



**EVALUATION OF A CLOUD DETECTION
TECHNIQUE USING SPATIAL AND
RADIOMETRIC THRESHOLDS
FOR NEAR INFRARED SATELLITE
IMAGERY**

THESIS

William J. Graff, Capt, USAF
AFIT-ENP-MS-18-M-083

**DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY**

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

DISTRIBUTION STATEMENT A
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views expressed in this document are those of the author and do not reflect the official policy or position of the United States Air Force, the United States Department of Defense or the United States Government. This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

AFIT-ENP-MS-18-M-083

EVALUATION OF A CLOUD DETECTION TECHNIQUE USING
SPATIAL AND RADIOMETRIC THRESHOLDS FOR NEAR INFRARED
SATELLITE IMAGERY

THESIS

Presented to the Faculty
Department of Engineering Physics
Graduate School of Engineering and Management
Air Force Institute of Technology
Air University
Air Education and Training Command
in Partial Fulfillment of the Requirements for the
Degree of Master of Science in Atmospheric Science

William J. Graff, B.S.

Capt, USAF

March 22, 2018

DISTRIBUTION STATEMENT A
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

EVALUATION OF A CLOUD DETECTION TECHNIQUE
USING SPATIAL AND RADIOMETRIC THRESHOLDS
FOR NEAR INFRARED SATELLITE IMAGERY

THESIS

William J. Graff, B.S.

Captain, USAF

Committee Membership:

Robert A. Stenger, Ph.D.

Chair

Maj H. Rose Tseng, Ph.D.

Member

Robert D. Loper, Ph.D.

Member

Abstract

Knowledge of cloud location in near infrared (NIR) imagery is of interest to the meteorological community given the wavelengths greater spatial resolution compared to longwave infrared and its potential nighttime applications. This method consists of an algorithm that can be employed by multiple instrument platforms. It analyzes changes between satellite image radiances and a seasonal synthetic background radiance image. The NIR sensor bands of the Sentinel-2 Multispectral Instrument (MSI) and Suomi National Polar-Orbiting Partnership (Suomi NPP) Visible Infrared Imaging Radiometer Suite (VIIRS) satellite instruments are used to evaluate the performance of a monochromatic change detection technique designed to locate areas of cloud cover. VIIRS imagery is examined for its higher temporal resolution compared to MSI; whereas, MSI imagery is examined for its greater spatial resolution. Background images are constructed either manually or algorithmically using a first-guess image. Observed and background images are differenced based on user-defined radiance, size, and shape thresholds. Pixels that meet these thresholds in the first algorithm are flagged as cloud cover. Output is compared to operational cloud masks that rely on multispectral techniques. Findings indicate that the developed algorithm identifies cloud cover above the specified size threshold well, but optically thin clouds and fresh snow still present limitations.

Acknowledgements

Dedicated to my parents for their unwavering support and love.

Acknowledgements

The scope and complexity of this project required me to seek advice from a handful of individuals to whom I owe a great deal of thanks. Firstly, I would like to thank my advisor Dr. Robert Stenger for his guidance and patience. His knowledge of remote sensing and imagery was invaluable. Additionally, I wish to thank the members of my committee, Dr. Robert Loper and Maj Rose Tseng. Maj Tseng provided helpful suggestions for study sites, exploration topics, and figures. Dr. Loper furthered my understanding of orbital mechanics and threats faced by sensors in the space environment. Additionally, I extend gratitude to my friend and cousin for their writing suggestions. This document also contains modified Copernicus Sentinel data from 2016 and 2017. I extend my gratitude to the European Space Agency for allowing me to access and use it. Finally, I owe the utmost appreciation to my parents for their encouragement through the long hours it took to complete this enormous project. Thank you.

William J. Graff

TABLE OF CONTENTS

	Page
Abstract	iv
Acknowledgements	v
List of Figures	xi
List of Tables	xxvi
Acronym and Abbreviation List	xxviii
I. Introduction	1
A. Motivation	1
1. Definition of Near Infrared Satellite Imagery	3
2. Definition of Satellite-Based Change Detection	6
3. Definition of Cloud Cover	7
B. Research Objective and Technique Preview	10
1. Research Objective	10
2. Technique Preview	11
II. Related Research and Background	17
A. Cloud Detection Techniques	18
1. Threshold Techniques	18
2. Histogram Techniques	19
3. Pattern Recognition Techniques	22
4. Multispectral Techniques	22
5. Spatial Coherence Techniques	24
6. Geometric Techniques	25
7. Hybrid Techniques	25
B. Standard Cloud Masks	27
1. Multispectral Instrument (MSI) Cloud Masks	27
2. Visible Infrared Imaging Radiometer Suite (VIIRS) Cloud Masks	29
C. Related Research	35

	Page
III. Methodology	38
A. Imagery Attributes	38
1. VIIRS Datasets	38
1.1 Daytime Imagery	39
1.2 Daytime Cloud Masks	39
1.3 Nighttime Imagery	40
1.4 Nighttime Cloud Masks	41
2. MSI Datasets	42
2.1 Imagery	42
2.2 Cloud Masks	43
3. Digital Elevation Models	43
B. Hardware and Specialized Software	45
1. Hardware and Operating Systems	45
2. Specialized Software	45
2.1 Polar2Grid	46
2.2 ERDAS Imagine	46
2.3 ArcMap	47
2.4 ArcScene	47
C. Case Study Locations	48
1. VIIRS Sites	49
1.1 Fairbanks, Alaska	49
1.2 Miami, Florida	50
1.3 Pyongyang, North Korea	51
1.4 South Bend, Indiana	53
2. MSI-Only Site: Ambouli, Djibouti	54
3. VIIRS and MSI Sites	55
3.1 Bagram, Afghanistan	55
3.2 Los Angeles, California	56
D. Background Image Synthesis	57
1. Data Cleanup	57
2. First Guess Image	59
3. Background Image Synthesis Generation	60

	Page
3.1 Manual Background Synthesis Method	60
3.2 Primary Automated Synthesis Method	71
3.3 Alternate Automated Synthesis Method	77
3.4 Stereographic Terrain Analyses	82
E. Cloud Identification Workflow	85
1. Data Cleanup	85
2. VIIRS Settings	95
2.1 Daytime Settings	95
2.2 Nighttime Settings	96
3. MSI Settings	98
F. Qualitative and Quantitative Evaluations of Imagery	99
1. Truth Data	100
1.1 Operational Truth Data	100
1.2 Synthetic Truth Data	100
2. Chi-square Goodness-of-Fit Test	101
3. Change Detection Test	105
4. Contingency Tables	106
IV. Results and Analysis	109
A. Analysis of Synthetic Background Imagery	109
1. VIIRS Results	112
1.1 Afghanistan Results	112
1.2 Alaska Results	119
1.3 California May Results	122
1.4 California June Results	126
1.5 California July Results	129
1.6 Florida Results	132
1.7 Indiana February Results	137
1.8 Indiana July Results	141
1.9 North Korea May Results	144
2. MSI Results	148
2.1 Afghanistan Results	148
2.2 California Results	152

	Page
2.3 Djibouti Results	156
B. Analysis of VIIRS Results	160
1. Daytime Results	162
1.1 Afghanistan Results	162
1.2 Alaska Results	163
1.3 California Results	164
1.4 Indiana February Results	165
1.5 North Korea Results	166
2. Nighttime Results: Florida	167
C. Analysis of MSI Results	179
1. Bagram, Afghanistan Results	180
2. Ambouli, Djibouti Results	180
3. Los Angeles, California Results	181
D. Cross-Platform Background Test	186
V. Conclusion	193
A. Summary of Results	193
B. Recommendations for Future Research	197
Bibliography	200
Vita	205

List of Figures

Figure		Page
1.	Infrared subdivisions of the electromagnetic spectrum (EM) spectrum. Adapted from Tattersall (2016).	6
2.	Near infrared (NIR) view of a Los Angeles, CA marina on (a) June 18, 2016 and (b) the same location 10 days later.	12
3.	Image of change produced by subtracting (a) from (b). Shades of red indicate where pixel values have decreased. Shades of blue show where pixel values have increased. Dark shades indicate pixels where the most change has occurred, such as those associated with ship locations.	13
4.	Channel information for the MSI. The vertical axis displays spatial resolution in meters. The horizontal axis displays channel wavelength in nanometers. Rectangle width approximates the spectral width of each channel (ESA 2015a).	15
5.	An idealized two-dimensional histogram with infrared brightness temperature along the vertical axis and visible radiance along the horizontal axis. Pixel clusters are labeled according to their corresponding features. Adapted from Kidder and Vonder Haar (1995).	21
6.	Multispectral image of Los Angeles, California captured by the MSI on November 30, 2017. This false-color image assigns the visible green, visible red, and NIR channels to the monitor's blue, green, and red color guns, respectively. This channel arrangement reduces haze due to the absorption of incoming radiation by optically active gases in the atmosphere while highlighting features of interest. Clouds are shown as white, vegetation appears red, and bodies of water appear dark blue or black. Anthropogenic structures appear in light shades of blue. Identifying cloud cover with this enhancement must be done manually (Campbell and Wynne 2011).	23

Figure		Page
7.	Diagram of the MSI cloud detection workflow. The process utilizes information from multiple channels as well as external data (ESA 2017b).	28
8.	Operational cloud detection workflow for VIIRS daytime cloud masks. Nighttime cloud masks are produced from a simplified version of the same algorithm. Adapted from Hutchison et al. (2005).	32
9.	Example cloud mask overlaid atop an observed NIR image of California from May 26, 2017. The cloud mask is translucent and yellow.	33
10.	False-color image of an urban environment and surroundings from May 4, 2002 in (a). Vegetation appears in red and human infrastructure appears in shades of blue. Water is black. Synthetic image generated for the same date in (b). The numerical difference in pixel values between the first two panels in (c). Histogram in (d) is the differences between pixels from (a) and (b); (Zhu et al. 2015).	37
11.	NIR image of California from May 26, 2017 derived from HDF5 component images. California's border is shown in green (NOAA 2017c).	40
12.	NIR nighttime image of Florida and its surroundings on September 7, 2017. The US coastline is shown in yellow (NOAA 2017c).	41
13.	NIR image of the region surrounding Bagram, Afghanistan from June 16, 2016 (ESA 2017a).	43
14.	Synthetic background image of Afghanistan with three-dimensional terrain features generated from a digital elevation model (DEM) (Patterson 2017). Bagram is marked with a green flag. The Afghan border is highlighted in red.	44
15.	Map depicting all sites examined during this study. Stars mark each site. Their surrounding states or nations are highlighted in orange (Glickman 2000).	48

Figure	Page
16. NIR VIIRS image of Fairbanks, AK on a relatively clear day. The town is marked by a green dot. An inset map of Alaska is shown in the lower right corner of the image.	49
17. NIR VIIRS image of Florida on September 2, 2015. Major cities are marked with yellow dots. State borders are outlined with green.	51
18. NIR VIIRS image of North Korea May 18, 2016. Major cities are marked with yellow dots. North Korea's border is outlined in red. Other national borders are green.	52
19. NIR VIIRS image of Indiana on July 7, 2013. Major cities are marked with yellow dots. State borders are outline with green.	53
20. Synthetic NIR MSI background image of Ambouli, Djibouti. The city is marked with a green dot. An inset image of the country's borders is shown in the lower right of the figure. The country's border is outlined in yellow. Other national borders are green.	54
21. NIR VIIRS image of Afghanistan on May 5, 2015. Bagram is marked with a green dot. Afghanistan is outlined in red. Other national borders are yellow.	55
22. NIR MSI image of Los Angeles, CA on June 18, 2016. The city center is marked with a green dot. The coast is outlined in yellow.	56
23. Diagram of the model used to repair images with missing pixel data. Circles represent input and output files. The color blue indicates a file or directory variable, while light or dark green indicates an output value or file. Arrows with solid lines indicate the direction of the model. Arrows with dotted lines indicate the file is a precondition for an action that follows. Yellow rectangles indicate the execution of a process. Hexagons indicate do-loops. Any shape with a bold P over it is a parameter that can be defined in the graphical user interface for the model.	59
24. Original VIIRS image of North Korea on 25 May 2017 at 0407 UTC.	59

Figure	Page
25.	Cloud-free image snippet of North Korea on 8 May 2013 at 0417 UTC. The national border is outlined in red. 62
26.	Cloud-free image snippet of North Korea on 8 May 2013 at 0407 UTC merged with the initial synthetic background image. The grayscale color ramp causes each member of the composite to stand out more than in the final synthetic image. 63
27.	Cloud-free synthetic image of North Korea without water corrections. 64
28.	Simulated water raster for the vicinity of North Korea for May. The national border is outlined in red for spatial awareness. The land is colored white and water colored is black. 65
29.	Coastline (red) and 5 <i>km</i> coastline buffer (yellow) shapefiles overlaid atop the synthetic water image. 67
30.	Final cloud-free Synthetic image of North Korea for the month of May. The national border is outlined in red. Notice the increased contrast in land formations compared to Figure 27. 68
31.	An observed image of North Korea's northeast border. 70
32.	The yellow circle highlights an area of non-physical bumpy terrain. Notice how peaks and valleys stand out with high contrast inside the country's border (left portion of image), but outside the country's border the terrain quickly becomes rough in a manner not consistent with the terrain in Figure 31. 70
33.	The synthetic image generator graphical user interface (GUI). 72
34.	Diagram showing the workflow for the synthetic image generator. See Figure 23's caption for an explanation of the symbology. 73
35.	Input and output directories and files for the synthetic image generator. See Figure 23's caption for an explanation of the symbology. 73

Figure	Page
36.	Input variables for the synthetic image generator. See Figure 23's caption for an explanation of the symbology. 74
37.	Initial steps of the synthetic image generator's workflow. See Figure 23's caption for an explanation of the symbology..... 75
38.	Portion of synthetic image generator's workflow during which values for missing pixels are determined and cloud free portions of input file are merged with first-guess image. See Figure 23's caption for an explanation of the symbology. 76
39.	The final steps of the synthetic image generator. The algorithm looks for large negative changes in pixel values and replaces the identified pixels with first-guess pixels. See Figure 23's caption for an explanation of the symbology..... 76
40.	Subset of the synthetic background image resulting from the primary generation method. The mountains and landscape are blurry. Bagram, Afghanistan can be seen in the upper left of the image..... 78
41.	An observed image of the region surrounding Bagram on May 1, 2012. 79
42.	Subset of the final synthetic image using the alternate automated background generation method. Compare to Figures 41 and 40. 79
43.	Subset of the final synthetic image using the alternate automated background generation method with lakes highlighted in green (ESRI 2007a). 80
44.	The original synthetic background image for Afghanistan using the primary automated method. 81
45.	The image resulting from the pixel smoothing tool. Increased snow cover extent can be seen over the mountains on the right side of the image. 81
46.	Final synthetic background image for Afghanistan produced with the the alternate automated method. 82

Figure	Page
47. Three-dimensional view of Bagram generated by projecting the synthetic background onto the DEM's topography.	83
48. Original synthetic background image of Bagram. Large dark spots atop the mountains in the foreground interfere with nearly continuous snow cover.	84
49. Final synthetic background image for Bagram. Compare the evenly distributed Dendritic patterns over the mountains in the foreground to those seen in Figure 48.	84
50. The automated cloud analysis model (ACAM)'s primary GUI.	86
51. GUI for batch-processing multiple images using the ACAM.	86
52. Overview of the ACAM's image processing workflow as in Figure 23.	87
53. Symbolic overview of input and output variables for the ACAM as in Figure 23.	88
54. Symbolic overview of the initial raster calculator that conducts change detection math using both input images as in Figure 23.	90
55. Symbolic overview of tools that convert the binary mask into polygons as in Figure 23.	90
56. Selected polygons are used to extract cloud free portions of the original image that is output as a mask. The algorithm can be changed to output either a binary mask or mask retaining the original pixel values of all clouds from the input scene. Symbology as in Figure 23.	91
57. Observed VIIRS image over California from May 26, 2017 at 2042 UTC.	92
58. May synthetic background image for California.	93
59. Cloud cover detected over California on May 26, 2017 at 2042 UTC. The cloud mask is translucent and yellow.	94

Figure	Page
60.	Subset of the binary mask indicating surface lights generated for Florida and the surrounding region. 97
61.	Curve used to assess results of the chi-square test. The critical value is displayed on the horizontal axis. 103
62.	The workflow used to execute the chi-square goodness-of-fit test with symbology as in Figure 23. 103
63.	Input files and initial steps of the chi-square test workflow with symbology as in Figure 23. 104
64.	Input variables and random point generating tool used to extract data from sample points for the chi-square test with symbology as in Figure 23. 105
65.	Pixel values are extracted from the raster calculator's output at random points across the domain. The extracted values are saved to a shapefile along with the associated random points with symbology as in Figure 23. 106
66.	Comparison of operational and synthetic truth masks for VIIRS daytime imagery in (a). The observed image is given for reference in (b). The mask is translucent and overlaid atop an observed image from July 7, 2017. Locations that are identified mutually as cloud-free or cloudy are clear or white, respectively. Clouds identified by the operational truth mask only are colored red. Clouds identified by the synthetic truth mask are colored green. When compared to the observed image in (b), manual analysis revealed that the operational mask picks up more cirriform and optically thin cloud cover, but the synthetic truth mask picks up on more small cumulus clusters. 113
67.	Comparison of operational and synthetic truth masks for VIIRS nighttime imagery. The mask is translucent and overlaid atop an observed image from September 11, 2017. Colors as seen in Figure 66. When compared to the observed image in (b), manual analysis revealed the operational mask outperforms the synthetic truth mask with cirriform and optically thin clouds better, but underperforms with small areas of cloud cover. 114

Figure	Page
68.	Final synthetic background image for Afghanistan during May.
69.	Subset images comparing the clarity of dendritic patterns and rivers over eastern Afghanistan. The left subset was taken from an observed image. The right subset was taken from the final synthetic background. Yellow circles highlight differences between the clarity of rivers between the images. Red circles highlight changes in dendritic snow patterns between images. Notice that while contrast increases with the synthetic image, land-water boundaries are less distinct compared to the observed image. Snow covers more of the mountaintops in the synthetic image, but land snow boundaries are also blurred in comparison to the observed image.
70.	Plot of the random sample points measuring the magnitude of difference between the image observed on May 15, 2015 and the synthetic background image for Afghanistan. The horizontal red line represents the sample mean of $19.3 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. The majority of sample points fall below the detection threshold of $45 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$
71.	Image depicting the magnitude of pixel difference between the image observed on May 15, 2015, and the synthetic background image for Afghanistan. Differences range from 2.1 to $179.4 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. Tan shows the lowest values of change. Shades of red show areas of greater change, with the darkest reds associated with the highest values. Differences above the highest value observed from the sample points for the chi-square test can be seen in violet. Notice the darkest reds correspond frequently to areas where dendritic patterns can be found along mountains. The cluster of violet pixels in the upper left corner of the image is associated with cloud cover. Sample points are marked with dark green rhombi.
72.	Final synthetic background image for Fairbanks, Alaska and the surrounding region for April.

