokugawa, T. Matsunaga, J.S. Cothorn, S. Hook, and A.B. Kahle. A temperature and emissivity separation algorithm for advanced space-

- borne thermal emission and reflection radiometer (ASTER) images. *Geoscience and Remote Sensing, IEEE Transactions on*, 36(4):1113 – 1126, jul 1998.
- [28] C. Ottle, A. Kallel, G. Monteil, S. LeHegarat, and B. Coudert. Subpixel temperature estimation from low resolution thermal infrared remote sensing. In *Geoscience and Remote Sensing Symposium, 2008. IGARSS 2008. IEEE International*, 2008.
- [29] A description of RAMS. <http://rams.atmos.colostate.edu/rams-description.html>.
- [30] Atmospheric Technologies Group. Advanced atmospheric modeling. Technical report, Savannah River National Laboratory.
- [31] Vivek Verma, Rakesh Kumar, and Stephen Hsu. 3d building detection and modeling from aerial lidar data. In *IEEE, CVPR*, 2006.
- [32]
- [33] *CRC Handbook of Tables for Applied Engineering Science (Second Edition)*. CRC Press, 2nd edition, 1976.

Appendix A

Heat Transfer Study

This study was conducted to determine the importance of modeling heat transfer in the subpixel temperature extraction methodology. An indoor transformer on a concrete floor was used in this experiment. The transformer had a temperature of 36.3 °C at its base and a temperature of 47.6 °C at the top. Contact temperature measurements of the floor were taken every 0.15 meters starting at the base of the transformer and moving out to a distance of just over 2 meters (7 feet). The temperature measurements were plotted versus distance and the result is found in Figure A.1.




Figure A.1: Plot of measurements of ground temperature taken as a function of distance from an electrical transformer.

The plot in Figure A.1 shows an increase in ground temperature as the distance from the transformer decreases. There was a 3 degree difference from the base of the transformer to just over 2 meters away from the transformer. While ignoring this temperature gradient might not be the largest source of error in the methodology, it could introduce appreciable error.

Appendix B

Novel contributions to the field

Previous methods applied to the subpixel temperature problem have used multiple bands and perhaps the most notable contribution of this research was that it only used one thermal band to solve the problem. Not only that, but the single-band solution produced results that are comparable with the multi-band techniques. This research also employed a physics-based modeling approach that used MuSES, MODTRAN, and DIRSIG to model the heat transfer, atmosphere, and sensor-reaching radiance, respectively. Previous methods have been based on spatial downscaling or spectral mixture analysis. Finally, this research investigated subpixel temperature extraction with extremely low emissivity objects, which is something that has not been published to date.

Appendix C

Suggestions for future work

This research focused on targets with uniform surface temperatures. However, MuSES is capable of computing a temperature gradient across the surface of an object, so the next step should be to leverage that capability and extend this methodology to targets without uniform surface temperatures. It is also suggested that a new field experiment be conducted that incorporates the lessons learned from this project. More time to plan and more money for better supplies are highly encouraged. It would also be useful to incorporate a time-based, per facet sun shadow option to MuSES for cloud coverage computations.