Figure	Page
73.	Plot of the random sample points measuring the magnitude of difference between the image observed on April 21, 2017 and the synthetic background image for Fairbanks, AK. The horizontal red line represents the sample mean of $7.7 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. All sample points are below the detection threshold of $45 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$ 121
74.	The final background image for California during May. Snow cover can still be seen over the mountains near the center of the image. 123
75.	Difference between California's May synthetic background and the image observed on May 5, 2017. Sample points are marked by yellow rhombi. 124
76.	Plot of the random sample points measuring the magnitude of difference between the image observed on May 20, 2017 and the synthetic background image for California. The horizontal red line represents the sample mean of $9.9 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. The majority of sample points are below the detection threshold of $45 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$ 125
77.	The final synthetic background image for California during June. 127
78.	Plot of the random sample points measuring the magnitude of difference between the image observed on June 24, 2016 and the synthetic background image for California. The horizontal red line represents the sample mean of $6.4 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. All sample points are below the detection threshold of $45 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$ 128
79.	The final synthetic background image for California during July. Snow cover along the mountaintops at the center of the image has decreased from Figure 74, but still remains visible in the background image. 130
80.	Plot of the random sample points measuring the magnitude of difference between the image observed on July 28, 2017 and the synthetic background image for California. The horizontal red line represents the sample mean of $9.7 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. All sample points are below the detection threshold of $45 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$ 131

Figure	Page
81.	Final synthetic background image for Florida at night. Parent images consist of moonlit nights observed from July to October from 2012 to 2017. 132
82.	Subset of Florida’s synthetic background image. Lake Okeechobee can be seen near the center of the image. Differences between land cover types not in the immediate vicinity of lights can also be observed in the image. 133
83.	Difference between Florida’s synthetic background and the image observed at 0635 UTC on September 13, 2013. Sample points are marked by green rhombi. The largest differences are shown in shades of light gray and white and frequently near cities and other anthropogenic lights sources. 135
84.	Plot of the random sample points measuring the magnitude of difference between the image observed on September 13, 2013 at 0635 UTC and the synthetic background image of Florida. The horizontal red line represents the sample mean of $1.0 \times 10^{-9} \text{ W m}^{-2} \text{ sr}^{-1} \mu\text{m}^{-1}$. The sample with the largest difference is associated with changes in radiance near city lights. Most sample points are below the detection threshold of $5 \times 10^{-9} \text{ W m}^{-2} \text{ sr}^{-1} \mu\text{m}^{-1}$ 136
85.	The final synthetic background image for Indiana during the month of February. Snow can be seen along the Indiana-Michigan border in the uppermost parts of the image. 138
86.	Difference between Indiana’s February synthetic background and the image observed on February 17, 2017 at 1800 UTC. Sample points are marked by yellow rhombi. The white pixels near the top of the state are associated with snow. 139

Figure	Page	
87.	Plot of the random sample points measuring the magnitude of difference between the image observed on February 17, 2017 at 1800 UTC and the synthetic background image for Indiana. The horizontal red line represents the sample mean of $3.3 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. All sample points are below the detection threshold of $45 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$	140
88.	The final background image of Indiana for July.	142
89.	Plot of the random sample points measuring the magnitude of difference between the image observed on July 2, 2017 and the synthetic background image for Indiana. The horizontal red line represents the sample mean of $11.6 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. All sample points are below the detection threshold of $45 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$	143
90.	The final synthetic background image for North Korea during May.	145
91.	Difference between North Korea's May synthetic background and the image observed on May 19, 2017. Sample points are marked by yellow rhombi.	146
92.	Plot of the random sample points measuring the magnitude of difference between the image observed on May 19, 2017 and the synthetic background image for North Korea. The horizontal red line represents the sample mean of $16.0 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. The spread of values about the mean is wide with respect to the maximum and minimum values, but is not large enough for the images to be considered significantly different. All sample points are below the detection threshold of $45 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$	147
93.	The final synthetic background for Bagram, Afghanistan and the surrounding region for May and June.	149
94.	Difference between Bagram, Afghanistan's synthetic background and the image from June 1, 2016. Sample points are marked by yellow rhombi. The relatively bright area of land in the lower right corner of the image indicates the difference in spatial coverage spatial between the observed and synthetic images. Sample points are not taken from this area.	150

95.	Plot of the random sample points measuring the magnitude of difference between the image observed on June 1, 2016 and the synthetic background image for Bagram, Afghanistan. The horizontal red line represents the sample mean of $222.3 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. The spread of values about the mean is wide with respect to the maximum and minimum values, but is not large enough for the images to be considered significantly different. The majority of points are below the detection threshold of $1000 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$	151
96.	Final synthetic background image for Los Angeles, California during June and July.	153
97.	Difference between Los Angeles, California's synthetic background and the image observed on June 13, 2017. Sample points are marked by green rhombi. Large differences between the sea's synthetic and observed values are associated with bright pixels. Areas of urban landscape in the upper left corner of the image also appear in bright shades.	154
98.	Plot of the random sample points measuring the magnitude of the difference between the image observed on June 13, 2017 and the synthetic background image for Los Angeles, CA. The horizontal red line represents the sample mean of $234.2 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. The spread of values about the mean is wide with respect to the maximum and minimum values, but is not large enough for the images to be considered significantly different. The majority of points are below the cloud detection threshold of $1000 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$	155
99.	The final synthetic background image for Ambouli, Djibouti and its surroundings during June and July.	157

FigurePage

100.	Plot of the random sample points measuring the magnitude of the difference between the image observed on July 5, 2017 and the synthetic background image for Ambouli, Djibouti. The horizontal red line represents the sample mean of $113.3\text{ }Wm^{-2}sr^{-1}\mu m^{-1}$. The spread of values about the mean is wide with respect to the maximum and minimum values, but is not large enough for the images to be considered significantly different. All sample points are below the detection threshold of $1000\text{ }Wm^{-2}sr^{-1}\mu m^{-1}$	158
101.	Difference between Anbouli, Djibouti's synthetic background and the image observed on July 5, 2017. Sample points are marked by yellow rhombi. The white cluster of pixels in the upper left of the image is associated with cloud cover. Sample points are not taken from this area.	159
102.	Mask comparison image overlaid atop the observed image from May 9, 2017 in (a). The observed image is given for reference in (b). Mask comparison is translucent. Clouds identified by both masks are colored gray or white. Clouds identified only by the ACAM's mask are green. These are referred to as false positive pixels. Clouds identified only by the truth data mask are red. These are referred to as false negative pixels.	168
103.	Comparison of cloud masks on April 26, 2017 over Fairbanks, Alaska in (a) with colors as in Figure 102. The observed image is given for reference in (b).	169
104.	Mask comparison over California on May 30, 2017 in (a) with colors as in Figure 102. The observed image is given for reference in (b).	170
105.	Subset of mask comparison over the Sierra Nevada Mountains in eastern California on May 30, 2017 in (a) with colors as in Figure 102. The observed image is given for reference in (b).	171
106.	Subset of mask comparison over Los Angeles, California on July 24, 2017 in (a) with colors as in Figure 102. The observed image is given for reference in (b).	172

Figure	Page
107. Mask comparison for Indiana on February 23, 2017 in (a) with colors as in Figure 102. The observed image is given for reference in (b).	173
108. Mask comparison for July 23, 2017 over Indiana in (a) with colors as in Figure 102. The observed image is given for reference in (b).	174
109. Mask comparison for July 27, 2017 over Indiana in (a) with colors as in Figure 102. The observed image is given for reference in (b).	175
110. Mask comparison for May 5, 2017 over North Korea in (a) with colors as in Figure 102. The observed image is given for reference in (b).	176
111. Mask comparison for September 10, 2017 over Florida in (a) with colors as in Figure 102. The observed image is given for reference in (b).	177
112. Mask comparison for September 11, 2017 over Florida in (a) with colors as in Figure 102. The observed image is given for reference in (b).	178
113. Mask comparison for Bagram on May 12, 2017. Red-colored pixels at the bottom of the image are not all associated with cloud cover in (a) with colors as in Figure 102. The observed image is given for reference in (b).	182
114. Subset of the May 12, 2017 mask comparison near Bagram, Afghanistan in (a) with colors as in Figure 102. While both masks flag parts of the mountain snow cover as clouds, the ACAM's mask outperforms the operational mask. The operational mask struggles most with the optically thin clouds surrounding the cumulus clouds near the center of the image. The observed image is given for reference in (b).	183
115. Mask comparison for Ambouli, Djibouti on July 20, 2017 in (a) with colors as in Figure 102. The observed image is given for reference in (b).	184

Figure	Page
116.	Mask comparison for Los Angeles, California on June 8, 2016 in (a) with colors as in Figure 102. The observed image is given for reference in (b). Issues with the operational mask flagging non-cloudy pixels as cloudy can be seen in the upper half of the image. 185
117.	The mean image of Los Angeles determined by averaging pixel values from California's June and July VIIRS background images. 186
118.	The final cross-platform synthetic background image for Los Angeles, California is displayed in (a). It is comprised of VIIRS backgrounds from June and July. The MSI background for this time period is displayed in (b). 188
119.	Magnitude of difference between the original and cross-platform background images for Los Angeles, California. Sample points are marked with green rhombi. The largest differences are associated with shades of white and mostly occur over the cityscape. 190
120.	Mask comparison between ACAM runs using the cross-platform background and original MSI synthetic background from June 6, 2016 for Los Angeles, CA is displayed in (a). The mask comparison is translucent and overlaid atop the background image. Cloud-free pixels mutually identified by both masks are not assigned colors. Clouds identified only by the original background's mask are colored red. Clouds identified only by the cross-platform mask are colored green. The original image is displayed in (b). 192

List of Tables

Table		Page
1.	Channel information for VIIRS. There are two primary types of channels: moderate resolution bands (M) and imagery resolution bands (I). There is also one day-night band (DNB) that is optimized to sense imagery both day and night. Channels used to construct the operational cloud mask are marked with an x. Horizontal spatial resolution (Horizontal spatial resolution (HSR)) at nadir is the resolution sensed at the point directly below the sensor's field of view (field of view (FOV)); (NOAA 2017b).	14
2.	Pixel values for VIIRS cloud masks. Cloudy and cloud-free flags are grouped into separate classes to produce the binary cloud mask. Cloud-free pixels in the binary mask are assigned values of zero. Cloudy pixels are assigned values of one. Adapted from NASA (2014).	31
3.	List of resources used by the daytime VIIRS cloud classification algorithm. Applicable land surfaces for each scheme are marked with an 'x' (NASA 2014).	32
4.	List of resources used by the nighttime VIIRS cloud classification algorithm. Applicable land surfaces for each scheme are marked with an 'x' (NASA 2014).	34
5.	ACAM settings by instrument and time of day.	95
6.	Chi-Square results for all synthetic background images by site, month, instrument (INSTR), sample point number (PNTS), and degrees of freedom (DF). The chi-square test statistic (CHI-SQ) is given, along with the associated critical value (CRT VAL). The table also displays whether background images differed significantly from the observed image used in each test (SG DIF).	110
7.	The number of parent images used to generate the synthetic background for each site by location, month, and instrument (ISNTR).	110
8.	Summary of results for comparisons between ACAM and truth data masks.	160

Table		Page
9.	Results for comparisons between ACAM and VIIRS operational masks.	161
10.	Results for comparisons between ACAM and VIIRS synthetic masks.	161
11.	Results from comparisons between the ACAM and MSI masks.	179

EVALUATION OF A CLOUD DETECTION TECHNIQUE
USING SPATIAL AND RADIOMETRIC THRESHOLDS
FOR NEAR INFRARED SATELLITE IMAGERY

I. Introduction

A. Motivation

Reliable methods of automated cloud detection are desired by both the aviation and defense communities. Positive identification of clouds allows threats to aircraft safety or mission success to be rapidly mitigated. According to the American Meteorological Society (AMS) Glossary of Meteorology, clouds are any visible accumulation of frozen or liquid water droplets suspended above a surface (Glickman 2000). Clouds come in many sizes, shapes, textures, and phase combinations. Once identified, weather analysts can determine present and future hazards that may prove dangerous to human activities both at the surface and aloft.

NIR satellite imagery can reveal cloud locations at night at spatial resolutions near that of visible imagery. There is evidence, given sufficient illumination, that the technique presented in this paper can be used to identify daytime and nighttime cloud cover. The ability to delineate the boundary between cloud and clear sky is of particular interest to the Air Force. Activities that require cloud-free line of sight (CFLOS) depend on accurate knowledge of where visual limitations can occur. Aerial refueling, reconnaissance, and search and rescue missions all require high visibility to ensure success around the clock. Aircraft sensitive to icing conditions must also be able to avoid clouds, as well as those used for close air support in hostile environments. The need for accurate cloud boundary data only increases as mission length increases, driving the need for an automated system that can rapidly detect clouds.

Multiple techniques have been developed for detecting clouds in satellite imagery. An ACAM that incorporates visible imagery allows for high spatial resolution, but is not usable after sunset. ACAMs that work during both day and night run at coarse resolutions and generally require output from weather models to determine cloud boundaries. Those that do not rely on outside data are prone to identifying bright, non-meteorological features as clouds. Systems that rely on differences between multiple satellite channels to detect cloud cover within a scene have been developed and work well, but are complex and add to the computational resources necessary to detect cloud cover. They also demand multiple channels from the same satellite, restricting the number of platforms to which they can be applied or causing them to be platform-specific (Kidder and Vonder Haar 1995). These are called multispectral techniques and are available at multiple resolutions, but they are limited to the resolution of the most coarsely resolved channel. These techniques will be discussed in further detail in Chapter II.

For this paper, any resolution equal to or below 400 m will be considered high-resolution, while any resolution above 400 m will be considered low-resolution. The technique presented in this paper can be applied to high or low-resolution NIR images. The ACAM presented in this paper is designed to supply scene information to decision makers as part of intelligence, surveillance, and reconnaissance (ISR). The technique uses NIR satellite imagery in conjunction with spectral and geometric thresholds to determine cloud cover. Threshold-based cloud detection has been used in the past with success, but thresholds require frequent adjustments to resolve cloud cover correctly. This requires regular threshold adjustments to optimize cloud detection, limiting this technique's practicality (Kidder and Vonder Haar 1995). The ability of the system presented in this paper to rely on seasonally set thresholds will be investigated. The ACAM is evaluated with imagery from two radiometers to

demonstrate its performance with different platforms and multiple resolutions. The underlying algorithm is scalable and flexible. It can be operated using a GUI or by command line. Its modular design also increases its adaptability, making it open to include other feature identification modules to be developed in future work. A cloud detection technique that relies on data derived from one satellite channel and seasonal threshold parameters is desirable because it is also less complicated and more computationally efficient than common multispectral techniques. This drives down overall cost as well.

1. Definition of Near Infrared Satellite Imagery

NIR satellite imagery is commonly used to analyze the atmosphere. Space-based radiometers supply scientists with vantage points not possible with other in-situ and remote instrumentation. Specialized weather satellites monitor carefully selected channels of the electromagnetic spectrum (EM) tuned to best resolve atmospheric features. Wavelengths of light within the NIR spectrum reveal details about atmospheric elements that are not detected in other bands, such as the phases of cloud particles. The lower sensitivity of this part of the spectrum to certain atmospheric gases also allows the borders between clouds, land, and water to stand out more readily. Computer-based quantitative analyses of these datasets can be used to increase the quality of the analysis process or reduce the amount of labor involved. Given the right conditions, such algorithms can also automate the detection of clouds.

Sources of NIR imagery with meteorological value are available from many different platforms that span several scientific disciplines. Widespread use of infrared imagery began with the introduction of Kodak’s revolutionary NIR-sensitive film during the Second World War (Rogalski 2010). The film’s success paved the way for electronic scanning radiometers capable of capturing and preserving image data

in the early 1950s and launched into orbit during the 1960s. Flown aboard the Defense Meteorological Satellite Program (DMSP) and Television Infrared Observing Satellite (TIROS) polar-orbiting systems, they captured NIR radiance from Earth and its atmosphere both day and night. These satellites offered daytime-like images with resolution near that of visible light during moonlit nights (Hall 2001).

All physical matter exhibits radiance, assuming its temperature is above absolute zero. The magnitude of radiance depends on each object's temperature and emissivity at a certain wavelength. Elements that act as perfect absorbers and emitters of radiation are referred to as blackbodies. They are assumed to give off the maximum value of radiation possible for their unique makeup across all wavelengths. Many natural features approximate blackbodies at certain wavelengths. By employing radiometric observations, brightness value emitted at specific wavelengths can be determined through Planck's Law given as

$$B_{\lambda}(T) = \frac{2hc^2\lambda^{-5}}{\exp(\frac{hc}{\lambda kT}) - 1}, \quad (1)$$

where emitted blackbody radiation (B) is given as $Wm^{-2}\mu m^{-1}sr^{-1}$ and will be referred to as observed brightness value. The brightness value of a feature at one wavelength is a function of temperature (T) observed at that wavelength. The equation also relies on Planck's constant ($h = 6.63 \times 10^{-34}Js$), the speed of light in a vacuum ($c = 3 \times 10^8 ms^{-1}$), the wavelength of interest (λ), and Boltzmann's constant ($k = 1.38 \times 10^{23}JK^{-1}$); (Evans et al. 2014).

To measure total radiance of a blackbody over an interval of wavelengths, the Stefan-Boltzmann Law is employed using

$$M_{BB} = \int_0^{\infty} \pi B_{\lambda}(T) d\lambda, \quad (2)$$

which measures radiant exitance over all possible wavelengths ($\int_0^{\infty}$). It determines radiant exitance by employing π and brightness value ($B_{\lambda}(T)$), integrated over all

possible wavelengths ($d\lambda$). Emissivity is the unitless ratio of radiation emitted by an object to that emitted by its ideal blackbody counterpart. Emissivity is given by

$$\epsilon_\lambda = \frac{M_\lambda}{B_\lambda(T)}, \quad (3)$$

where M_λ is the observed radiance of an object and $B_\lambda(T)$ is the ideal blackbody radiance value, with units of $Wm^{-2}\mu m^{-1}$. Kirchoff's Law ($\alpha_\lambda = \epsilon_\lambda$), relates absorption (α_λ) to emission (ϵ_λ) and defines them as equal for a blackbody (Kidder and Vonder Haar 1995).

In the NIR spectrum, observed radiance values are comprised of both reflected and emitted radiation. The degree to which both variables contribute to the total radiance value depends upon a feature's composition. Radiance also varies with wavelength. For the purpose of this paper, NIR satellite imagery is defined as any scene captured by a space-based radiometer where the instrument's field of view (FOV) is parallel to the Earth's surface at nadir and each pixel in the resulting image corresponds to an observed radiance value.

Understanding what surfaces look like at a certain wavelength is valuable when interpreting NIR imagery, but understanding where NIR fits within the EM spectrum is equally as important. The infrared portion of the light spectrum spans wavelengths from $0.7 \mu m$ to $1 mm$ (Campbell and Wynne 2011). Subdivisions are generally characterized as being NIR ($0.7 \mu m$ to $3.9 \mu m$), middle infrared (MIR) ($3.9 \mu m$ to $7.0 \mu m$), or far infrared ($7.0 \mu m$ to $1 mm$). Far infrared may also be referred to as thermal infrared (TIR). In this paper, MIR and TIR will be treated as two neighboring categories. Additionally, shortwave infrared may constitute a portion of the spectrum generally called NIR, occupying in the portion of the spectrum from $2 \mu m$ to $3.9 \mu m$. These divisions vary between disciplines and their application to different wavelengths has also changed over time (Hecker et al. 2010). The locations where different types of infrared light fit within the EM spectrum can be seen in

Figure 1. In this paper NIR will refer to wavelengths from $0.70\ \mu\text{m}$ to $3.9\ \mu\text{m}$ in the electromagnetic spectrum.




Figure 1: Infrared subdivisions of the EM spectrum. Adapted from Tattersall (2016).

2. Definition of Satellite-Based Change Detection

Satellite-based change detection is the process by which two or more images of the same scene are compared to reveal changes over time. One image is defined as the normal or background state for a given scene. This background image is then compared to images of the same scene for other points in time. Comparing two images is known as bitemporal change detection. Comparing more than two images is known as multi-temporal change detection. Pixels values that have changed are identified either by a computer system or manually by an analyst (Campbell and Wynne 2011).

There are five main types of change detection (Kidder and Vonder Haar 1995). The first method is visual interpretation by an analyst. It is the oldest and most labor intensive. Image algebra is the second technique. Pixels values from the same location across the entire scene in both images go through mathematical treatment.

Output values from each pixel-to-pixel calculation are recorded as a new image. The resulting pixel values can then be used to determine how features in the scene have changed. Transformation and data reduction is the third type of technique. It involves altering the images using mathematical functions, then assessing the resulting images to determine what changes have occurred. Classification is the next technique. It involves creating bins for all possible pixel values in a scene. Each bin corresponds to a category for some type feature from the scene, such as a roadway or forest. The method identifies areas of change from the scene and attributes the change one of the categories. The statistical method is the fifth type of technique. The most common statistical techniques involve comparing individual pixels or clusters between images. A correlation equation is used to determine how similar each point in some image is to the background image. The standard deviation for each pixel location between the two images can also be used to detect changes within a given scene. In the current research, the technique developed by the author relies on algebraic change detection, but visual interpretation and statistical techniques were used throughout the development of the methodology, as well as for each case study.

3. Definition of Cloud Cover

Detecting cloud cover presents many challenges to both algorithms and analysts. Clouds are composed of water droplets in some form that are visible and suspended in the air. There are many possible ways for this definition to be met and not all of them can be readily observed by radiometers. Clouds come in different shapes, sizes, phases, thicknesses, and opacities. All of these factors contribute to their ability to be distinguished by a radiometer, given its wavelength and spatial resolution. The main types of clouds that will be referred to throughout this paper are cumulus, stratus, and cirrus.

Cumulus clouds are usually white, optically bright clouds when viewed in NIR imagery because they are highly reflective. They have a rough, cauliflower-like appearance. They may be small or large and can occur singularly or in clusters on an image. Stratus clouds are clouds that can appear less optically bright than cumulus, often exhibiting a gray surface but can also be white in NIR imagery. They are not usually as reflective as cumulus clouds. Stratus clouds are large with smooth surfaces and often occur in layers. Cumulus and stratus clouds are mostly composed of liquid water droplets exhibiting a high enough density to ensure that their cloud appears opaque to a given radiometer. With NIR imagery, clouds that reflect the most radiation appear bright white.

Cirrus clouds are composed primarily of ice crystals. They have smooth surfaces and may exhibit an elongated shape, with their longest axis oriented parallel to the flow of wind. Their edges are usually thin and fibrous. They can be opaque, but are often partially translucent which allows radiation from below to be transmitted through the cloud. This causes cirrus clouds to blend in with the features below and complicates the process of locating them using automation. In this paper, clouds that are not opaque will be referred to as optically thin (Kidder and Vonder Haar 1995). Because optically thin clouds can supply information about the scene below as part of ISR, identifying of these features as cloud cover does not add value in all scenarios. The ability for this study's ACAM to detect opaque cloud cover is desirable, but an inability to detect optically thin clouds does not necessarily constitute a weakness. If elements of a scene below the cloud base could be manually distinguished by an analyst and the corresponding pixels are flagged as cloud cover, end-users of the product could overlook valuable information in an image.

Clouds observed in NIR imagery can fit into one or more of the of the categories above. Cloud brightness in this segment of the EM spectrum is attributed primarily

to reflected sunlight. Emitted radiation from clouds contributes to their observed radiance to a lesser extent. This is opposite of what is observed in MIR and TIR imagery where the brightness of objects is related primarily to their thermal properties. These types of images are inverted so that the coldest clouds appear white instead of black (Kidder and Vonder Haar 1995).

Clouds appear in shades of gray and white in NIR imagery because they reflect more sunlight at this wavelength than at MIR and TIR wavelengths (Rogalski and Chrzanowski 2014). Cloud surfaces may not always be continuous and can have boundaries that form complex shapes. Like surface textures, optical brightness may not always appear uniform either. When measured by radiometers, clouds may completely or partially fill a pixel corresponding to an area observed by the instrument. Partially filled pixels are referred to as contaminated and cannot be used to reliably determine the spatial extent of cloud cover. Contamination may also refer to the transmission of radiation from below a cloud up to a radiometer. Clusters of small clouds may or may not be easily delineated from their surroundings in this way. Just after forming, low-lying cumulus clouds with shallow vertical extents can also be difficult to distinguish from their surroundings, like optically thin cirrus, and suffer from contamination. The observed radiance values of those pixels are not representative of the cloud's radiative properties.

All three types of clouds can occur at different sizes, including those sizes below a radiometer's spatial resolution. It is important to delineate the spatial size at which identifying something as a cloud adds value to an analysis. While a one-meter square area of visible moisture constitutes a cloud by definition, its meteorological significance or value to an aviation hazard chart would be small. Depending on the nature of the product's application, inclusion of the small cloud would decrease the value of the product by introducing clutter. For this paper, any clouds below the minimum

cloud area threshold (MCAT) will not be considered operationally significant. While these small clouds may be resolved by masks used for truth data, the ACAM is only required to resolve clouds above the MCAT. The ACAM’s capabilities and settings will be discussed further in Chapter III.

B. Research Objective and Technique Preview

1. Research Objective

The goal of this document is to demonstrate a new, relatively simple ACAM that works with NIR wavelengths. The final product of the system will be a simple cloud mask file delineating cloudy from cloud-free pixels. A byproduct of this process will be a synthetic cloud-free background image for each area of interest (AOI). Cloud masks will be verified against the standard cloud masks provided by each instrument’s sponsoring agency. This project is part of a larger initiative that aims to resolve issues with detecting and characterizing phenomena in infrared imagery, such as sun glint and thermal crossover. An accurate and reliable ACAM system is necessary to delineate pixels from a scene where these and other phenomena are occurring from those where they are not. The system presented in this document is modular by design so that it can be further developed to detect, then diagnose cloud types within the scene along with other transient non-cloud features meeting MCAT requirements.

The MCAT is defined both by end-user requirements and physical limitations of the sensor. The higher the spatial resolution is, the less influenced by resolution the MCAT becomes and the more user-driven the variable will be. Image sets used to demonstrate the technique were selected to demonstrate this principle and highlight the flexibility it brings to the analyst. If all features exceeding the MCAT can be identified within a scene, the problem of detecting and characterizing the remaining

features is simplified. The modular design of the system allows new detection capabilities to be added over time. The system iterates through each connected module, with each component adding further detail to the output mask for each detected feature. The final iteration of the project results in a monochromatic feature identification system capable of characterizing each pixel in an observed image and outputting masks for each phenomenon.

2. Technique Preview

The method for detecting cloud cover described in this document fuses imagery analysis tools and techniques from the fields of meteorology and geographic information science (GIS). The system detects cloud cover in a NIR image using spatial and spectral thresholds. It uses these thresholds in conjunction with a synthetic cloud-free reference image to detect changes from a scene’s background state. The reference image depicts a cloud-free version for a given AOI during a certain time of the year and is synthesized using a technique that is also described in this document. The method is also monochromatic, meaning that it only requires one channel to work, and handles resolution as a variable (Campbell and Wynne 2011). Pixels in the observed image that exceed the values of the co-located reference image by a defined value are flagged as possible cloud cover by the system, then clusters of pixels that exceed a user-defined size threshold are labeled as cloud cover.

The observed radiance of a scene can vary with the time of year due to seasonal changes tied to weather patterns, growing seasons, and sun angle variations (Campbell and Wynne 2011). Background images represent one-month periods in this document unless otherwise stated. This avoids issues tied to seasonal differences in radiance. It allows features that exhibit large spatial and spectral changes over different seasons, such as water levels and crop cover, to be captured in a background image while

minimizing the impact of features that vary on a daily basis. The more representative the background image for a given period of time becomes, the more reliable the output from change detection techniques will be as a result.

The type of change detection used in this study is bitemporal. Two images taken at different times are used to infer details about the evolution of features within a scene. An initial image is chosen where elements of interest, such as crops or construction sites, have been identified by an analyst. The second image is from either an earlier or subsequent period in time. The first image is subtracted from the second, and the resulting shifts in pixel values across the scene can be used to infer what has changed between the images. Usually, a color ramp is applied to the resulting image to aid in the identification of changes across the scene (Campbell and Wynne 2011). An example of bitemporal change detection method is displayed in Figures 2 and 3. The example highlights changes in ship locations across the depicted marina. In the current research, this approach is used to determine where clouds are located.




Figure 2: NIR view of a Los Angeles, CA marina on (a) June 18, 2016 and (b) the same location 10 days later.

The technique presented in this document is demonstrated with three types of passive NIR image products from two infrared radiometers. The first two products are daytime and nighttime images from the VIIRS sensor. This instrument flies on board




Figure 3: Image of change produced by subtracting (a) from (b). Shades of red indicate where pixel values have decreased. Shades of blue show where pixel values have increased. Dark shades indicate pixels where the most change has occurred, such as those associated with ship locations.

the Suomi National Polar-orbiting Partnership (Suomi-NPP) satellite (Wolfe et al. 2013). The third product is daytime imagery from the Sentinel-2 MSI (ESA 2015b). There are two operational Sentinel-2 satellites flying with the instrument on board. The twin satellites orbit at the same altitude, but have orbital positions 180° out of phase relative to each other. All three satellites are polar-orbiting, sun-synchronous, and sense multiple infrared channels. VIIRS offers imagery both day and night, while the MSI only offered imagery during the day at the time this document was written. NIR channels for both daytime datasets were chosen to have the same central wavelength of $0.865 \mu m$. This allows for cross-platform data exploration. VIIRS was chosen for its high temporal resolution, despite its relatively low spatial resolution. This ensures a large set of imagery is available for synthesizing background imagery.

Its images are available for most landmasses on a daily basis at a spatial resolution ranging from 375 to 750 *m*. The list of channels and resolutions for VIIRS are given in Table 1. The channels used to demonstrate the technique developed for this paper are I2 and DNB. They will be discussed during Chapter III.

Table 1: Channel information for VIIRS. There are two primary types of channels: moderate resolution bands (M) and imagery resolution bands (I). There is also one DNB that is optimized to sense imagery both day and night. Channels used to construct the operational cloud mask are marked with an x. Horizontal spatial resolution (HSR) at nadir is the resolution sensed at the point directly below the sensor’s field of view (FOV); (NOAA 2017b).

CHANNEL	CENTRAL WAVELENGTH (μm)	HSR AT NADIR (<i>m</i>)	CLOUD MASK
M1	0.412	750	X
M2	0.445	750	
M4	0.555	750	
I1	0.64	375	X
M5	0.672	750	X
DNB	0.7	750	X
M6	0.746	750	
I2	0.865	375	
M7	0.865	750	X
M8	1.24	750	X
M9	1.378	750	
I3	1.61	375	X
M10	1.61	750	
M11	2.25	750	X
M12	3.7	750	X
I4	3.74	375	X
M13	4.05	750	X
M14	8.55	750	X
M15	10.763	750	X
I5	11.45	375	X
M16	12.013	750	X




Figure 4: Channel information for the MSI. The vertical axis displays spatial resolution in meters. The horizontal axis displays channel wavelength in nanometers. Rectangle width approximates the spectral width of each channel (ESA 2015a).

MSI imagery is available at least twice a month for most landmasses across the planet. It was chosen for analysis of the cloud detection technique due to its high spatial resolution, despite its low temporal resolution. While the background images synthesized using this dataset come from a smaller sample set, the significantly higher spatial resolution allows the new technique’s ability to handle varying resolutions to be demonstrated. The wavelengths and resolutions for each MSI channel is shown in Figure 4. Channel 8a was chosen to demonstrate the technique because of its location within the NIR band. Its narrow spectral width also decreases the number of undesired environmental absorbers, such as atmospheric gases, that negatively impact transmission (Evans et al. 2014).

Background images were synthesized for each test site during months where cloud cover was minimal to ensure a sufficient sample set from which to generate each background image. The sites include locations with snowcapped mountains, snow-covered ground, and various water bodies during either warm or cold seasons

to demonstrate the flexibility of the technique. All automated tasks were designed, tested, run on an off-the-shelf computer system using GIS and imagery analysis software. Hardware and software will be covered in further detail in Chapter III. The imagery from both VIIRS (NOAA 2017a) and MSI (ESA 2017a) were obtained from publicly available data archives.

II. Related Research and Background

Spaced-based NIR instruments come in two varieties: those that face the Earth and those that face away from it. Terrestrial systems monitor the planet’s environment, while celestial systems monitor other elements of the universe. One example of a space-based terrestrial sensor is the NIR radiometer on board a weather satellite. These are used to measure properties of planet’s atmosphere. They reveal details about the state of the atmosphere, as well as underlying processes supporting each feature within a sensor’s FOV. Satellites capture images for different channels that correspond to different bands within the electromagnetic spectrum. Channels are chosen to highlight various atmospheric features and can also be tuned to focus on specific levels of the atmosphere.

Different NIR channels are used to highlight various aspects of the planet’s environment or atmosphere while suppressing features or phenomena that are not of interest to an analyst. Manual analyses have a high degree of accuracy, but can be time-consuming. The task of analyzing can also be laborious and repetitive which can eventually lead to a decrease in analysis quality. To increase the speed at which imagery can be analyzed and decrease analyst fatigue, a number of ACAMs have been developed. Most fall within one of seven method types: threshold, histogram, pattern recognition, multispectral, spatial coherence, geometric, and hybrid (Kidder and Vonder Haar 1995). Two more categories exist: radiative transfer and platform-specific. These methods generally rely on external data sources and will be treated as hybrid techniques. Because concepts from many of these methods were examined and adopted during the design of the ACAM presented in this document, each category will be reviewed. Standard VIIRS and MSI cloud masks will also be discussed. Additionally, any examples pertinent to the development of the technique presented in this document will also be presented.

A. Cloud Detection Techniques

1. Threshold Techniques

The threshold technique is one of the simplest methods of automating cloud cover identification within an image. It requires little computing power, which made it easy to program and realistic to implement on early computer systems (Kidder and Vonder Haar 1995). It is also the oldest ACAM. The technique behind this method is simple. Software compares the brightness values of each pixel within an image to a threshold value. Those at or above the threshold are flagged as cloud cover. Those that fall below are given a non-cloud designator. For channels measuring thermal properties, pixels below a certain brightness temperature are flagged as cloud cover. With these channels, multiple thresholds can be defined to create a simple cloud characterization system.

Cloud heights can also be determined by comparing each pixel to temperature values from a reference sounding. The height in the sounding at which each pixel's temperature occurs corresponds to the approximate height of an observed cloud top. Representative soundings for the area need to be available for this technique to work. Additionally, thresholds must be tuned by the user over time to preserve output reliability. If thresholds are not properly set, cloud cover could be misdiagnosed. These systems require little other input to function, which reduces the computational cost of running them.

Threshold ACAMs can be run with many various resolutions and channels, but are associated with two primary weaknesses: clouds smaller than one pixel are difficult to resolve and non-cloud features in some scenes can be flagged as cloud cover due to the limitations of setting a numerical threshold. For NIR imagery, thresholds are especially problematic in cloudy scenes with other high-albedo surfaces. Knowledge of clear-sky radiances can allow thresholds to be set more easily. Seasonal

variations in parameters such as sun angle, snow cover, vegetation can still lead to problems even when cloud-free radiance values are known. The more complex the scene becomes, the more difficult it is to set thresholds. Averaging co-located pixels from observed images similar in time to the analysis period can be used to determine clear-sky radiance values also, but finding this values can be difficult to automate. If an operator is required to tune threshold values too often, the value added by automating the process can be diminished as well because of the labor associated with determining new threshold values.

Due to their susceptibility to contamination, optically thin cirrus clouds are not always classified correctly. For systems that classify cloud types, layered clouds can also result in contamination that leads to misclassification of cloud cover. Generally, these systems are used to determine locations of cloud cover and cloud top heights. This type of technique can be combined with others to reduce the amount of manual input and increase product quality. These approaches will be discussed later with hybrid techniques.

2. Histogram Techniques

Similar to threshold techniques, histogram techniques look for pixels clusters in imagery histograms that correspond to cloud cover (Kidder and Vonder Haar 1995). Clusters of pixels in each histogram are identified using one of several methods. Then, the positions of pixel clusters with similar values on the histogram can be used to determine what physical feature they correspond to. This technique generally uses at least two channels. For each histogram, pixel values from one or more different channels are placed along each axis. For instance, if the histogram uses visible and TIR channels, bright and cold clouds would appear in the upper right corner of the

diagram, while warm and dark features would appear as pixel clusters in the lower left. An idealized example can be seen in Figure 5.

Clusters corresponding to features can be identified manually or through automated means. Manual methods work well for small datasets, but are not desirable for large datasets or for real-time analysis of satellite data. Alternatively, thresholds can be used to define bins for separate clusters of pixels corresponding to one feature from another to automate the process. Bin ranges can also be derived from pixel clusters identified in previous images as well. The average temperature and brightness values from each cluster can be used to determine the types of clouds in the scene as well. Lastly, the number of cloudy pixels is divided by the total number of pixels in the scene to determine cloud coverage. This process does not reveal where clouds are located in the original imagery though. While specialized software can be used to automatically characterize pixels with the identified clusters in the original image, clusters and their thresholds must be reevaluated for each new image being processed. Histograms were used during development of the ACAM presented in this document to examine the radiance characteristics of clouds, but are not employed in the final system due to their inability to relate pixel clusters to spatial locations for large datasets.




Figure 5: An idealized two-dimensional histogram with infrared brightness temperature along the vertical axis and visible radiance along the horizontal axis. Pixel clusters are labeled according to their corresponding features. Adapted from Kidder and Vonder Haar (1995).

3. Pattern Recognition Techniques

Pattern recognition techniques examine imagery for cloud-like patterns in bright or cold pixels (Kidder and Vonder Haar 1995). They work by determining the spatial variance of pixels relative to one another across a given scene. A training dataset containing key values used to recognize cloud patterns is required for these techniques to work. A combination of spectral and texture parameters are then employed to identify patterns within each scene. Spectral parameters are those that refer to radiance measurements of pixels. Texture parameters refer to the spatial variance of pixel values for a given scene. These parameters are then used to identify clusters of pixels from within the scene and diagnose cloud cover. While pattern recognition techniques overcome many of the drawbacks associated with the first two methods, training data must be kept up to date and representative of the scenes being analyzed for them to be successful. Texture parameters were used to validate synthetic imagery generated for the ACAM presented in this document.

4. Multispectral Techniques

Multispectral techniques are another way to automatically determine cloud cover. They use pixel data gathered from at least two channels. These channels are not limited to any particular part of the spectrum, such as visible or infrared, but they must be from the same instrument and scan spot. Similar to how various features or atmospheric layers can be targeted by certain satellite channels, various combinations of channels can be used to highlight specific environmental details (Kidder and Vonder Haar 1995). Techniques that use more than one satellite will be discussed later. An example multispectral image can be seen in Figure 6.

ACAMs in this category can be run manually or through automation to identify cloud cover. Pixel differencing is one of the most common manual techniques. With




Figure 6: Multispectral image of Los Angeles, California captured by the MSI on November 30, 2017. This false-color image assigns the visible green, visible red, and NIR channels to the monitor's blue, green, and red color guns, respectively. This channel arrangement reduces haze due to the absorption of incoming radiation by optically active gases in the atmosphere while highlighting features of interest. Clouds are shown as white, vegetation appears red, and bodies of water appear dark blue or black. Anthropogenic structures appear in light shades of blue. Identifying cloud cover with this enhancement must be done manually (Campbell and Wynne 2011).

this approach, pixels from one or more channels are added to or subtracted from co-located pixels from another. Alternatively, monitor color guns can be assigned to instrument channels to create color patterns that allow various features to be recognized. False-color images use wavelengths of visible light to represent reflectivity or radiance from other wavelengths in the electromagnetic spectrum. Multispectral methods can also be used to produce gray-scale images. Pixel brightness and color

can be used to reveal the location and characteristics of features such as clouds. Aerosols such as airborne dust or volcanic ash can also be revealed using this approach. Depending upon the application of the imagery with manual analysis, this technique may also be referred to as multispectral enhancement.

The resulting image can be inspected by the analyst as a single product or multiple images can be used to create an animated loop to reveal further details about the features of interest. Automated methods may use pixel-to-pixel comparisons between channels or compare averaged values from groups of pixels to determine the location of cloudy pixels. Unlike the histogram technique, the advantage of this method is that clouds and their characteristics are assigned to geographic coordinates on a pixel-by-pixel basis. Drawbacks to these techniques vary with each approach, but the biggest constraints are the number of available channels for a particular instrument and the wavelengths the channels are centered on.

5. Spatial Coherence Techniques

Techniques from the spatial coherence category are the most limited in application compared to the other ACAMs. While their output is highly precise, each image analyzed must be completely clear and completely cloudy portions (Kidder and Vonder Haar 1995). This approach requires the use of automation, as multiple mathematical processes must be run on various areas of the supplied image must be compared. Other necessities include uniform cloud top heights and cloud top emittance, as well as homogeneous background characteristics beneath each area of cloud cover. This method is best suited for the characterization of marine cloud cover. It does not work with visible data or when a scene contains frontal features. It also cannot be used when cirrus clouds are present. Though the technique yields cloud cover data with cloudy pixel locations, the number of constraints limit its applicability.

6. Geometric Techniques

Geometric ACAMs provide accurate information on the location and tops of clouds by using two satellites that pass over the same location within a short period of time from different look angles (Kidder and Vonder Haar 1995). Knowledge the sun’s position along with each satellite’s position can be used to determine cloud information using trigonometry. In a similar way, stereoscopy can be employed to locate cloud cover. The displacement of cloud cover between the images from both satellites can be used to determine cloud height and position. When both images are viewed together with stereoscopic glasses, clouds can also be identified visually. The glasses enable the images to mimic three-dimensional figures making the clouds easy to distinguish from the background and are used during manual analysis. The biggest limitation of this type of technique is that each observed images can be separated by no more than 10 seconds. While stereoscopic analyses were employed during this project, no other elements from geometric ACAMs were used.

7. Hybrid Techniques

There are also ACAMs that combine different techniques to identify cloud cover. These use numerical weather prediction or radiative transfer models in conjunction with imagery to detect clouds or may rely on output from one type of technique to provide the information required for the rest of the technique (Kidder and Vonder Haar 1995). Platform-specific techniques also exist. Though reliable, they are usually designed to only work with a specific instrument or satellite. Hybrid techniques come in both manual and automated varieties. They have also become increasingly complex as technology has advanced.

Radiative transfer ACAMs are hybrid methods that use the output from threshold analyses in conjunction with radiative transfer calculations or models to determine

cloud properties. Additionally, these techniques require knowledge of cloud location before they can function. Other hybrid methods can use the output from weather models to determine likely cloud cover locations in conjunction with imagery data. Some techniques rely on two or more instruments or satellites to determine cloud location, such as those that use radar imagery in conjunction with infrared imagery to locate cloud cover. These techniques may employ geometric or spatial calculations in conjunction with imagery to determine cloud location as well. Readers should refer to Kidder and Vonder Haar (1995) for further information on hybrid or any of the other techniques listed above. The type of ACAM developed for this project qualifies as a hybrid technique and will be discussed in further detail in Chapter III.

B. Standard Instrument Cloud Mask Products

Cloud masks from both the VIIRS and MSI instruments were used during the course of this project for verification purposes. The MSI’s cloud mask is generated using a relatively simple multispectral technique (ESA 2017b). The method used by VIIRS is more complex (NASA 2014). Both techniques are discussed below.

1. MSI Cloud Masks

The multispectral technique employed by the MSI determines whether pixels are cloudy or cloud-free (ESA 2017b). It does not characterize cloud cover. All channels used for this technique come from the MSI. Pixels are flagged according to cloud type at 60 *m* resolution for all images, then radiometrically interpolated to higher resolutions. The final masks available from the MSI archive used in this paper are shapefiles that are later converted to images. The algorithm uses single and multi-channel threshold tests, along with differencing and ratio tests to detect possible areas of cloud cover. Differencing tests subtract two or more channels and compare the output to reference values. Ratio tests divide pixel values between two or more channels and compare the output values to reference values. The algorithm checks possible areas of cloud cover against snow mask data and removes areas of surface snow from the final mask. Radiometric interpolation is used to resample each cloud mask to higher resolutions. The algorithm also removes isolated pixels denoted as cloud cover and fills in all cloud-free gaps inside the identified clouds with geometric filters. Cloud masks produced when data from any of the channels used is missing or corrupted contain only null value pixels. Pixels valued as zero represent clear skies. Pixels flagged with values of one or two represent opaque clouds.

Opaque clouds can occur at various heights either in single or multiple layers. Cumuliform and stratiform clouds both belong to this category. They are defined by

high radiance in the visible blue band, which corresponds to channel B2 on the MSI (see Figure 4). High-albedo pixels from channels B11 and B12 are cross-referenced with co-located pixels in B2 to ensure that features that only reflect well in the blue band do not get labeled as cloud cover. Clouds exhibit high albedos at each of these wavelengths, but surface-level features such as snow do not.

Cirrus clouds generally occur at high altitudes, typically above 6 *km*. While they can be optically thick and detectable at certain visible wavelengths, they do not reflect highly in blue. Those clouds that are thin and wispy due to their ice crystal composition may not be perceived with visible imagery at all. They show up well on certain NIR wavelengths such as with the MSI B10 channel. When dense cloud cover is not detected in a pixel, it is checked with channel B10 for cirrus. While cirrus can be observed in B2 images, the clouds have lower reflectance than with channel B10. Pixels that meet these criteria are labeled as cirrus cloud cover. A workflow diagram of the operational MSI cloud detection algorithm can be seen in Figure 7. Further information on the algorithm can be found in the Sentinel-2 MSI Technical Guide (ESA 2017b).




Figure 7: Diagram of the MSI cloud detection workflow. The process utilizes information from multiple channels as well as external data (ESA 2017b).

2. VIIRS Cloud Masks

The ACAM utilized by VIIRS provides more detail about each cloudy pixel than the system employed by the MSI. It is a hybrid technique that incorporates information from multiple channels, reference datasets, and models (NASA 2014). It uses threshold, ratio, and differencing tests in conjunction with external reference and model data to determine the likelihood a given pixel is cloudy. Cloudy and cloud-free pixels are determined from likelihood values. In addition to flagging pixels as cloudy or not cloudy, the algorithm gives information on cloud particle phase, cloud opacity, and the fraction of each pixel covered by clouds. The system also derives aerosol content and detects cloud shadows. Each pixel flagged as cloud cover is assigned a confidence level. Categories include probably cloudy, probably clear, confidently cloudy, and confidently clear. Once completed each value assigned in the mask also undergoes a quality check. The same architecture is used for day and night cloud detection, but the night cloud detection algorithm uses a simplified version of the daytime method with fewer input channels (Hutchison et al. 2005). Mask pixel assignments are rated poor, low, medium or high quality. The workflow that determines these values can be seen in Figure 8. Both cloud detection methods involve complex data processing. A simple overview is presented in this document. Further details about both algorithms can be found in the Joint Polar Satellite System VIIRS Cloud Mask Algorithm Theoretical Basis Document (NASA 2014).

The system flags each pixel as one of six classes referred to as bytes. Each byte has eight possible values, also called bits. There are a total of 48 possible values each mask can contain as a result. Pixel value assignments for the VIIRS cloud mask system can be seen in Table 2.

Cloud masks can only be produced if brightness temperature and radiance data are available for the whole scene. For the system to function, the type of surface below the cloud cover must also be known. VIIRS cloud masks contain details on pixels flagged as non-cloud cover as well. The processes that characterize non-cloud pixels are beyond the scope of this project and will not be examined in detail. The methodology for detecting pixels with cloud cover will be discussed below. These are grouped according to cloud and cloud-free pixels to produce the binary cloud mask used in this paper.

The cloud detection portion of the VIIRS ACAM is a multispectral method. Pixel brightness values are compared between different band combinations to determine whether they are associated with clear or cloudy skies. Each differencing test targets a different type or level of cloud cover. Information derived from these tests is also used to characterize cloud cover. Bytes one and two contain cloudy and non-cloud pixel flags (see Figure 2). Tables 3 and 4 detail the methods used to determine cloud cover, given by channel, for both daytime and nighttime, respectively. An example VIIRS cloud mask can be seen in Figure 9.

Table 2: Pixel values for VIIRS cloud masks. Cloudy and cloud-free flags are grouped into separate classes to produce the binary cloud mask. Cloud-free pixels in the binary mask are assigned values of zero. Cloudy pixels are assigned values of one. Adapted from NASA (2014).

BYTE	Bit	Flag Description Key	Result
0	0-1	Cloud Mask Quality	00= Poor 01= Low 10= Medium 11= High
			00 = Confident Clear 01 = Probably Clear 10 = Probably Cloudy 11 = Confident Cloudy
	2-3	Cloud Detection Result & Confidence Indicator	0 = Night 1 = Day
	4	Day / Night	1 = Snow/Ice 0 = No Snow/Ice
	5	Snow / Ice Surface	00 = None 01 = Geometry Based 10 = Wind Speed Based 11 = Geometry & Wind
	6-7	Sun Glint	000 = Land & Desert 001 = Land no Desert 010 = Inland Water 011 = Sea Water (Oceans) 101 = Coastal
			1 = Yes 0 = No
1	0-2	Land / Water Background	1 = Yes 0 = No
	3	Shadow Detected	1 = Yes 0 = No
	4	Non Cloud Obstruction (Heavy Aerosol)	1 = Yes 0 = No
	5	Fire Detected	1 = Cloud 0 = No Cloud
	6	Cirrus Detection (Solar) (RM9)	1 = Cloud 0 = No Cloud
	7	Cirrus Detection (IR) (BTM15-BTM16)	1 = Cloud 0 = No Cloud
	0	IR Threshold Cloud Test (BTM15)	1 = Cloud 0 = No Cloud
2	1	High Cloud (BTM12 - BTM16) Test	1 = Cloud 0 = No Cloud
	2	IR Temperature Difference Test (BTM14 - BTM15 & BTM15 - BTM16)	1 = Cloud 0 = No Cloud
	3	Temperature Difference Test (BTM15 - BTM12)	1 = Cloud 0 = No Cloud
	4	Temperature Difference Test (BTM12 - BTM13)	1 = Cloud 0 = No Cloud
	5	Visible Reflectance Test (RM5)	1 = Cloud 0 = No Cloud
	6	Visible Reflectance Test (RM7), also used for Visible Reflectance Test (RM1)	1 = Cloud 0 = No Cloud
	7	Visible Ratio Test (RM7/RM5)	1 = Cloud 0 = No Cloud
3	0-1	Adjacent Pixel Cloud Confident Value	00 = Confident Clear 01 = Probably Clear 10 = Probably Cloudy 11 = Confident Cloudy
	2	Conifer Boreal Forest	1 = Yes 0 = No
	3	Spatial Uniformity Test	1 = Yes 0 = No
	4	Dust candidate	1 = Yes 0 = No
	5	Smoke candidate	1 = Yes 0 = No
	6	Dust/Volcanic Ash	1 = Yes 0 = No
	7	SPARE	
4	0-7	SPARE	
5	0-2	Cloud Phase	000 = Not Executed 001 = Clear 010 = Partly Cloudy 011 = Water 100 = Mixed Phase 101 = Opaque Ice 110 = Cirrus 111 = Cloud Overlap
	3	Thin Cirrus Flag	1 = Thin Cirrus 0 = None
	4	Ephemeral Water Flag	1 = Yes 0 = No
	5	Degraded TOC NDVI Flag	1 = Yes 0 = No
	6	Degraded Sun Glint Flag	1 = Yes 0 = No
	7	Degraded Polar Night Flag	1 = Yes 0 = No




Figure 8: Operational cloud detection workflow for VIIRS daytime cloud masks. Nighttime cloud masks are produced from a simplified version of the same algorithm. Adapted from Hutchison et al. (2005).

Table 3: List of resources used by the daytime VIIRS cloud classification algorithm. Applicable land surfaces for each scheme are marked with an 'x' (NASA 2014).

Cloud Spectral Tests Used to Determine Cloud Confidence	Water	Land	Desert	Coast	Snow
1. M9 (1.38 μm) Reflectance Test	X	X (if TPW > cutoff cm)	X (if TPW > cutoff cm)	X (if TPW > cutoff cm)	X
2. M15-M16 (10.76 μm – 12.01 μm) Brightness Temperature Difference (BTD)	X	X	X	X	X
3. Tri-Spectral M14, M15, M16 (8.55, 10.76, 12.01 μm) BTD Test	X				
4. M15-M12 (10.76 μm -3.70 μm) BTD Test	X (if no sun glint)	X (if TOC NDVI > 0.2)	X (if Lat > 60° or < -60°)	X (if no sun glint and if TOC NDVI > 0.2)	X
5. M12-M13 (3.70 μm -4.05 μm) BTD Test	X (if no sun glint)	X (if TOC NDVI > 0.2)			X
6. M1 (0.412 μm) Reflectance Test			X (if -60° < Lat < 60°)		
7. M5 (0.672 μm), M1 (0.412 μm) Reflectance Tests		X (M5 if TOC NDVI $\geq$ 0.2, M1 otherwise)		X (M5 if TOC NDVI $\geq$ 0.2, M1 otherwise)	
8. M7 (0.865 μm) Reflectance Test	X				
9. M7/M5 (0.865 μm / 0.672 μm) Reflectance Ratio Test	X	X (if RefM5 $\geq$ LD_M5_Gemi Thresh)			
Cloud Spatial Tests Used to Modify the Final Cloud Confidence Classification	Water	Land	Desert	Coast	Snow
10. I5 (11.45 μm) Spatial Test	X				
11. I2 (0.865 μm) Reflectance Test	X				




Figure 9: Example cloud mask overlaid atop an observed NIR image of California from May 26, 2017. The cloud mask is translucent and yellow.

Table 4: List of resources used by the nighttime VIIRS cloud classification algorithm. Applicable land surfaces for each scheme are marked with an 'x' (NASA 2014).

Cloud Tests	Water	Land	Desert	Coast	Snow
1. M15-M16 (10.76 μm – 12.01 μm) Brightness Temperature Difference (BTD)	X	X	X	X	X
2. M15 (10.8 μm) Brightness Temperature (BT) Test	X	X	X	X	X
3. M12-M16 (3.70 μm – 10.76 μm) BTD Test		X M12 > BTM12 limit	X M12 > BTM12 limit	X M12 > BTM12 limit	X M12 > BTM12 limit
4. Tri-Spectral M14, M15, M16 (8.55, 10.76, 12.01 μm) BTD Test	X				
5. M15-M12 (10.76 μm -3.70 μm) BTD Test	X	X if TOC NDVI > NDVI limit & M12 > BTM12 limit	X if TOC NDVI > NDVI limit & M12 > BTM12 limit	X if TOC NDVI > NDVI limit & M12 > BTM12 limit	X
Cloud Spatial Tests Used to Modify the Final Cloud Confidence Classification	Water	Land	Desert	Coast	Snow
6. I4 (3.74 μm) Spatial Test	X				
7. I5 (11.45 μm) Spatial Test	X				

C. Related Research

Accurate background images are pivotal to the technique presented in this document. A cloud-free background image representative of the average state of a scene cannot be captured by a single remotely sensed image. It must be derived from many images. This type of image will be referred to as a synthetic background image. For the new method presented in this document to work, each synthetic image must be free of features that vary significantly on a daily basis. To increase the value of each background image, attempts were also made to create reference images that could be adjusted to work with multiple platforms and for multiple months.

Zhu et al. (2015) demonstrated a process for generating multispectral synthetic background images using observed data from the Landsat thematic mapping satellite. The satellite used for the study were sun synchronous and carried radiometers capable of sensing multiple channels including visible and NIR. The generated images were used to detect surface cover changes. The imagery was generated by determining the average clear-sky spectral values for each pixel across the domain of interest. The output imagery is free of clouds, cloud shadows, and missing or damaged pixels. The system removes scene variations due to changing sun angles, vegetation life cycles, snow, and other seasonally varying parameters. A reference image for each wavelength is derived from multiple clear sky observations for each pixel, consistent with the normal background state of the scene. It uses radiative transfer calculations and multispectral data to generate a synthetic background image representative of a user-specified day of the year.

The equation uses a series of coefficients to determine daily differences for each pixel. For each channel used by the system, a series of spectral coefficients, date values, numerical thresholds determine average pixel values. Three models were devised

for generating synthetic imagery. The simplest determined final pixel values from the reference image using

$$\rho(i, \chi) = a_{0,i} + a_{1,i}\cos(\frac{2\chi\pi}{T}) + b_{1,i}\sin(\frac{2\chi\pi}{T}) + c_{1,i}\chi, \quad (4)$$

where the output pixel value ρ is a function of Julian date χ and Landsat channel i . Coefficients for intra-annual radiance change at the i th band are represented by $a_{1,i}$ and $b_{1,i}$. The overall value for each reference pixel value is represented by $a_{0,i}$. The fourth coefficient, $c_{1,i}$, represents the slope of inter-annual radiance change for the pixel. There are two other versions of the model that use mathematical variants of Equation 4. These eliminate $c_{1,i}\chi$ from the equation, add additional modes to each radiation coefficient, and alter the scalars in each trigonometric function. The document revealed that the model that used three modes for each coefficient was most reliable. A comparison between an observed Landsat multispectral image and synthetic image can be seen in Figure 10. The ACAM presented in this document will utilize similar techniques to those published by Zhu et al. (2015), but only for a single channel. Unlike Zhu et al. (2015), the method presented in this document aims to preserve information about features such as seasonal snow cover to aid in delineating clouds from non-cloud features rather than to suppress them. This will be discussed further in Chapter III.




Figure 10: False-color image of an urban environment and surroundings from May 4, 2002 in (a). Vegetation appears in red and human infrastructure appears in shades of blue. Water is black. Synthetic image generated for the same date in (b). The numerical difference in pixel values between the first two panels in (c). Histogram in (d) is the differences between pixels from (a) and (b); (Zhu et al. 2015).

III. Methodology

A. Dataset Overview

Images generated from VIIRS and MSI measurements have different file formats. VIIRS imagery must be generated from two Hierarchical Data Format 5 (HDF5)-formatted component files (Campbell and Wynne 2011). These files are obtained from the Comprehensive Large Array-data Stewardship System (CLASS); (NOAA 2017c). The first file contains radiance data for a specific VIIRS channel. The second file contains spatial information about the radiance file. These files are used to generate georeferenced tagged image file format (GEOTIFF) files, where each pixel is associated with geographic coordinates. Sentinel-2 data is available as georeferenced JP2 image files from the European Space Agency (ESA) via the Copernicus Open Access Hub (ESA 2017a). No special operations are required to view Joint Photographic Experts Group 2000 (JP2) files, but they are converted to GEOTIFF files before being processed by the ACAM. Clouds detected by the ACAM using VIIRS data is verified using standard VIIRS cloud mask files formatted as files. MSI cloud mask files are formatted as geography markup language (GML) vector files (ESA 2017b). All images and masks analyzed in this study ACAM are converted to the GEOTIFF format. Operational cloud masks identify clouds no smaller than the horizontal resolution of the coarsest channel used during mask production. Masks created for use as synthetic truth data will be addressed in section F of this chapter.

1. VIIRS Datasets

VIIRS images are synthesized using POLAR2GRID, a command-line-based Linux program. The synthesis was performed in batches using a shell script that searches a specified directory for image component files with matching date-time des-

ignators and inputs them into a command line statement readable by POLAR2GRID. Additionally, the script automatically detects whether daytime or nighttime imagery is contained in the directory of files to be batch processed. It then generates a GEO-TIFF file from each pair of image component files. Rapidly varying terrain combined with the high speed with which the satellite passes over a scan spot can result in warped images. For most areas National Atmospheric and Oceanic Administration (NOAA) provides terrain-corrected component files that minimizes this distortion (Wolfe et al. 2013). Terrain-corrected files are available for all case study sites except Alaska (NOAA 2017c). An example daytime image generated using POLAR2GRID can be seen in Figure 11.

1.1 Daytime Imagery

VIIRS daytime images in this study are all taken from the instrument’s NIR channel, I02 (NOAA 2017c). The channel’s central wavelength is $0.865 \mu m$. The component files are available at $350 m$ resolution, but the gridding process reduces resolution by $50 m$. Images are produced with $400 m$ resolution using POLAR2GRID’s High-Resolution Lambert Conic Conformal Dynamic Fit projection, the program’s highest resolution grid (Space Science and Engineering Center 2017).

1.2 Daytime Cloud Masks

Daytime cloud masks (see Figure 9) are available as $1 km$ resolution NetCDF-4 files (NOAA 2017c). They must be converted into GEOTIFF files for analysis. They are radiometrically interpolated from $750 m$ to $400 m$ resolution. They delineate cloudy from clear-sky pixels. Mask conversion will be discussed in Chapter IV.




Figure 11: NIR image of California from May 26, 2017 derived from HDF5 component images. California's border is shown in green (NOAA 2017c).

1.3 Nighttime Imagery

VIIRS nighttime imagery (see Figure 12) is available as two-component HDF5 files with 750 *m* resolution. It is taken from the DNB channel. This channel has a central wavelength of 0.7 μm . Gridding the data with POLAR2GRID converts the file to a 0.0057 deg resolution GEOTIFF. Though the day and night VIIRS cloud masks are determined from different workflows, they are made available at the same resolution and are treated as part of the same dataset by NOAA (NOAA 2017c).




Figure 12: NIR nighttime image of Florida and its surroundings on September 7, 2017. The US coastline is shown in yellow (NOAA 2017c).

1.4 Nighttime Cloud Masks

Nighttime cloud masks are available as 1 *km* resolution NetCDF-4 files (NOAA 2017c). These files contain both the categorized and binary cloud masks. The binary masks delineate cloudy from cloud-free pixels and are converted to GEOTIFF images for further analysis. Then the masks are radiometrically interpolated to match the 0.0057 deg resolution of DNB imagery.

2. MSI Datasets

MSI data is available in file formats that require minimal processing compared to VIIRS data. All images are distributed as gridded JP2 files with 20 *m* horizontal resolution (ESA 2017a). Cloud masks are available as vector files with 60 *m* resolution.

2.1 Imagery

MSI imagery (see Figure 13) is available for all major landmasses across the planet. The instrument captures imagery during day-lit hours at high spatial resolution. NIR channel 8A was used for this study. It is available at a resolution of 20 *m* with the same 0.865 μm wavelength as the VIIRS daytime dataset (ESA 2017a). JP2 image files are converted to GEOTIFFs before any processing by the ACAM is done.




Figure 13: NIR image of the region surrounding Bagram, Afghanistan from June 16, 2016 (ESA 2017a).

2.2 Cloud Masks

MSI cloud masks are available as GML-formatted vector files that must be converted to GEOTIFF images before use in the study. These are converted to GEOTIFF images with a horizontal resolution of 20 *m* (ESA 2017a).

3. Digital Elevation Models

Stereographic inspections are used during the course of this study for quality assurance during generation of synthetic background imagery. For locations with

rough terrain, synthetic images are stretched over digitally rendered three-dimensional topography derived from DEM files. These files contain elevation data used to create digital representations of topography that are used to give shape to other types of data (Patterson 2017). DEM files (see Figure 14) used for this study come from multiple vendors at different resolution and will be addressed individually as necessary.




Figure 14: Synthetic background image of Afghanistan with three-dimensional terrain features generated from a DEM (Patterson 2017). Bagram is marked with a green flag. The Afghan border is highlighted in red.

B. Hardware and Specialized Software

1. Hardware and Operating Systems

Two off-the-shelf computer systems are used to conduct the study. The first is a 2.2GHz Intel Core i7 MacBook Pro with four cores, 16GB of RAM, and a 250GB solid-state hard drive. Processing is done on this computer using the virtualized version of Windows 8.1. Eight processing cores and 12GB of RAM are made available to run Windows-based processes. This computer is primarily for raster analysis and running the ACAM. Onboard storage is supplemented with a 2 TB external hard drive.

The second computer is a 2.9 GHz Intel Core i5 iMac with four cores, 24 GB of RAM, and a 1 TB hard disk hard drive. It runs a virtualized version of the Fedora 20 Linux operating system that has access to all four processing cores and 16 GB of RAM, with 8 GB of RAM reserved for the stability of the computer's native operating system. This computer is used to synthesize VIIRS imagery with POLAR2GRID.

2. Specialized Software

Findings in this study are derived using a combination of imagery and geospatial analysis tools. These tools are leveraged primarily through two software programs: ArcMap and ERDAS Imagine. ArcMap is a geospatial analysis program, also referred to as a GIS. ArcMap allows remote sensing imagery and other types of geo-referenced data to be viewed and manipulated. Two additional programs are also used during the study: POLAR2GRID and ArcScene. POLAR2GRID is a standalone program for image synthesis. ArcScene is the three-dimensional counterpart to ArcMap (ESRI 2007b). ArcScene allows the analyst to build three-dimensional representations of a scene using data imported or generated within ArcMap. It is used for quality assurance. ERDAS Imagine offers an advanced analysis toolkit for satellite and radar images, also called rasters. This program will be referred to as Imagine for the remain-

der of this paper. Both programs allow their users to create automated workflows that allow for batch processing of raster datasets.

2.1 Polar2Grid

VIIRS imagery was synthesized using POLAR2GRID version 2.1, a Linux program that ingests satellite image component files and transforms them into commonly used image formats, such as JPEG and TIFF. The command-line-based program can ingest raw, preprocessed, or processed data from multiple U.S.-based instruments. It generates gridded, imagery for use in specialized software. The terms gridding and georeferencing both mean that data has been assigned geographic coordinates (Kidder and Vonder Haar 1995; Campbell and Wynne 2011). The program can also perform simple visualization operations, such as overlaying international borders, that produce monochromatic or multispectral image sets. It supports single-image synthesis and batch production (Space Science and Engineering Center 2017).

2.2 ERDAS Imagine

ERDAS Imagine version 16.0 is a Windows-based raster analysis program designed to interrogate georeferenced satellite and radar imagery. It provides radiometric, statistical, and pixel cataloging tools. The software also has image manipulation and transformation functions. It contains a library of spectral reference data and environmental characterization toolkits. The software’s workflow and script building functions can also be used to create advanced image processing models and batch edit imagery datasets. It can ingest raw data directly from multiple platforms and convert it into standard gridded image formats. The software also hosts its own console where commands can be executed via command line using the Imagine pseudo-language. More information on Imagine can be found in Hexagon Geospatial (2017).

2.3 ArcMap

ArcMap 10.5 is a Windows-based GIS designed to analyze and manipulate georeferenced data with a focus on two-dimensional visualization. It can be used to interrogate vector-based and raster-based files either separately or in conjunction with one another. It provides algebraic, statistical, and spatial analytics tools. It also supports the use of Python-based scripts. The program allows for automated workflows that can be built either in the program's command line interface with Python or using a tool called ModelBuilder.

ModelBuilder enables construction of complex data processing models through a GUI. It also supports the use of python scripts within models and model nesting. These models can be saved as Modelbuilder files or converted into Python scripts. The program's Python support also allows it to be run through an external Python terminal (ESRI 2007b).

2.4 ArcScene

ArcScene 10.5 is a Windows-based GIS designed to analyze and manipulate georeferenced data with a focus on three-dimensional visualization. It boasts the same analytical capabilities and toolkits as ArcMap, but with additional parameters to account for the additional dimension. The program allows the user to adjust solar intensity, angle, and azimuth to create realistic scenes as static images or fly-through animations. Like ArcMap, it supports Python scripts and hosts the ModelBuilder toolkit (ESRI 2007b). It is used for stereographic analyses of synthetic background images in this paper.

C. Case Study Locations

Seven sites are chosen to evaluate the performance of the ACAM developed for this study. They are selected based on their latitude, terrain features, land cover differences, and seasonal characteristics. Additional factors include human infrastructure and satellite overflight frequency. All of the sites have at least one major city within the sensor’s FOV. The sites fall into one of three categories based on satellite platforms: VIIRS-only sites, MSI-only sites, and dual-instrument sites.

VIIRS imagery was examined over one-month periods spanning the entirety of the instrument’s six-year data archive for a given month. Each site is examined for at least one month, but some sites are examined for multiple months. The MSI imagery is only examined over two-month periods due to small sample sets from the instrument’s two-year data archive. These months are chosen to minimize changes in scene radiance due to seasonal differences in sun angle. Datasets are composed of imagery observed between May and July from 2016 to 2017. A map of all sites is displayed in Figure 15.



Figure 15: Map depicting all sites examined during this study. Stars mark each site. Their surrounding states or nations are highlighted in orange (Glickman 2000).

1. VIIRS Sites

1.1 Fairbanks, Alaska

The region around Fairbanks, Alaska is selected for a one-month examination. It is chosen for its high latitude, mountainous terrain, and snowy landscapes. The region is examined during the month of April to ensure as much surface snow as possible and enough light to distinguish surface features throughout the scene. Terrain-corrected component files are not available for this location. Imagery from the mountainous region appears distorted as a result. See Figure 16 for an example image of this site.




Figure 16: NIR VIIRS image of Fairbanks, AK on a relatively clear day. The town is marked by a green dot. An inset map of Alaska is shown in the lower right corner of the image.

1.2 Miami, Florida

The region surrounding Miami, Florida is chosen due to its low extra-tropical latitude, long coastline, and large anthropogenic presence. Unlike the other sites, Florida is examined using the VIIRS DNB. It is evaluated during moonlit nights between the months of August and October from 2012 to 2017. Images are selected around dates of full moons. Data is discarded if clouds cannot be manually distinguished in a dataset due to low lunar illumination. While the focus of the location is on the area around Miami, the dataset provides sufficient cloud-free views of the landscape to produce an analytical domain spanning the entire state of Florida. The examined domain is expanded to the state borders. An example image of this site is displayed in Figure 17.




Figure 17: NIR VIIRS image of Florida on September 2, 2015. Major cities are marked with yellow dots. State borders are outlined with green.

1.3 Pyongyang, North Korea

Pyongyang, North Korea is chosen due to the site's midlatitude location and climate, as well as the surrounding country's diverse land cover types and terrain features. It is evaluated during the month of May due to the lower frequency of completely cloudy images relative to other months examined. The May dataset provides sufficient cloud-free views of the landscape to expand the analytical domain to North Korea's borders. An example image of this site is displayed in Figure 18.




Figure 18: NIR VIIRS image of North Korea May 18, 2016. Major cities are marked with yellow dots. North Korea's border is outlined in red. Other national borders are green.

1.4 South Bend, Indiana

South Bend, Indiana is chosen due to its location in the midlatitudes and seasonal weather differences. It is evaluated during the months of February and July from 2012 to 2016 to examine the ACAM's performance during summer and winter. The dataset provides sufficient cloud-free views of the landscape to expand the analytical domain to the state's borders for both seasons. An example image is displayed in Figure 19.




Figure 19: NIR VIIRS image of Indiana on July 7, 2013. Major cities are marked with yellow dots. State borders are outline with green.

2. MSI-Only Site: Ambouli, Djibouti

Ambouli, Djibouti is selected for its tropical latitude and climate. The city includes a small peninsula that adds to the diversity of the area's landcover. It is evaluated during the months of June and July from 2016 to 2017 using the MSI's high-resolution B8a channel. An example image of this site is displayed in Figure 20.




Figure 20: Synthetic NIR MSI background image of Ambouli, Djibouti. The city is marked with a green dot. An inset image of the country's borders is shown in the lower right of the figure. The country's border is outlined in yellow. Other national borders are green.

3. Dual Instrument Sites

3.1 Bagram, Afghanistan

Bagram, Afghanistan is chosen for its relatively low latitude, mountainous terrain, and diverse land cover types. The dataset provides sufficient cloud-free views of the landscape to expand the analytical domain for this paper to the nation's borders. The country is evaluated during May from 2012 to 2017 with the VIIRS radiometer. The area near Bagram is evaluated between May and June from 2016 to 2017 with the MSI radiometer. An example image of this site is displayed in Figure 21.




Figure 21: NIR VIIRS image of Afghanistan on May 5, 2015. Bagram is marked with a green dot. Afghanistan is outlined in red. Other national borders are yellow.

3.2 Los Angeles, California

Los Angeles, California is chosen for its minimal cloud cover and extensive civil infrastructure. It is examined for the months of May, June, and July from 2012 to 2017 using VIIRS data. The dataset provides sufficient cloud-free views of the landscape to expand the analytical domain to the state's borders for low-resolution imagery. The area near Los Angeles is examined during May and June from 2016 to 2017 using MSI data. An example image of this site is displayed in Figure 22.



Figure 22: NIR MSI image of Los Angeles, CA on June 18, 2016. The city center is marked with a green dot. The coast is outlined in yellow.

D. Background Image Synthesis

Two images are required to run the ACAM developed for this paper. The first is the image detected by the radiometer. The second is the synthetic background image used during change detection to determine where clouds have developed. The synthetic background image is created by determining an average value for cloud-free pixels in a scene over some period of time Zhu et al. (2015). For this paper, the synthesis of imagery representing one-month or two-month periods was attempted. The synthetic image approximates a natural scene, but reduces the appearance of short-lived optical absorbers in the atmosphere. Averaging pixel values over a one-month period also reduces the impact of changes to scene radiance due to varying solar angles, wet surfaces, and other short-lived surface phenomena.

1. Data Cleanup

Images synthesized using POLAR2GRID require quality checks for missing or unreadable pixels from the component files. Once the GEOTIFFs have been synthesized, they can be loaded into ArcMap for checking. Each file is manually inspected for missing data (null data) pixels and other image quality issues. If large swaths of pixels are missing due to dropped scan lines or other satellite hardware issues that degrade image quality, the file is removed from the dataset. Individual missing pixels can be filled in manually using the average radiance of surrounding pixels and do not require the image to be deleted. Once all files have been inspected manually, they can be fed into a model that fills in missing pixels.

The model uses Raster Calculator, an ArcMap tool that allows math to be performed using georeferenced raster files in a manner analogous to matrix algebra. Brightness values for each pixel can be manipulated using mathematical functions on single or across multiple images. The calculator identifies areas of each raster

where pixels exhibit null values and identifies the closest five pixels in every direction surrounding the missing pixel. The mean value of the closest five surrounding pixels is used to fill in the missing value. This is repeated for every missing value across each image. Because areas of missing data can vary in size, the Raster Calculator process is performed twice to ensure the final image is entirely free of missing or unreadable pixels. Images identified as having missing or unreadable pixels extending over large areas should be removed from the dataset, as the repaired pixel values may not faithfully represent a given part of the scene.

The model also truncates pixels in the image that fall more than five degrees latitude or longitude outside of the user-defined AOI. Next, the repaired and cropped images are saved as GEOTIFFs for further analysis. If no repairs are necessary for an image, it is simply cropped by the model before being saved to a new location. A diagram of the model is displayed in Figure 23. An untreated image is displayed in Figure 24. For imagery corrections beyond isolated clusters of null value pixels, such as dropped scan lines or errors due to satellite navigation problems, Imagines Radiometric Toolkit was used for repairs on an individual basis (Campbell and Wynne 2011). Imagery that could not be corrected due to the nature or scale of the bad pixels were deleted.




Figure 23: Diagram of the model used to repair images with missing pixel data. Circles represent input and output files. The color blue indicates a file or directory variable, while light or dark green indicates an output value or file. Arrows with solid lines indicate the direction of the model. Arrows with dotted lines indicate the file is a precondition for an action that follows. Yellow rectangles indicate the execution of a process. Hexagons indicate do-loops. Any shape with a bold P over it is a parameter that can be defined in the graphical user interface for the model.




Figure 24: Original VIIRS image of North Korea on 25 May 2017 at 0407 UTC.

2. First Guess Image

Once all files have been processed by the ArcMap image repair model, each image is manually inspected. If the majority of an image is cloud-free, the image

is moved to a directory for mostly clear images. Those identified as mostly cloudy are moved to a directory for cloudy images. This part of the technique is subjective, but unintended user influence on results from this technique is minimized because any bright pixels from clouds that might exist in a nearly cloud-free image will be minimized when averaged with pixels from other images. After all the images have been sorted into their respective directories, the first synthetic image needed for the authors technique is generated.

This first synthetic image serves as a foundation for the final image to be built upon and will be used as a first guess for what cloud-free pixels should resemble. This image is created by averaging together all cells sharing a common location across all images in the cloud-free directory using the ArcGIS Raster Calculator. Averaging together these pixel stacks eliminates minor daily variations in radiance due to scene changes and atmospheric effects. This uses the following equation:

$$S_1 = Con((I_1 \& I_2 \& \dots \& I_n), (Sum(I_1, I_2, \dots, I_n) / n), null), \quad (5)$$

where S_1 is the resulting synthetic image. The equation uses conditional logic, *Con*, to determine which stacks of cells will be averaged together. Co-located pixels from images 1 to n are selected by $(I_1 \& I_2 \& \dots \& I_n)$. The values from each stack of co-located cells are added together and divided by the total number of images, n. Stacks of pixels that do not have members in all images are given null values in the final image to ensure the final product contains only cells located in each of the original images. The resulting synthetic image is saved for later use.

3. Background Image Synthesis Generation

3.1 Manual Background Synthesis Method

The next step in creating a synthetic background image requires editing each image from the sample set in Imagine. This is the first of two methods for generating

synthetic background imagery. It is most suitable for datasets consisting of 30 images or less, like those originating from the MSI in this paper. The examined images are referred to as parent images. Imagines AOI creation tool is used to draw polygons around all land-based cloud-free areas for each parent image. For North Korea, optically thin cirrus is excluded from these polygons whenever it was prominent enough to be distinguished from the background. Low-level haze due to moisture or pollution is ignored during the creation of each polygon. Cloud cover that occurs over bodies of water is dealt with in a separate part of this technique. It does not have to be removed from the AOI polygons as a result.

After all cloud-free areas for a given image have been enclosed in AOI polygons, Imagines Create Subset Image tool is used to extract the pixels from each polygon-enclosed area of the image and output them into a new image file. Care should be taken to ensure each polygon excludes cloud shadows as well. Failing to do this can lead to overly-dark average pixel values in the final background. Areas of the raster outside the cloud-free polygons are assigned null values. Imagine preserves the coordinate data for each pixel in the resulting image so that georeferencing in ArcMap can be performed by the software automatically. The resulting file is then saved to a separate directory of cloud-free image snippets and loaded into ArcMap. An example of a cloud-free snippet is displayed in Figure 25.

Before the final synthetic raster can be generated, the series of cloud-free images created in the previous step must be merged with the first-guess synthetic raster. This simplifies the math involved in generating the final image and ensures pixels not encompassed by the cloud-free snippets are represented with reasonable radiance values. The ArcGIS Append tool is used to create a new image for each cloud-free snippet. The cloud-free cells replace co-located cells from the original synthetic image. The resulting image is saved to a separate directory. An example of output from this




Figure 25: Cloud-free image snippet of North Korea on 8 May 2013 at 0417 UTC. The national border is outlined in red.

step is displayed in Figure 26. Once all cloud-free snippets have been merged with the synthetic image to create new images, all of the resulting files are averaged together using Equation (5) with the Raster Calculator tool. The result is a synthetic image of North Korea where all areas of land are cloud-free (see Figure 27).

The next step in the process eliminates the reflections of light off rough and turbid areas of water. The focus of this step is on bodies of water where the surface texture varies regularly and the body of water is much larger than the MCAT. This process is not performed on small ponds and rivers. Manual inspections of observed




Figure 26: Cloud-free image snippet of North Korea on 8 May 2013 at 0407 UTC merged with the initial synthetic background image. The grayscale color ramp causes each member of the composite to stand out more than in the final synthetic image.

imagery indicate water levels and flow paths of these features vary too rapidly to simulate. For bodies of water large enough and static enough to warrant simulation, pixel values are generated by creating a raster image the size of the AOI with a constant value of 5, which simulates a radiance value of $5 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$, the unit of measurement used for radiance by both VIIRS and MSI (Cao et al. 2013; ESA 2015b). This value represents the mean radiance of seawater determined from small cloud-free areas of the sea near North Korea. Similar checks are done for bodies of




Figure 27: Cloud-free synthetic image of North Korea without water corrections.

water near the other VIIRS sites. Mean pixel values near $5 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$ are observed and used to simulate water in these background images as well. Seawater at night is sampled around Florida. A simulation value of $0.6 \times 10^{-9} \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$ is chosen after manual inspection of seawater pixels. These measurements are taken using Imagine's Compute Pyramids and Statistics tool. Due to the small sample set and the high resolution of the MSI imagery, this step is not required for any of the MSI synthetic background images. A raster encompassing the same domain as the draft synthetic image is created with the value of $5 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$ using the ArcGIS Create Constant Raster tool. Next the ArcGIS Extract by Mask tool is used to create

an image of simulated cloud-free areas of water. The tool utilizes a shapefile outlining bodies of water in and around North Korea to define the areas of the constant value raster to extract. All other pixels are assigned null values.

The water body shapefile is comprised of two separate 10 *m* resolution shapefiles that are merged together. One shapefile is made up of polygons representing all large areas of water across the planet such as seas, while the other contains features such as rivers and lakes (Kelso et al. 2017). The shapefiles are merged using ArcMaps Union tool to form a new shapefile showing all water-land boundaries in high detail. The simulated water image for the vicinity of North Korea can be seen in Figure 28.




Figure 28: Simulated water raster for the vicinity of North Korea for May. The national border is outlined in red for spatial awareness. The land is colored white and water colored is black.

Once the simulated water raster has been completed, it is merged with the synthetic image of North Korea using ArcMaps Append tool. A visual inspection of

the synthetic image reveals shallow areas of water along the coastlines that is not be properly represented by the simulated seawater value. To compensate for this, ArcMaps Buffer tool is used in conjunction with the water body shapefile from the previous step to generate a 5 *km* buffer around all coastlines. This ensures no areas of shallow water are misrepresented along the country’s coast. The original shapefile and the resulting buffer shapefile are displayed in Figure 29. The buffer shapefile is used in conjunction with ArcMaps Append tool to replace areas of coastal water with synthetic radiance pixels from the second synthetic raster. The resulting product is the final cloud-free synthetic image of North Korea for May (see Figure 30).

Although areas of cloud cover are identified manually and excluded during the generation of the background image, a visual analysis of the final image is still performed to ensure the products integrity. The finished image can be used with the ACAM in this paper or as a standalone product as a chart of terrain features. If radiance for any given pixel varies too greatly across the sample set radiance values in the final background image may not represent a cloud-free environment. This will also cause portions of the resulting background image to be significantly brighter or darker than their surroundings, such as in Figure 26.

The composite image resembles clothing mended with patches. If the technique used in this report is not successful in avoiding cloudy pixels, the final synthetic image will have a patchy appearance. Patchy areas can also indicate that the incoming radiation from the sun over the sample period has varied significantly throughout the parent images, requiring the final background image to cover a smaller period of time. Pixel variations when viewed across the scene should form natural looking terrain features that line up with DEM data.

The effects of optically thin cirrus are absent from the scene, leaving instead a crisper image of North Korea that closely resembles cloud-free portions of its parent




Figure 29: Coastline (red) and 5 *km* coastline buffer (yellow) shapefiles overlaid atop the synthetic water image.

images. Pixels that were identified retrospectively as haze are also noticeably absent from the image. Topographical variations stand out more clearly, with peaks and valleys exhibiting greater contrast compared to their surroundings than in parent images. Additionally, small lakes and rivers that were not previously discernible or lacked clear boundaries can now be identified without radiometric enhancement. Shallow areas of water along North Korea's coasts are left intact. Small islands off the coast of North Korea can also be clearly distinguished. An inspection of ridge lines in the final image revealed dendritic patterns are present as well. Given the number of parent images used, it makes sense that seasonally present snow cover remains pronounced over time while snow that melts soon after the scenes radiance values are collected is not present in the final image.




Figure 30: Final cloud-free Synthetic image of North Korea for the month of May. The national border is outlined in red. Notice the increased contrast in land formations compared to Figure 27.

The process also diminishes any effects of short-term or annual variations in water levels for the nation's lakes and rivers. These changes in water level and dendritic patterns could cause issues in change detection software if not properly addressed in within a detection algorithm. For features that cause problems, new polygons can be drawn around them in a new shapefile and simulated radiance values can then be added to the scene using the same approach applied to simulate the radiance of deep water. Seasonal sea ice can also be handled in this fashion, but due to the movement of ice, the polygons could not be based on seasonal values. Rather all sea ice has to

be removed from the synthetic image. Then, their positions have to be determined from an observed image used for change detection and added to the synthetic image for analysis of that same observed image. These corrections should be investigated in future research and are discussed further in Chapter V.

When selecting cloud-free areas in parent images, cloud cover outside the area of interest is occasionally included in the AOI polygons. The influence of cloudy areas from the parent images can be seen in a few locations outside North Korea's borders. These areas shift rapidly from bright to dark shades, giving the appearance of rough terrain that is not present in parent images. This can be seen by comparing Figures 31 and 32.

The unphysical bumpy terrain highlights the importance of extracting only cloud-free areas from the parent images used for this technique. Pixels influenced by cloud cover should be considered contaminated and eliminated from the final background imagery by either adjusting pixel values manually or by replacing the pixels in question with those from the first-guess image. Additionally, the resulting synthetic terrain will appear rough if cloud shadows are not avoided during pixel extraction. Contamination of pixels due to cloud shadows will result in unphysical features resembling valleys. An example of this issue is displayed in Figure 32. Pixel values in the final background can also be influenced by features such as standing clouds and fog. These features could cause background pixels to be too bright. This issue is avoided over large areas of water where pixels are assigned simulated values.




Figure 31: An observed image of North Korea's northeast border.




Figure 32: The yellow circle highlights an area of non-physical bumpy terrain. Notice how peaks and valleys stand out with high contrast inside the country's border (left portion of image), but outside the country's border the terrain quickly becomes rough in a manner not consistent with the terrain in Figure 31.

3.2 Primary Automated Synthesis Method

The second method for generating synthetic background imagery developed for this paper relies on an ArcMap model. This method is intended for large datasets of 30 or more images. The user supplies the model with the first-guess background image and defines the directory containing the checked and corrected observed imagery from which the background image will be generated. To simplify the analysis process, a GUI was developed. The GUI is displayed in Figure 33 and diagram of the full model is shown in Figure 34. The model relies on a nested version of the ACAM developed for this paper. One of the ACAM's output files is used by the synthetic image generator. The synthetic image generator can be configured to generate the final synthetic background image in one model run, but due to hardware memory constraints must be broken up into two components. The model saves the cloud-free synthetic images from each observed image to an output directory for later use.

To run the model, the user defines the location of the input files, a background first-guess image, minimum cloud size in square meters (MCAT), scratch workspace, and output directory. The scratch workspace is the directory where ArcMap processes the files used by the model. For very large datasets using a scratch workspace speeds up processing speed by moving the current files to the computer's onboard storage. Using a scratch workspace does not alter the input files. If onboard storage is limited this setting can be safely set to the output directory as well. The cloud brightness threshold (CBT) and vertex filter number (VFN) must also be set by the user. The workflow symbology for the background image first-guess image and three directories can be in Figure 35. The CBT is the difference in radiance between a cloudy and cloud-free pixel for a given scene. For daytime VIIRS imagery, the brightness threshold is set to $20 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. For nighttime VIIRS imagery, the CBT is set to $2 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. MSI data is not processed with this model.




Figure 33: The synthetic image generator GUI.

These values are chosen by manually examining variations in pixel values between multiple types of surface features and clouds from imagery generated by subtracting the first guess image from an observed image. ArcMap's Identify tool was used to display pixel values across the differenced scene. The shadow filter compensates for pixel values much lower than the first-guess image's and will be discussed later. The daytime CBT is applied to all daylit VIIRS imagery evaluated in this paper. The nighttime CBT is only used for Florida's synthetic background image.

The VFN defines the number of vertices that must make up a polygon identified by the ACAM for it to be labeled as cloud cover. This filter prevents plots of land or other elements of the scene with less complex shapes valued at or above the CBT from being labeled as cloud cover. It is used by the ACAM for identifying clouds and




Figure 34: Diagram showing the workflow for the synthetic image generator. See Figure 23's caption for an explanation of the symbology.

helps to ensure all cloud features are eliminated from the output synthetic image. The MCAT is set to $640,000 \text{ m}^2$ for daytime imagery and $2,560,000 \text{ m}^2$ for nighttime imagery. These values are chosen because they are approximately twice the height and length of one pixel from each VIIRS image type. In conjunction with the VFN, the MCAT threshold ensures that clusters of pixels identified as cloud cover are large enough to be operationally significant clouds and exhibit the correct shape complexity. Future studies should investigate more specific settings for each threshold. The workflow symbology for the three threshold settings are displayed in Figure 36. The model is executed once all of the values in the GUI have been set.




Figure 35: Input and output directories and files for the synthetic image generator. See Figure 23's caption for an explanation of the symbology.




Figure 36: Input variables for the synthetic image generator. See Figure 23’s caption for an explanation of the symbology.

After execution, the model indexes the imagery inside the input directory. Each file is run through the model workflow individually. Both the full filename and shortened version of the filename are extracted as strings of text. The full versions of the filenames are used to edit each file in the workflow. The shortened versions are used to name final output. The observed image is ingested by the ACAM nested in the workflow, which uses threshold settings and the background image to determine cloud-free pixel clusters. It generates a file containing only cloud-free pixels. The ACAM will be discussed in detail in a subsequent section. This is displayed in Figure 37.

During preliminary model tests, the files produced by this part of the workflow contained missing pixels. This was overcome by adding a tool that replaces any missing pixels within the array in the output image with pixels from the first-guess image as a precaution. The tool consists of a raster calculator with a conditional statement that searches for empty pixels in the image. It replaces pixels flagged by the ACAM as cloud cover with first-guess cloud-free pixels. Missing pixel values would otherwise be preserved throughout the workflow and lead to gaps in the final image. Once checked and corrected, the cloud-free image is appended to a copy of




Figure 37: Initial steps of the synthetic image generator's workflow. See Figure 23's caption for an explanation of the symbology.

the first-guess image. A copy of this file is then created using ArcMap's Copy Raster tool (as displayed in Figure 38).

The appended image is then ingested by an ArcMap tool that filters out cloud shadow pixels. If a given pixel value from the synthetic image has a lower value than the first guess image, it is replaced with the first-guess pixel value using ArcMap's Raster Calculator tool with a conditional statement. This approach is chosen so that pixel values on the bright end of possible background values are favored. This helps minimize the difference between bright, shortlived non-cloud features such as snowfall and the final synthetic background image. The exception to this filter application is to Afghanistan's background image. Due to the high frequency and extent of cloud cover in the sample set, pixel values in the synthetic image are allowed to fall 25 percent lower than their first-guess counterparts' values before being replaced. Afghanistan's




Figure 38: Portion of synthetic image generator’s workflow during which values for missing pixels are determined and cloud free portions of input file are merged with first-guess image. See Figure 23’s caption for an explanation of the symbology.




Figure 39: The final steps of the synthetic image generator. The algorithm looks for large negative changes in pixel values and replaces the identified pixels with first-guess pixels. See Figure 23’s caption for an explanation of the symbology.

synthetic background is generated using the alternate automated method and will be discussed to greater depths in subsection 3.3 of this chapter. The resulting cloud-free synthetic image is saved to the output directory using a shortened version of the

original observed file's name (as displayed in Figure 39). The model run is complete after each input file has been used to generate a synthetic image.

After the run is finished, the mean value for each pixel is determined from all the synthetic images in the output directory with the Raster Calculator using Equation 5. For images completely obscured by cloud cover, all pixels are replaced by their first-guess values. This weights the synthetic background image toward cloud-free pixels for each day of complete cloud cover present in the sample set. Once completed, the resulting mean synthetic image requires the same treatments as images generated using the manual method. The remaining steps can be found in subsection 3.1 of this chapter. Each final image is manually compared to a cloud-free or relatively cloud-free observed image to check for accuracy. A qualitative verification technique is chosen because completely cloud-free images for a given month are not available for all locations.

Certain sites exhibit noticeable differences over large areas between the background and observed images. This is primarily due to constantly changing surface features, such as mountain dendric patterns, surface wetness, agricultural activity, and urban activity. These differences are plausible and acceptable for the background image. In areas with frequent cloud cover and rough terrain, stereographic terrain inspections are used as additional checks. This technique will be discussed in a subsection 3.4 of this chapter.

3.3 Alternate Automated Synthesis Method

It is found that the synthetic image produced for Afghanistan with the primary automated method does not compare well with observed imagery. It exhibits dark elliptical features that cover large areas. They resemble craters or lakes. Other features resemble rivers that run either over flat terrain or along mountain ridge lines. None

of these are found in the observed images. Examples of these features are displayed in Figure 40. A manual examination reveals the synthetic images used to generate this background image contained widespread and frequent landcover obscurations. The dark pixel clusters generally correspond to shadows from cloud cover in observed images. Valleys are found to also appear substantially darker in the background. These issues can be observed by comparing Figures 41 and 42.




Figure 40: Subset of the synthetic background image resulting from the primary generation method. The mountains and landscape are blurry. Bagram, Afghanistan can be seen in the upper left of the image.

Manual inspections are used in conjunction with DEM data and vector-format files containing locations of local water bodies to ensure background quality (ESRI 2007a; Patterson 2017). Waterbody polygons are overlaid atop the suspect image and assigned bright colors against the grayscale background (as displayed in Figure




Figure 41: An observed image of the region surrounding Bagram on May 1, 2012.




Figure 42: Subset of the final synthetic image using the alternate automated background generation method. Compare to Figures 41 and 40.

43). These checks indicate anomalous pixel values are present in the final synthetic background image. To overcome this issue, the synthetic background image is edited




Figure 43: Subset of the final synthetic image using the alternate automated background generation method with lakes highlighted in green (ESRI 2007a).

using ArcMap’s Focal Statistics tool. This tool analyzes all cell values within a given radius of each pixel. It applies a mathematical operation to the cell to assign it a value closer to the surrounding pixels. For Afghanistan, the tool is set to evaluate rectangular-shaped pixel clusters that extend five pixels in each direction around the pixel in question. The rectangle moves across the input image and applies the maximum detected pixel value from the cluster to each pixel at the rectangle’s center. The background image of the entire country is used as input. The results are displayed in Figure 44.

The output image from the tool eliminates most of the anomalously low pixel values from the scene (as displayed in Figure 45). It is used as the first guess image in the Synthetic Image Generator. The shadow filter’s threshold is set 25 percent below the CBT value. The resulting synthetic background image does not exhibit the widespread clusters of dark pixels. The final image is displayed in Figure 46.




Figure 44: The original synthetic background image for Afghanistan using the primary automated method.




Figure 45: The image resulting from the pixel smoothing tool. Increased snow cover extent can be seen over the mountains on the right side of the image.




Figure 46: Final synthetic background image for Afghanistan produced with the the alternate automated method.

3.4 Stereographic Terrain Analyses

Stereographic checks involve stretching a synthetic background image over a DEM using ArcScene. The topography of the DEM acts as truth data. In this paper, stretching refers to mapping each pixel in an image to a position on the DEM's three-dimensional topography. The checks rely on common knowledge of what natural formations should look like and observed images from each site. The appearance of lakes is one example. They are associated with planar surfaces. If lake-like features are observed along the slope of a large hill or mountainside, the representativeness of the synthetic image is called into question. Further comparisons to the same location with observed imagery are then made. If the feature is present in observed images, it is possibly not a lake and the background image is not altered. If it is not present in the observed imagery, the background image requires adjustment. Three-dimensional views of each image in question are created with ArcScene's Anaglyph

Display and Perspective view tools. Perspective views of the area surrounding Bagram are displayed in Figures 47, 48, and 49.

Once the issue with the imagery is recognized and a fix implemented, the final image is checked once more with the DEM (Patterson 2017). If the aforementioned check is completed and no issues are found, the image is ready for use with the ACAM. Stereographic terrain checks are best applied in areas with rough or mountainous topography. In addition to Afghanistan, these checks are used in Alaska and North Korea. This approach is not used in areas with relatively smooth terrain such as Indiana and Florida.




Figure 47: Three-dimensional view of Bagram generated by projecting the synthetic background onto the DEM's topography.




Figure 48: Original synthetic background image of Bagram. Large dark spots atop the mountains in the foreground interfere with nearly continuous snow cover.




Figure 49: Final synthetic background image for Bagram. Compare the evenly distributed Dendritic patterns over the mountains in the foreground to those seen in Figure 48.

E. Cloud Identification Workflow

Construction of the ACAM used in this paper is done using ArcMap’s ModelBuilder Toolkit. This approach allows models to be nested and duplicated quickly. The nesting capabilities of ModelBuilder also allows for the ACAM’s modular design. In section D of this chapter it was mentioned that the synthetic image generator relies on a nested version of the ACAM. In the same way that it became a part of a large workflow, the ACAM can be linked with variances of itself that have different settings or to a detection model that searches for non-meteorological features. After the ACAM has masked out cloud cover the task of identifying other scene elements is less complicated. New tools and functions can easily be added to the ACAM to transform or filter the output data via ModelBuilder. It can also incorporate python scripts to leverage resources from other software (ESRI 2007b).

1. Model Description

The ACAM developed for this paper is a features a workflow requiring minimal input aside from the observed and synthetic background images. It features a GUI that can analyze images individually and a batch analysis mode. The GUIs for the single image and batch analysis modes can be seen in Figures 50 and 51. A diagram of the entire model workflow can be seen in Figure 52. The GUI resembles the Synthetic Image Generator’s closely because most of its settings control the nested ACAM (see Figure 33).

To operate the model, the user opens the GUI and selects an observed image to be analyzed for cloud cover, a synthetic background image, and the output location of the analysis files. The CBT and VFN must also be set. Optionally, the user can edit the ACAM in ModelBuilder and save new default values for the GUI settings. The brightness scaling value is used to adjust the values of the observed image to




Figure 50: The ACAM's primary GUI.




Figure 51: GUI for batch-processing multiple images using the ACAM.

match the order of magnitude of the synthetic background image. This setting is primarily used to scale up the values of VIIRS DNB imagery. These images can contain floating point values smaller than 1×10^{-4} that must be scaled so that the smallest pixel value in each image corresponds to a digit, otherwise ArcMap's Raster




Figure 52: Overview of the ACAM's image processing workflow as in Figure 23.

Calculator will not run properly. According to ArcMap documentation, these values do not require scaling for algebraic functions to function, but testing demonstrates that with large images, such as those exceeding 100 MB, floating point values had to be rounded up or down to the nearest integer. Scaling up values much less than zero prevented this problem. ModelBuilder's edit mode also allows the user to adjust the number of cores used by the ACAM. The default value is set to four. The symbology for these settings in the model workflow are displayed in Figures 53-56.




Figure 53: Symbolic overview of input and output variables for the ACAM as in Figure 23.

The model is ready to be executed after all settings have been entered. The first step in the analysis creates a binary change detection image composed made up of differences between the observed and background images. It relies on a series of conditional statements. Figures 57 and 58 display examples of observed and synthetic images. The following equation is used by ArcMap's Raster Calculator to detect changes between the two files and generate the binary image:

$$I_M = Con(I_O \& I_S, Con(((Int(I_O) - Int(I_S)) \geq Int(B)), 1, 0), 0), \quad (6)$$

where I_M is the generated binary image, I_O is the observed image, I_S is the synthetic background image, and B is the brightness threshold. All pixel values are rounded to the nearest integer to conserve memory using the Int function. Conditional statements, corresponding to the $Con(X, Y, Z)$, are nested to reduce the number of tools needed in the model to create the binary image. With conditional statements, if some input X is true then Y is output, otherwise Z is output for each pixel value. The outermost conditional statement examines all pixels across both images. When pixels are present for the same location in both images, the inner conditional statement is

executed. Otherwise, the pixel is assigned a value of zero in the output raster. If two co-located pixels are detected, the synthetic pixel value I_S is subtracted from the observed pixel value I_O . When the difference between the two pixels exceeds the brightness threshold value B , the corresponding pixel in the output image is assigned a value of one. Otherwise, the output pixel is assigned a value of zero. Figure 54 displays this part of the workflow.

The resulting binary image is analyzed by ArcMap's Raster to Polygon Tool. The tool converts clusters of like-valued pixels into polygons and saves them to a shapefile. These polygons are assigned values of zero or one that correspond to their original pixel values. The perimeters of each polygon correspond to boundaries between clusters of pixels valued zero and one. ArcMap's Add Geometry Attributes Tool analyzes the shapefile output during the previous step. It determines the geodesic area of each polygon, which allows the surface area to be calculated regardless of map projection type for gridded imagery. This portion of the workflow is displayed in Figure 55.

ArcMap's Select Tool is used to find polygons that meet the MCAT threshold and the minimum number of allowed vertices for shapes coded with the value one. ArcMap's Extract by Mask Tool uses the selected polygons to extract clusters of pixels corresponding to areas of detected cloud cover. The Parse Path Tool is used to store the original filename as a string. Lastly, the image of extracted cloud cover is saved under a name derived from the original file. The final portion of the workflow is displayed in Figure 56. The areas of cloud cover detected by the ACAM using Figures 57 and 58 are displayed in Figure 59.




Figure 54: Symbolic overview of the initial raster calculator that conducts change detection math using both input images as in Figure 23.




Figure 55: Symbolic overview of tools that convert the binary mask into polygons as in Figure 23.




Figure 56: Selected polygons are used to extract cloud free portions of the original image that is output as a mask. The algorithm can be changed to output either a binary mask or mask retaining the original pixel values of all clouds from the input scene. Symbology as in Figure 23.




Figure 57: Observed VIIRS image over California from May 26, 2017 at 2042 UTC.




Figure 58: May synthetic background image for California.




Figure 59: Cloud cover detected over California on May 26, 2017 at 2042 UTC. The cloud mask is translucent and yellow.

2. VIIRS Settings

One goal of this paper is to build an ACAM capable of detecting cloud cover with few or no changes to cloud detection thresholds. To accomplish this ACAM settings are divided into three groups: VIIRS daytime, VIIRS nighttime, and MSI daytime. Except in special situations these settings are applied to all samples respective to their platform and channel. The settings that will be used for each category are based on experiences with early ACAM experimentation by the author. These settings are to compare the ACAM's detection capabilities to standard platform cloud masks in Chapter IV. The settings used for this paper are displayed in Table 5.

Table 5: ACAM settings by instrument and time of day.

Variable Name	VIIRS Day	VIIRS Night	MSI Day
MCAT (m^2)	640,000	2,560,000	2,600
CBT	45	5	1000
VFN	10	10	25
Brightness Scalar	1	10^9	1

2.1 Daytime Settings

The three settings required for cloud detection in daylight imagery in the GUI are the MCAT), CBT, and VFN. For this image category, the cloud brightness threshold is set to $45 \text{ W m}^{-2} \text{ sr}^{-1} \mu\text{m}^{-1}$. The GUI interface is not platform-specific, so the units assigned to pixel radiance or brightness is sensor-specific. The MCAT is set to 640,000 m^2 . This corresponds to the area covered by a cluster of four pixels from the VIIRS NIR channel. Plots of land or other surface features small enough to only span one or two pixels could be bright enough to meet the CBT. The MCAT minimizes the chances these features will be flagged as cloud cover. The threshold value is determined by doubling the horizontal length and width of the channel's original spatial resolution.

The VFN is set to ten vertices to complement the number of pixels used to set the MCAT. The least complex shape that four rectangular pixels can create is a rectangle. ArcMap determines the number of vertices comprising each polygon by starting arbitrarily choosing one and counting each subsequent vertex. The final vertex counted is the first vertex. This method ensures each shape is closed so that the software does not incorrectly classify a multi-point line as a shape. Each shape will always have a vertex count one unit higher than the number of vertices it has. For example, triangles have four vertices and squares have five vertices in ArcMap (ESRI 2007b). The brightness scaling value is left at the default value of one to prevent scaling of observed pixel values.

2.2 Nighttime Settings

Imagery analyzed from the VIIRS DNB requires an additional tool be added to the ACAM and an additional setting in the GUI. The additional tool reflects the necessity of masking out pixels corresponding to bright city lights in the analysis workflow. City lights exhibit varying intensities throughout the observed images. The mask of city lights is generated by manually examining the background image and using conditional statements with Raster Calculator to assign values of zero to pixels valued less than those corresponding to the lowest city light value in the background image. Manual analysis indicates the lowest values associated with lights occur at approximately $250 \times 10^{-9} \text{ Wm}^{-2} \text{ sr}^{-1} \mu\text{m}^{-1}$ for the VIIRS DNB channel. The resulting image is converted to a binary mask, where one corresponded to city light pixels and zero corresponded to all other pixels. The city lights mask is displayed in Figure 60.

The mask ensures the majority of city lights are not flagged as cloud cover by the ACAM. The model uses this image in conjunction with Raster Calculator to mask out areas exceeding the CBT in the binary image output after the ACAM compares




Figure 60: Subset of the binary mask indicating surface lights generated for Florida and the surrounding region.

the observed image to the background image. The rest of the model workflow runs as discussed in section D of this chapter. A city lights image is required to run the ACAM for nighttime imagery. Evaluation of areas with no surface lights can be accomplished by supplying the ACAM an image of all zeros as the city lights mask. Next, the MCAT is set to $2,560,000 \text{ m}^2$. This reflects twice the height and width of an 800m by 800m pixel, which approximates the DNB's 5.7×10^{-3} degree resolution. The CBT is set to $5 \times 10^{-9} \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. This setting is determined by manual analysis of multiple images from the sample set as in subsection 2.1 of this section. The VFN is set to ten. The brightness scalar is set to 10^9 . This setting is determined manually after inspecting the minimum values of observed images. It scales the brightest pixel values in night scenes into the hundreds or thousands to

ease interpretation and reduces potential issues with Raster Calculator. This setting can be left at the default value of one if pixel values in an observed imagery have already been scaled. Because missing pixels were repaired for the entire sample set, scaling of each nighttime image was accomplished during that step as well. The DNB images analyzed from this point forward in the paper are scaled by the software but unscaled radiance values will be given.

3. MSI Settings

Standard ACAM settings for daytime MSI imagery differ from VIIRS settings due to the instruments's high resolution. For change detection to work at this level of detail, delineating daily changes due to non-cloud objects, such as anthropogenic activity and atmospheric moisture puffs, requires the MCAT be set to a value that excludes these details while ensuring the ACAM can identify operationally significant small clouds. The value of 2600 m^2 is chosen. This approximates the size of an American football field. Other than agricultural activity, changes to the background scene at this scale that are bright enough to meet the CBT should be due to the development of cloud cover.

This MCAT value also reduces the influence of construction in cityscapes where structures have highly reflective roofs. The spatial extent of newly built structures in a scene should be minimal so long as that month's background image is updated annually. When new buildings are in close proximity to others with bright features, the older features should be well-represented in the synthetic background image and not be flagged as cloud cover. This prevents the polygon associated with new, bright features from approaching the MCAT. The CBT is set to $1000 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. This value is chosen based on multiple manual analyses of observed imagery. It approximates the lowest pixel values associated with cumulus and stratus clouds

when an observed image is subtracted from a synthetic background image. The VFN is set to 25 based on findings from manual analyses of the intermediate shapefile produced by the ACAM.

F. Qualitative and Quantitative Evaluations of Imagery

Success of the ACAM developed for this paper depends on synthetic background quality. For a synthetic background image to be useful it must not differ significantly from cloud-free images of the same location. Manual analyses of these images can be accomplished quickly, but are subjective and difficult to quantify. Statistical tests are used in addition to manual inspections to evaluate ACAM performance. Manual interpretation can also be used to create binary cloud masks for validation purposes. Manually generated cloud masks are often used to establish truth data for measuring the accuracy of algorithmically constructed masks. The time required to produce each mask constrains the volume available for validation (Hutchison et al. 2014). Operational cloud masks from both radiometers are used as truth for masks produced by this paper’s ACAM. Due to limited availability of operational cloud masks from VIIRS prior to July 2017, synthetic truth data is generated to supplement the locations available for validation in the dataset (NOAA 2017c). Synthetic truth data will be defined as masks produced manually with the aid of automated processes to accelerate production. Operational cloud masks are chosen for validation based on cloud type and extent. Imagery with minimal cloud cover is avoided where possible, but a range of cloud cover is evaluated to check for patterns of false positive or false negative pixels in the ACAM masks. Imagery with various cloud types is favored over imagery with homogeneous cloud cover.

1. Truth Data

Truth data for this study is obtained from the CLASS data archive or produced manually (NOAA 2017c). Manually generated masks are referred to as synthetic truth data. Operational truth data refers to the cloud masks produced for operational use by each radiometer.

1.1 Operational Truth Data

Recall, operational truth data is comprised of NETCDF-4 formatted VIIRS Cloud Mask Environmental Data Records (NOAA 2017c). Binary cloud mask layers are loaded into ArcMap as layers of georeferenced points and are converted to GEO-TIFFs with resolutions matching that of the NIR channel being evaluated. ArcMap's Point to Raster Tool is used to perform conversions. Gaps between pixels during the conversion are eliminated using the most frequently encountered pixel value around each gap to minimize value bias. The resulting binary mask image can be compared to the ACAM cloud masks.

1.2 Synthetic Truth Data

Synthetic truth data is generated by employing an ERDAS Imagine tool that allows pixel values to be recoded. Recoding is the process of overwriting pixel values. The Unsupervised Classification Tool automatically determines bins to which like-valued pixels are assigned. The tool does not alter pixel values until the image is recoded after the tool's final iteration. The K-means method is used for this paper. It generates a large, arbitrary quantity of bins that cover the input image's entire range of pixel values. The tool determines center points for each bin. It calculates the distance between each the center point. Next, it combines the set of bins separated by the smallest center point distance. The center point for the new bin is calculated

before the next iteration begins. The algorithm recodes the image once it reaches the user-specified number of bins.

Bin values used for this paper range from 10 to 36 $Wm^{-2}sr^{-1}\mu m^{-1}$. The highest values are associated with imagery containing snow or cirrus clouds. Recoded image results are manually inspected. If the highest-valued bin for a recoded image contains scene elements other than cloud cover, the image is run through the recoding tool again with a higher bin value. Images loaded into ArcMap are converted into binary masks using Raster Calculator. Bin values associated with cloud-cover are assigned values of one. All other bins are assigned values of zero. This process results in binary masks exhibiting the source image’s resolution, but can also result in pixels along the edge of cloud formations being misclassified. Misclassification of optically thin clouds or small regions of bright, non-cloud features is avoidable, but cannot be eliminated completely. Limited presence of these elements is deemed acceptable in synthetic truth data because they are analogous to shortcomings observed in the operational cloud masks also used as truth in this paper. The resulting cloud mask is used for sites where truth data is limited or unavailable.

2. Chi-square Goodness-of-Fit Test

Synthetic background images developed during this study are generated by averaging together a series of co-located cloud-free pixels from observed imagery for a given AOI. Pixel values in each synthetic image correspond to the mean value of all cloud-free co-located pixels from the entire sample set for a given time period. Minor changes in clear sky radiance are observed between certain images in each AOI as well. This is attributed to changes in the sun-synchronous orbit of each satellite and sun angle over time (Kidder and Vonder Haar 1995; Campbell and Wynne 2011). Images are selected at approximately the same time of day to minimize these types

of radiance variations. Averaging pixel values mitigates the influence of these variations on background images and also ensures that short-lived, large changes in pixel radiance do not influence background images. The synthetic image is not intended to be an exact match of each observed image. It only needs to be an approximation. Background images are deemed acceptable if they do not differ significantly from an image observed captured during the associated season or month.

The chi-square goodness of fit test is used to evaluate each synthetic background image (McGrew et al. 2014). The test quantifies differences between two histograms. These differences determine whether the author will reject or fail to reject the null hypothesis for this paper. The null hypothesis states that there is no significant difference between the observed and synthetic images. The alternate hypothesis is that the two images are significantly different. The test relies on a statistical value that designates a position along the horizontal axis of an approximately Gaussian distribution. The process given in this paper is for the one-tail version of the chi-square test only, so only the rightmost tail of the curve is used to determine the outcome of the test. This end of the curve is also referred to as the upper tail. It is displayed in Figure 61. If the test statistic is positioned too close to a specified location near the right tail of the curve, the null hypothesis is rejected and the alternate hypothesis fails to be rejected. This threshold is defined as the critical value. If the test statistic exceeds this value the null hypothesis is rejected, and the synthetic image in question is not acceptable for use in the study. The chi-square test uses the following formula

$$\chi^2 = \sum_{i=1}^k \frac{(O_i - E_i)^2}{E_i}, \quad (7)$$

where O_i is the observed radiance value for a given pixel, E_i is the synthetic pixel's radiance value, and k is the number of pixels evaluated. Small differences in pixel values will result in a small value for χ^2 , but any large difference between pixels results

in a large value for χ^2 . A ModelBuilder workflow is created to automate the majority of this test. An overview of the model can be seen in Figure 62.




Figure 61: Curve used to assess results of the chi-square test. The critical value is displayed on the horizontal axis.




Figure 62: The workflow used to execute the chi-square goodness-of-fit test with symbology as in Figure 23.

The chi-square workflow relies on a Raster Calculator equation that squares the difference of each pixel between the two compared images and divides the resulting quantity by the synthetic pixel value. Next, the filenames of each image are extracted and stored as strings to name final output. This part of the model is displayed in Figure 63. A pixel-by-pixel comparison of each image is not feasible due to the large spatial domains the and quantity of short-lived features inherent to the observed images. Rather than mask out all non-representative features, the model randomly selects points within the spatial domain of each analysis. Next, it extracts pixel values from the intermediate chi-square image created by the workflow. This portion of the model is displayed in Figures 64 and 65. The extracted points are saved to a shapefile named after both input images. Points that occur over pixels associated with cloud cover, snow, or anthropogenic activities (e.g., construction) are deemed invalid for the test and removed from the shapefile. The remaining samples are used to determine the test statistic.




Figure 63: Input files and initial steps of the chi-square test workflow with symbology as in Figure 23.

The final results of the test are displayed in ArcMap in the shapefile's attribute table. The column labeled "RASTERVALU" contains values that must be added together to determine the test statistic. The χ^2 value is obtained by right-clicking the column header. A window with the sum of the entire column is displayed. The




Figure 64: Input variables and random point generating tool used to extract data from sample points for the chi-square test with symbology as in Figure 23.

degrees of freedom used for this test is determined by subtracting one from the number of sample points. Test results can be obtained by comparing the test statistic and degrees of freedom value to a chi-square lookup table (National Institute of Standards and Technology and Semiconductor Manufacturing Technology 2013). For an image to be useful, there must be at least 97.5 percent certainty that the two images do not differ significantly. This percentage is associated with a significance level of 0.025. If the test statistic determined from the comparison is less than the value given by the lookup table, there is not enough evidence to reject the null hypothesis. The synthetic image is determined to not differ significantly from the observed image. The smaller the test statistic is in a one-tail chi-square test, the better the goodness-of-fit is between two images. Test results for each background image will be given in Chapter IV.

3. Change Detection Test

Cloud masks are compared using a simple change detection method. Truth data masks delineate cloudy from non-cloudy pixels with either ones or zeros, respectively. The masks generated by this study's ACAM are recoded to denote cloud and non-




Figure 65: Pixel values are extracted from the raster calculator's output at random points across the domain. The extracted values are saved to a shapefile along with the associated random points with symbology as in Figure 23.

cloud pixels with values of two or zero, respectively. Adding the two cloud masks together using Raster Calculator results in an image with pixels values that correspond to one of four possible categories. Pixels with values of three correspond to points where both masks detect cloud cover. Pixels valued at one or two correspond to points where only one mask identifies cloud cover. Pixels valued as zero correspond to areas identified as cloud-free by both masks. Each image is checked manually against the respective observed image to assure the integrity of each comparison. A basic subjective analysis is also performed. Sources of error are highlighted along with the strengths and weaknesses observed with each input mask.

4. Contingency Tables

Contingency tables are one method for examining error in each change detection image (McGrew et al. 2014). These tables display differences between pixel categories using numerical values. In this paper, the proportion of pixels found in each ordinal category is compared to the total quantity of pixels per scene used for each mask

comparison. Differences are expressed as percentages and are determined by the following equation

$$P_{Cat} = \frac{x}{n} \times 100, \quad (8)$$

where P_{Cat} is the percent of total pixels in each category, x is the ordinal category value, and n is the total number of pixels in the image. Prior to these calculations, the amounts of clear and cloudy pixels detected in each scene are inspected to determine if either cloudy or cloud-free pixel categories required weighting. Weighting increases the influence of one or more element in a calculation to reduce the influence of bias on the resulting value (McGrew et al. 2014). The percentage of cloud-free pixels in each image is compared to the total pixels from each truth mask. Cloud-free pixels range from 6.3 percent (mostly cloudy) to 95.6 percent (mostly clear) across the evaluated images. The mean percentage of cloud-free pixels per scene is 61.5 percent with a standard deviation of 25.5 percent. This indicates the dataset contains on average more cloud-free than cloudy pixels, but the spread of values indicated by the standard deviation indicates that weighting any of the categories would exaggerate the ACAM's performance. The decision is made not to weight the data. This approach to validation is different than what was used by Zhu et al. (2015), because synthetic imagery in that paper was tested with relatively homogeneous images. The examined AOIs were also much smaller than most of the images examined in this study. Their background images were used to determine land cover and ecological information about a scene rather than find cloud cover. They obtained their results by minimizing the presence features such as snow. Synthetic images generated for this study lead to better cloud detection when they are able to capture features that present themselves on a monthly or seasonal basis. This allows the ACAM to better detect the changes in scene radiance that reveal cloud cover.

Manual checks of operational cloud masks reveal that not all pixels classified as cloud cover correspond to clouds in observed imagery. Cloud masks generated by this study's ACAM are deemed acceptable if at least 80 percent of the pixels composing a scene are mutually identified by both cloud masks. This value is chosen because the ACAM masks are being compared to imperfect truth data. The performance standard for VIIRS operational cloud masks over any given land cover type is at least 90 percent during the day and 85 percent at night (Hutchison et al. 2014). The accuracy of MSI masks is at least 80 percent for a given land cover type according to Mueller-Wilm et al. (2016). Differences in pixel classifications between the two masks do not always represent misclassified areas of cloud cover as a result. This drives the decision to use both quantitative and qualitative evaluations of mask comparisons for gauging the ACAM's overall performance.

IV. Results and Analysis

A. Analysis of Synthetic Background Imagery

Synthetic backgrounds are synthesized for each location using the maximum number of images available for the location of interest. Only images chosen for ACAM validation are excluded from background synthesis. Each synthetic background image is evaluated quantitatively and qualitatively. Qualitative evaluations highlight any strengths or weaknesses difficult to quantify. Statistical tests have final authority on whether or not background images are acceptable for use with the ACAM. Differences between the pixel values of both images are depicted as images as necessary. Scatter plots of differences between the two images at the sample points used for the chi-square test are also given. It is desirable for plotted points to be less than the threshold used for change detection for each instrument platform. These values are 45, 5×10^{-9} , and $1000 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$ for VIIRS day, VIIRS night, and MSI images, respectively. Pixels that exceed the CBT threshold are permissible in the background image if isolated. Clusters of bright pixels must still meet the MCAT requirement before being flagged as cloud cover. While false positives are possible, the MCAT acts as a filter that minimizes false positives in the final mask. Chi-square test results are displayed in Table 6. The total number of parent images used to create each background image is displayed in Table 7.

Table 6: Chi-Square results for all synthetic background images by site, month, instrument (INSTR), sample point number (PNTS), and degrees of freedom (DF). The chi-square test statistic (CHI-SQ) is given, along with the associated critical value (CRT VAL). The table also displays whether background images differed significantly from the observed image used in each test (SG DIF).

SITE	MONTH	INSTR	PNTS	DF	CHI-SQ	CRT VAL	SG DIF
AFG	MAY-JUN	MSI	150	149	2.6	118.0	NO
DJB	JUN-JUL	MSI	176	175	0.8	140.3	NO
LA	JUN-JUL	MSI	200	199	7.9	161.8	NO
AFG	MAY	VIIRS	147	146	23.4	114.4	NO
AL	APR	VIIRS	182	181	3.3	145.6	NO
FL	MULTI	VIIRS	118	117	63.2	89.0	NO
IN	FEB	VIIRS	150	149	1.2	118.0	NO
IN	JUL	VIIRS	134	133	2.2	103.0	NO
LA	MAY	VIIRS	150	149	5.3	118.0	NO
LA	JUN	VIIRS	150	149	2.0	118.0	NO
LA	JUL	VIIRS	146	145	3.4	113.6	NO
NK	MAY	VIIRS	150	149	10.1	118.0	NO

Table 7: The number of parent images used to generate the synthetic background for each site by location, month, and instrument (ISNTR).

SITE	MONTH	INSTR	IMAGES
AFG	MAY-JUN	MSI	8
DJB	JUN-JUL	MSI	5
LA	JUN-JUL	MSI	6
AFG	MAY	VIIRS	93
AL	APR	VIIRS	178
FL	MULTI	VIIRS	454
IN	FEB	VIIRS	247
IN	JUL	VIIRS	238
LA	MAY	VIIRS	225
LA	JUN	VIIRS	268
LA	JUL	VIIRS	225
NK	MAY	VIIRS	151

The equivalence of the synthetic and observed truth data masks is examined before accepting synthetic truth masks as a means of validating ACAM masks generated for VIIRS imagery. Synthetic masks are not required for MSI-based masks. Two types of images are compared to the synthetic truth masks. The first represents daytime imagery and the second represents nighttime imagery. An operational mask from each category is compared to a synthetic truth mask. The daytime synthetic truth mask is compared to an operational cloud mask observed on the date of interest. It is cropped to share the same spatial domain. The observed mask is coded with zeros and ones for cloud-free and cloudy pixels, respectively. The synthetic truth mask is coded with zeros and twos for cloud-free and cloudy pixels, respectively. The masks are added using Raster Calculator. The output is displayed in Figure 66.

The percentage of pixels classified as cloud or cloud-free is 74.2 percent. The proportion of cloudy pixels mutually identified by both masks is 53.3 percent. The validation image contains widespread cirriform cloud cover with varying optical thickness, which the synthetic truth mask does not detect well, but optically thick clouds are well-represented in the mask. The synthetic mask also detects small clusters of cumulus clouds the operational mask does not. Because the ACAM developed in this paper is intended for ISR, the weaker detection of optically thin clouds by the synthetic truth mask is tolerable. Synthetic truth masks are not synthesized for scenes where the majority of clouds in the observed imagery are optically-thin. The comparison indicates that daytime synthetic truth masks are acceptable for validation.

The nighttime image used to validate synthetic truth masks was captured on September 11, 2017. The domain of the synthetic truth mask is cropped to match the operational truth mask. The comparison image is displayed in Figure 67. City lights identified by the ACAM are removed from the synthetic truth mask. The proportion of pixels classified either as cloudy or cloud-free is 86.8 percent. The proportion of

only cloudy pixels identified by both masks is 94.7 percent. Subjectively both masks perform well. Quantitative evaluation reveals only minor differences between the masks. The inability of the synthetic truth masks to detect optically thin cloud cover is more pronounced at night. The nighttime synthetic truth masks are determined to be acceptable for validation.

1. VIIRS Results

1.1 Afghanistan Results

The VIIRS-based background image for Afghanistan originally encompassed only the area within three geographic degrees of Bagram. The spatial extent of imagery available for this country allowed for the background image domain to be expanded across the entire country. The image is created from 93 parent images generated by the automated synthetic image workflow. Afghanistan is the only VIIRS site where less than 100 parent images are used. Observed imagery over this region exhibited frequent navigation issues that constrained the number of possible parent images. Widespread cloud shadows are present in the parent images that are usable. The final background is generated using the alternate automated method to offset these elements (see Figure 68). The alternate method results in a background image that is slightly blurred, but this does not impact the image's usability with the ACAM. The loss in visual clarity can be observed in Figure 69.

The loss in visual clarity complicates manual interpretation of the background. A stereographic analysis is employed to ensure pixel patterns line up correctly with topographic features from the data for Afghanistan. Terrain features appear realistic when compared to examples from cloud-free observed images mapped to the same DEM. Dendritic patterns and valleys also line up with peaks and valleys in the DEM. The alternate automated method expands the dendritic patterns associated




Figure 66: Comparison of operational and synthetic truth masks for VIIRS daytime imagery in (a). The observed image is given for reference in (b). The mask is translucent and overlaid atop an observed image from July 7, 2017. Locations that are identified mutually as cloud-free or cloudy are clear or white, respectively. Clouds identified by the operational truth mask only are colored red. Clouds identified by the synthetic truth mask are colored green. When compared to the observed image in (b), manual analysis revealed that the operational mask picks up more cirriform and optically thin cloud cover, but the synthetic truth mask picks up on more small cumulus clusters.




Figure 67: Comparison of operational and synthetic truth masks for VIIRS nighttime imagery. The mask is translucent and overlaid atop an observed image from September 11, 2017. Colors as seen in Figure 66. When compared to the observed image in (b), manual analysis revealed the operational mask outperforms the synthetic truth mask with cirriform and optically thin clouds better, but underperforms with small areas of cloud cover.




Figure 68: Final synthetic background image for Afghanistan during May.

with mountains across the country over larger spatial areas. This drives some differences from observed imagery, but better accounts for all possible locations of snow cover as surface snow evolves over time.

The chi-square test is run on the VIIRS Afghanistan background image using 147 randomly generated points across the image domain. The test has an associated 146 degrees of freedom. It is compared to a daytime image observed at 0909 UTC on May 15, 2015. Points associated with cloud cover or outside the AOI are removed from the sample set before calculating the test statistic. The associated critical value with this sample number is $114.4 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. The null hypothesis that the images are significantly different fails to be rejected after the test statistic is determined to be $23.4 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. The test statistic is small compared to the critical value. This




Figure 69: Subset images comparing the clarity of dendritic patterns and rivers over eastern Afghanistan. The left subset was taken from an observed image. The right subset was taken from the final synthetic background. Yellow circles highlight differences between the clarity of rivers between the images. Red circles highlight changes in dendritic snow patterns between images. Notice that while contrast increases with the synthetic image, land-water boundaries are less distinct compared to the observed image. Snow covers more of the mountaintops in the synthetic image, but land snow boundaries are also blurred in comparison to the observed image.

indicates a strong goodness-of-fit between the compared histograms. A comparison of both histograms can be seen in Figure 70. The background image has radiance values ranging from 2.1 to $179.4 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. While certain sample pixels are separated by more than $50 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$, 56 percent of the sample points fall below the sample mean of $19.3 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. It is determined that 90 percent of the sample points are less than $32 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$ apart. This is well below the $78.5 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$ mean radiance value of the background image. The highest values in the comparison are associated with areas of snow cover change or water level change near mountains or bodies of water, respectively (see Figure 71). Variations observed during these evaluations are expected because the background image only needs to approximate the associated observed image. This background image is determined to be acceptable for use with the ACAM.




Figure 70: Plot of the random sample points measuring the magnitude of difference between the image observed on May 15, 2015 and the synthetic background image for Afghanistan. The horizontal red line represents the sample mean of 19.3 $W m^{-2} sr^{-1} \mu m^{-1}$. The majority of sample points fall below the detection threshold of 45 $W m^{-2} sr^{-1} \mu m^{-1}$.




Figure 71: Image depicting the magnitude of pixel difference between the image observed on May 15, 2015, and the synthetic background image for Afghanistan. Differences range from 2.1 to $179.4 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. Tan shows the lowest values of change. Shades of red show areas of greater change, with the darkest reds associated with the highest values. Differences above the highest value observed from the sample points for the chi-square test can be seen in violet. Notice the darkest reds correspond frequently to areas where dendritic patterns can be found along mountains. The cluster of violet pixels in the upper left corner of the image is associated with cloud cover. Sample points are marked with dark green rhombi.

1.2 Alaska Results

The synthetic background image generated for Alaska encompasses a three-degree radius around Fairbanks. The final image is a polygon rather than circle due to the locations of co-located pixels across the parent imagery for this domain. Fairbanks is located at the highest latitude of all the sites examined in this paper. Due to the low sun angle during the winter months, April is chosen for analysis to ensure sufficient daylight is available to evaluate the ACAM. The final image is determined from the mean of 178 parent images created with the primary automated method. The background is displayed in Figure 72. The spatial extent of snow cover changes frequently across the 178 parent images, raising concerns that the ACAM would not be able to differentiate fresh snowfall from cloud cover due to the magnitude of difference between fresh snowfall in observed images and underrepresented surface snow values in the background. The influence of shadows cast by the mountains south of Fairbanks during early April is also seen near the center of the final image. Manual inspections of the parent images reveal the mountain shadow diminishes with time. The only evidence of this shadow in the background is the dark swath of pixels that run horizontally along the center of the image. Details of the valley below the mountains toward the bottom of the image can still be discerned manually, but fresh snowfall during late April will be bright enough that misclassification of pixels in this area is likely.

The chi-square test performed on this background image uses 182 sample points with an associated 181 degrees of freedom. It is compared to an observed image from April 21, 2017 at 2006 UTC. The resulting test statistic is determined to be $3.3 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$ with a critical value of $145.6 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. This indicates a strong goodness-of-fit between the two histograms. The null hypothesis that the images are significantly different fails to be rejected. The largest contributors to the




Figure 72: Final synthetic background image for Fairbanks, Alaska and the surrounding region for April.

chi-square statistic are samples taken over and near the mountains surrounding Fairbanks. Stereographic analysis reveals the terrain depicted in the image lines up well with the peaks and valleys from a DEM of the area (Patterson 2017). The differences between the histograms are displayed in Figure 73. The mean pixel difference is $7.7 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. Pixel differences range from 0.1 to $37.8 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. This background image is determined to be acceptable for use with the ACAM.




Figure 73: Plot of the random sample points measuring the magnitude of difference between the image observed on April 21,2017 and the synthetic background image for Fairbanks, AK. The horizontal red line represents the sample mean of $7.7 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. All sample points are below the detection threshold of $45 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$.

1.3 California May Results

The synthetic background image for the month of May for California uses images observed between 2012 and 2017. The domain extends to the state's boundaries as allowed by parent images. The same spatial domain is applied to the background images for June and July at this location as well. The May background is determined from the mean of 225 parent images. The background image is displayed in Figure 74. Details such as the dendritic snow patterns over California's mountains are faithfully represented when compared to observed images. No significant differences are noted when the background is qualitatively compared to the observed image from May 5, 2017 at 2051 UTC.

The chi-square test compares this background image to an observed image from May 5, 2017. It uses 150 sample points with an associated 149 degrees of freedom. The same set of random points is used to evaluate the June and July background images of California as well. The test statistic is determined to be $5.3 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$ with an associated critical value of $118.0 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. The largest differences between the two images are associated with areas of snow cover. The sample points and pixel differences can be seen in Figure 75. The differences between pixel values across the sample points are typical of other background images from this paper where the histograms exhibit strong goodness-of-fit. These differences are displayed in Figure 76. The background image is determined to be acceptable for use with the ACAM.




Figure 74: The final background image for California during May. Snow cover can still be seen over the mountains near the center of the image.




Figure 75: Difference between California's May synthetic background and the image observed on May 5, 2017. Sample points are marked by yellow rhombi.




Figure 76: Plot of the random sample points measuring the magnitude of difference between the image observed on May 20, 2017 and the synthetic background image for California. The horizontal red line represents the sample mean of $9.9 Wm^{-2}sr^{-1}\mu m^{-1}$. The majority of sample points are below the detection threshold of $45 Wm^{-2}sr^{-1}\mu m^{-1}$.

1.4 California June Results

The synthetic background image (see Figure 77) for the month of June for California uses images observed between 2012 and 2017. It covers the same domain as the California background for May and is determined using 268 parent images. Other than minor pixel changes associated with land cover, inland water level changes, or dendritic snow pattern evolution there are few differences to note between the June and May background images. The chi-square test compares the background to an observed image from June 24, 2016 at 2045 UTC. The test uses the same 150 sample points as the state's May background. The resulting test statistic is determined to be $2.0 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. It has an associated critical value of $118.0 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. The pixel differences between pixel values across the sample points are typical of other background images from this paper that exhibit strong goodness-of-fit. These differences extracted by the sample points are displayed in Figure 78. The background image is determined to be acceptable for use with the ACAM.




Figure 77: The final synthetic background image for California during June.




Figure 78: Plot of the random sample points measuring the magnitude of difference between the image observed on June 24, 2016 and the synthetic background image for California. The horizontal red line represents the sample mean of $6.4 Wm^{-2}sr^{-1}\mu m^{-1}$. All sample points are below the detection threshold of $45 Wm^{-2}sr^{-1}\mu m^{-1}$.

1.5 California July Results

The synthetic background image (see Figure 79) for the month of July for California uses images observed between 2012 and 2017. It covers the same domain as the California background for May and is determined using 225 parent images. Other than minor pixel changes associated with the evolution of land cover, inland water levels, or dendritic snow patterns there are few noticeable differences between this background image of California and those for May and June. The chi-square test uses an observed image from July 28, 2017 at 2101 UTC. The test only uses 146 of the 150 random sample points used to evaluate the May image. Four points are removed due to the presence of cloud cover. The resulting test statistic is determined to be $3.4 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. It has an associated critical value of $113.6 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. The differences between pixel values across the sample points are typical of other background images from this paper that exhibit strong goodness-of-fit. The mean difference of the compared pixels is $9.7 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. Sample values range from 0.1 to $44.1 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. Pixel differences at each point are displayed in Figure 80. The image is determined to be acceptable for use with the ACAM.




Figure 79: The final synthetic background image for California during July. Snow cover along the mountaintops at the center of the image has decreased from Figure 74, but still remains visible in the background image.




Figure 80: Plot of the random sample points measuring the magnitude of difference between the image observed on July 28, 2017 and the synthetic background image for California. The horizontal red line represents the sample mean of $9.7 Wm^{-2}sr^{-1}\mu m^{-1}$. All sample points are below the detection threshold of $45 Wm^{-2}sr^{-1}\mu m^{-1}$.

1.6 Florida Results

The synthetic background image generated for Florida using VIIRS DNB data leverages images observed on nights where lunar illumination allowed cloud cover and cloud-free surfaces to be manually distinguished. This occurred when lunar illumination was approximately 90 percent. Parent imagery allows the original spatial domain to be expanded across the entire state and into parts of the surrounding southeast United States. The imagery was observed between the months of July and October from 2012 to 2017 (NOAA 2017c). The resulting synthetic image is determined from the mean of 454 parent images. The spatial domain evaluated encompasses only Florida. The resulting image (see Figure 81) reveals an illuminated nighttime surface of the entire state.




Figure 81: Final synthetic background image for Florida at night. Parent images consist of moonlit nights observed from July to October from 2012 to 2017.

Manual inspection reveals details about the surface at night that extends beyond elements usually revealed by surface lights at night. Divisions between large parcels of land can be seen as well as land-water boundaries associated with inland bodies of water. A subset of the background image containing Lake Okeechobee is displayed in Figure 82.




Figure 82: Subset of Florida's synthetic background image. Lake Okeechobee can be seen near the center of the image. Differences between land cover types not in the immediate vicinity of lights can also be observed in the image.

The chi-square test utilizes 118 randomly selected sample points with an associated 117 degrees of freedom. The synthetic background is compared to an image of Florida observed on September 13, 2013 at 0635 UTC. No sample points are allowed outside state boundaries. Sample points are also not allowed over artificial

light sources. The difference between the two images and the positions of the sample points is displayed in Figure 83. The test statistic is determined to be $63.2 \times 10^{-9} \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$ with a critical value of $89.0 \times 10^{-9} \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. This indicates an acceptable goodness-of-fit between the image histograms. It is the weakest observed test statistic from all synthetic backgrounds evaluated in this paper. The weak goodness-of-fit is expected due to large shifts in radiance values associated with nightly variations in surface lights. Varying amounts of lunar illumination over time also influence the visibility of surface details and clouds. Differences between the observed and synthetic image histograms are displayed in Figure 84. The mean pixel difference is $1.0 \times 10^{-9} \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. Unscaled radiance values range from 4.0×10^7 to $19.9 \times 10^{-9} \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. The background image is determined to be acceptable for use with the ACAM.




Figure 83: Difference between Florida's synthetic background and the image observed at 0635 UTC on September 13, 2013. Sample points are marked by green rhombi. The largest differences are shown in shades of light gray and white and frequently near cities and other anthropogenic lights sources.




Figure 84: Plot of the random sample points measuring the magnitude of difference between the image observed on September 13, 2013 at 0635 UTC and the synthetic background image of Florida. The horizontal red line represents the sample mean of $1.0 \times 10^{-9} \text{ W m}^{-2} \text{ sr}^{-1} \mu\text{m}^{-1}$. The sample with the largest difference is associated with changes in radiance near city lights. Most sample points are below the detection threshold of $5 \times 10^{-9} \text{ W m}^{-2} \text{ sr}^{-1} \mu\text{m}^{-1}$.

1.7 Indiana February Results

The synthetic background image generated for Indiana during the month of February uses images observed between 2012 and 2017. No efforts are made to minimize the influence of snow cover in the background image. The spatial extent of synthetic parent images allows the evaluated domain to encompass the entire state. The area validated only includes pixels within state borders. The resulting background image was the product of 247 parent images and depicts both Indiana and the surrounding area (as seen in Figure 85). Seasonal snow cover is observed in the uppermost portion of the image. This is desirable because it reduces the influence of high-valued seasonal and fresh snow pixels on cloud detection.

The chi-square test is conducted using an observed image from February 17, 2017 at 1800 UTC. The test also uses 150 sample points with an associated 149 degrees of freedom. The test statistic is determined to be $1.2 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. The critical value is $118.0 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. The largest differences between the observed and synthetic image are associated with surface snow. The sample points and pixel differences can be seen in Figure 86. The differences between the observed and synthetic image histograms are displayed in Figure 87. The mean difference between pixels is $3.3 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. Values range from 8.0×10^{-3} to $24.8 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. These values indicate a strong goodness-of-fit. This image is determined to be acceptable for use with the ACAM.




Figure 85: The final synthetic background image for Indiana during the month of February. Snow can be seen along the Indiana-Michigan border in the uppermost parts of the image.




Figure 86: Difference between Indiana's February synthetic background and the image observed on February 17, 2017 at 1800 UTC. Sample points are marked by yellow rhombi. The white pixels near the top of the state are associated with snow.




Figure 87: Plot of the random sample points measuring the magnitude of difference between the image observed on February 17, 2017 at 1800 UTC and the synthetic background image for Indiana. The horizontal red line represents the sample mean of $3.3 Wm^{-2}sr^{-1}\mu m^{-1}$. All sample points are below the detection threshold of $45 Wm^{-2}sr^{-1}\mu m^{-1}$.

1.8 Indiana July Results

The synthetic background image of Indiana for the month of July uses images observed between 2012 and 2017. The domain validated encompasses all pixels within state borders. The resulting background is the result of averaging together 238 parent images (as seen in Figure 88). The city of Indianapolis stands out as a large cluster of dark pixels near the center of the image. A similar pixel pattern can be observed near South Bend.

The chi-square test is conducted using an observed image from July 02, 2017 captured at 1907 UTC. The observed image includes multiple small clusters of clouds throughout the bottom half of the state. Sample points are not allowed over these locations. There are 134 random sample points used for the test with an associated 133 degrees of freedom. The test statistic is determined to be $2.2 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. The critical value is $103.0 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. The test statistic indicates a strong goodness-of-fit between the histograms. Differences between the observed and synthetic images at each sample point are displayed in Figure 89. The mean magnitude of difference between the sampled pixels is $11.6 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. Difference values range from 4.7×10^{-2} to $41.6 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. Despite the spread in pixel values, the test indicates the images do not differ significantly. The background image is determined to be acceptable for use with the ACAM.




Figure 88: The final background image of Indiana for July.




Figure 89: Plot of the random sample points measuring the magnitude of difference between the image observed on July 2, 2017 and the synthetic background image for Indiana. The horizontal red line represents the sample mean of $11.6 Wm^{-2}sr^{-1}\mu m^{-1}$. All sample points are below the detection threshold of $45 Wm^{-2}sr^{-1}\mu m^{-1}$.

1.9 North Korea May Results

The synthetic background image of North Korea during May uses images observed between 2012 and 2017. The spatial domain encompasses the entire country. The background image is the result of averaging 151 parent images (as seen in Figure 90). Peaks and valleys are clearly observed throughout the image. The dendritic snow pattern atop Mount Peaktu is observed in the upper right corner of the image. Surface features from cloud-free portions of observed images are preserved well in the background image. Surface features in the vicinity of Pyongyang also appear less hazy in the background when compared to the background image.

The chi-square test compares the background image to an observed image of the country from May 19, 2017 at 0421 UTC. It uses 150 sample points. The test statistic is determined to be $10.8 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. The associated critical value is $118.0 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. The points and difference between pixel values are displayed in Figure 91. The test statistic indicates a strong goodness-of-fit. The mean magnitude of pixel difference across all sample points is determined to be $16.0 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$, with values ranging from 0.2 to $40.5 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. These differences are displayed in Figure 92. Despite the spread in pixel values, the images are determined not to be significantly different. The null hypothesis fails to be rejected. The background image is determined to be acceptable for use with the ACAM.




Figure 90: The final synthetic background image for North Korea during May.




Figure 91: Difference between North Korea's May synthetic background and the image observed on May 19, 2017. Sample points are marked by yellow rhombi.




Figure 92: Plot of the random sample points measuring the magnitude of difference between the image observed on May 19, 2017 and the synthetic background image for North Korea. The horizontal red line represents the sample mean of $16.0 W m^{-2} sr^{-1} \mu m^{-1}$. The spread of values about the mean is wide with respect to the maximum and minimum values, but is not large enough for the images to be considered significantly different. All sample points are below the detection threshold of $45 W m^{-2} sr^{-1} \mu m^{-1}$.

2. MSI Results

2.1 Afghanistan Results

The synthetic background image developed for Afghanistan using the MSI encompasses the region immediately around and to the southeast of Bagram. Parent images are based on data collected between May and June from 2016 to 2017. The background image is the product of eight parent images that enclose the maximum spatial extent common to inputs. The final image is cropped to a rectangle. Qualitative validation of the image reveals only minor differences between the synthetic background and observed images over cloud-free locations. The final background is displayed in Figure 93. Observed features that differ most from the synthetic background are the bodies of water near Bagram and dendritic snow patterns atop the mountains observed near the upper right corner of the background image.

The chi-square test compares the background image to an observed image of the country from June 1, 2016 at 0631 UTC. It uses 150 sample points. The test statistic is determined to be $2.6 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$, with an associated critical value of $118.0 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. The distribution of points, as well as the pixel difference between the two images, is displayed in Figure 94. While sample values exceed $100 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$, the calculated test statistic indicates a strong goodness of fit. The mean magnitude of pixel differences extracted from sample points is $222.3 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$, with sample values ranging from 1.6 to $3140.1 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. These differences are displayed in Figure 95. Some points differ from the background by more than $1000 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$, but 120 of the 150 sample points fall below the sample mean. The large variation relative to samples obtained from VIIRS imagery is expected due to the higher spatial resolution of MSI imagery. The null hypothesis that the images are significantly different fails to be rejected. This image is determined to be acceptable for use with the ACAM.




Figure 93: The final synthetic background for Bagram, Afghanistan and the surrounding region for May and June.




Figure 94: Difference between Bagram, Afghanistan's synthetic background and the image from June 1, 2016. Sample points are marked by yellow rhombi. The relatively bright area of land in the lower right corner of the image indicates the difference in spatial coverage spatial between the observed and synthetic images. Sample points are not taken from this area.




Figure 95: Plot of the random sample points measuring the magnitude of difference between the image observed on June 1, 2016 and the synthetic background image for Bagram, Afghanistan. The horizontal red line represents the sample mean of $222.3 Wm^{-2}sr^{-1}\mu m^{-1}$. The spread of values about the mean is wide with respect to the maximum and minimum values, but is not large enough for the images to be considered significantly different. The majority of points are below the detection threshold of $1000 Wm^{-2}sr^{-1}\mu m^{-1}$.

2.2 California Results

The synthetic background image (see Figure 96) created for Los Angeles, California is the result of averaging six parent images. Parent images are based on data collected between June and July from 2016 to 2017. Subjective validation reveals small areas with major variations between pixels from the observed and background images. These differences are attributed to changes in the cityscape. Pixel differences are less substantial further away from urban areas.

The chi-square test for this site uses an observed image captured on June 13, 2017 at 1829 UTC. The test also uses 200 sample points. The test statistic is determined to be $7.9 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$ with a critical value of $118.0 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. This indicates a strong goodness-of-fit between the compared histograms. The mean magnitude of pixel difference across the sampled pixels is $234.2 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$, with values ranging from 0.5 to $1879.5 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. The distribution of sample points across the image is displayed in Figure 97. The difference in sample values between the observed and synthetic image histograms are displayed in Figure 98. Despite the spread of values extracted by the samples, 134 of 205 points fall below the sample mean value. The null hypothesis that the images differ significantly fails to be rejected. This image is determined to be acceptable for use with the ACAM.



Figure 96: Final synthetic background image for Los Angeles, California during June and July.




Figure 97: Difference between Los Angeles, California's synthetic background and the image observed on June 13, 2017. Sample points are marked by green rhombi. Large differences between the sea's synthetic and observed values are associated with bright pixels. Areas of urban landscape in the upper left corner of the image also appear in bright shades.




Figure 98: Plot of the random sample points measuring the magnitude of the difference between the image observed on June 13, 2017 and the synthetic background image for Los Angeles, CA. The horizontal red line represents the sample mean of $234.2 Wm^{-2}sr^{-1}\mu m^{-1}$. The spread of values about the mean is wide with respect to the maximum and minimum values, but is not large enough for the images to be considered significantly different. The majority of points are below the cloud detection threshold of $1000 Wm^{-2}sr^{-1}\mu m^{-1}$.

2.3 Djibouti Results

The synthetic background image (see Figure 99) developed for Ambouli, Djibouti results from averaging five parent images. The parent images are based on data captured during June and July between 2016 and 2017. The final spatial domain for this site encompasses the region around Ambouli. Qualitative evaluation of the image reveals few significant differences between synthetic and observed images over cloud-free locations. The primary differences are attributed to evolving dendritic patterns along mountain ridges and changing water levels. Construction along Ambouli's coast is also observed.

The chi-square test compares the synthetic background image to an observed image of the region from July 5, 2017 at 0742 UTC. It also uses 176 sample points. The test statistic is determined to be $0.8 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$, with an associated critical value of $140.3 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. This is the smallest test statistic determined of all the background images evaluated in this paper. The magnitude of difference between the pixels at each sample point is displayed in Figure 100. The distribution of points across the image, as well as pixel differences, are displayed in Figure 101. The mean magnitude of difference between the images at the sample points is $113.3 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. Sample values range from 0.8 to $873.8 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. The null hypothesis fails to be rejected. The background image is determined to be acceptable for use with the ACAM.




Figure 99: The final synthetic background image for Ambouli, Djibouti and its surroundings during June and July.




Figure 100: Plot of the random sample points measuring the magnitude of the difference between the image observed on July 5, 2017 and the synthetic background image for Ambouli, Djibouti. The horizontal red line represents the sample mean of $113.3 Wm^{-2}sr^{-1}\mu m^{-1}$. The spread of values about the mean is wide with respect to the maximum and minimum values, but is not large enough for the images to be considered significantly different. All sample points are below the detection threshold of $1000 Wm^{-2}sr^{-1}\mu m^{-1}$.




Figure 101: Difference between Anbouli, Djibouti's synthetic background and the image observed on July 5, 2017. Sample points are marked by yellow rhombi. The white cluster of pixels in the upper left of the image is associated with cloud cover. Sample points are not taken from this area.

B. Analysis of VIIRS Cloud Mask Results

The ACAM's performance for the VIIRS platform is examined during both day and night. Overall performance of the ACAM is displayed in Table 8. Performance is given for each the masks compared to operational masks in Table 9. Results for comparisons using synthetic truth masks are given in Figure 10. Pixels are correctly classified by the ACAM relative to the VIIRS truth data for 84.3 percent of the pixels evaluated, while 2.8 percent and 12.9 percent of pixels are false positives and false negatives, respectively. False positive and negative pixels are broken down by synthetic and operational truth mask comparisons in Table 8 as well. Manual analysis revealed all mask types occasionally detected cloud-free pixels as cloud cover, but only extreme cases for any mask are noted below. This occurred more frequently with truth data than with the ACAM's masks. The cloud-free pixels most frequently misidentified by the ACAM masks in manual analyses were associated with fresh snow. Common weakness across the different mask types are difficult to quantify and therefore only discussed qualitatively.

Table 8: Summary of results for comparisons between ACAM and truth data masks.

Results	VIIRS	MSI
% CORRECT	84.3	91.8
% FALSE POSITIVE	2.8	3.1
% FALSE NEGATIVE	12.9	5.1
% OPER FALSE POSITIVE	25.7	5.1
% OPER FALSE NEGATIVE	1.2	3.1
% SYNTH FALSE POSITIVE	5.0	N/A
% SYNTH FALSE NEGATIVE	3.4	N/A

Table 9: Results for comparisons between ACAM and VIIRS operational masks.

SITE	DATE	% CORRECT	% FALSE POS	% FALSE NEG
CA	7/6/17	88.6	1.5	10.0
CA	7/19/17	93.5	6.3	0.3
CA	7/24/17	93.8	6.0	0.1
CA	7/31/17	96.8	3.0	0.2
FL	9/5/17	92.0	7.8	0.1
FL	9/6/17	92.7	7.1	0.2
FL	9/10/17	93.7	0.0	6.3
IN	7/11/17	77.7	0.9	21.3
IN	7/12/17	74.5	0.8	24.7
IN	7/13/17	79.9	4.4	15.7
IN	7/17/17	70.8	0.4	28.8
IN	7/19/17	60.6	0.1	39.4
IN	7/23/17	94.3	4.1	1.6
IN	7/27/17	96.1	1.5	2.4
IN	7/10/17	87.6	12.1	0.3

Table 10: Results for comparisons between ACAM and VIIRS synthetic masks.

SITE	DATE	% CORRECT	% FALSE POS	% FALSE NEG
AFG	5/9/17	92.3	5.5	2.2
AFG	5/13/17	95.3	3.2	1.5
AFG	5/16/17	94.8	3.8	1.3
FL	9/11/17	80.5	2.6	16.9
IN	2/23/17	87.5	0.8	11.7
IN	2/25/17	50.9	1.1	47.9
CA	5/5/17	93.9	0.7	5.5
CA	5/26/17	75.6	0.1	24.3
CA	5/30/17	70.4	0.3	29.3
NK	5/5/17	90.2	9.4	0.4
NK	5/14/17	91.9	7.9	0.2
NK	5/25/17	96.0	3.7	0.3

1. Daytime Results

Masks generated by the paper’s ACAM were compared against VIIRS truth masks for five daytime locations: Afghanistan, Alaska, California, Indiana, and North Korea. Results from each test can be seen in Tables 9 and 10. All images were validated quantitatively and subjectively. Select images from each location will be reviewed. The observed image is given for each example. An image with the mask comparisons overlaid atop the scene will follow.

1.1 Afghanistan Results

Afghanistan is validated using only synthetic truth data. Three images taken during May 2017 are evaluated. All three classify pixels correctly at least 92.3 percent of the time, with false positives occurring no more than 5.5 percent of the time and false negatives occurring no more than 2.2 percent of the time. The clouds detected on May 9, 2017 are displayed in Figure 102. The scene contains mostly cumulus and cirrus cloud types. An examination of the dendritic patterns reveals that both the ACAM’s mask and synthetic have issues handling some areas of snow, but the ACAM’s mask outperforms the truth mask. When compared to the observed image most of the clouds identified only by the truth mask in the upper right part of the image are associated with snow cover. The synthetic truth data picks up on some small areas of optically thin cloud better than the ACAM’s mask, but the ACAM detects more small cumulus clouds and picks up more optically thin clouds along the bottom of the Figure 102. Qualitatively the ACAM’s mask outperforms the truth mask. Quantitatively it performs well against the truth mask.

1.2 Alaska Results

Fairbanks, Alaska is evaluated with one image captured on April 26, 2017. Qualitatively, the ACAM's mask performs marginally well against the synthetic truth mask. Manual inspection of images captured prior to and following the tested scene reveal frequent changes in fresh snowfall coverage. The ACAM mask correctly detects 61.7 percent of pixels compared to the truth mask. False positives and false negatives make up 12.0 percent and 26.3 percent of the scene, respectively. While the ACAM does not have major issues distinguishing snow from clouds in other tests from this paper, it does not discriminate cloud cover over large areas of snowfall well for Fairbanks. This degrades the ACAM's performance. The mask is displayed in Figure 103.

Clouds are resolved relatively well by the truth data mask, but truth data still misidentifies some areas of river ice and snow as cloud cover. This can be seen in the upper left portion of Figure 103a. The truth mask detects the majority of the thin clouds, while the ACAM's mask mistakes most of these clouds as clear-sky pixels. This is attributed to the high reflectivity of snow present in the synthetic background image. The ACAM mask outperforms truth data near the center of the image. This is either indicates the background image represents this part of the scene well or that the bright surface snow the observed image is contaminating the optically thin clouds above. The mountain shadow's influence on the background image could also be influencing the difference in pixel values between the two images. The shadow is difficult to discern manually in the observed image, but the clouds are well lit above the center of the image. It is possible to adjust the input scene for the mountain shadow in ArcMap and build such a feature into the ACAM's workflow. ArcMap offers tools that can create artificial shadows over imagery using DEM data, but the influence of fresh snow on cloud detection must also be mitigated. Snow cover

enhances the ACAM’s performance over some locations and degrades it over others. The decision is made to exclude Alaska mask results and masks from other scenes with widespread snow from performance assessments for the ACAM. Until a future iteration of the model can account for scene changes due to fresh snowfall, the ACAM should not be used for scenes containing widespread or fresh snow cover.

1.3 California Results

Three synthetic truth masks and four operational truth masks are used to evaluate the ACAM’s performance over California during May and July. The truth masks are based on data from May and July of 2017. Results for each image are displayed in Tables 9 and 10. For the May masks, pixels were mutually identified as cloudy or cloud free at least 70.4 percent of the time in each scene. False positive and false negative pixels occur no more than 0.7 percent and 29.3 percent, respectively. The majority of differences between the two masks are attributed to the presence of optically thin cloud cover. This can be observed in Figure 104. Differences between mask quality can also be attributed to false positive or false negative pixels influenced by changes in the scene’s snow cover (see Figure 105).

Another source of difference between the masks is the quality of the VIIRS operational masks. The operational mask is observed to misclassify cloud-free pixels as cloudy over various locations of the scene displayed in Figure 106. The false positive pixels can be partly attributed to the lower resolution of the truth mask before resampling, but manual analysis of other operational masks reveals the operational algorithm occasionally classifies cloud-free pixels as cloudy. This frequently occurs near areas of optically thin cirrus or widespread cloud cover, but there are instances where no cloud cover is observed in the original imagery. The ACAM masks resolve optically thin clouds more frequently than synthetic truth masks, but the degraded

performance of synthetic truth masks with these cloud types was identified in Chapter III. Given the known issues with truth data, not all misclassified pixels can be attributed to ACAM error. There are instances where the ACAM masks outperform truth data that are not captured quantitatively. The algorithm performs well against both truth mask types qualitatively and quantitatively, but performance decreases slightly against synthetic truth data.

1.4 Indiana Results

Indiana masks are validated with ten images. Individual results can be seen in Tables 9 and 10. Eight operational cloud masks and two synthetic cloud masks are used as truth. The synthetic masks are used to evaluate ACAM performance for Indiana during February. The operational masks are used for Indiana during July. Pixels are mutually identified 87.5 percent of the time for the image from February 23, 2017. False positive and false negative pixels occur 0.8 percent and 11.7 percent of the time in the scene, respectively. This mask comparison is displayed in Figure 108a. Both compared masks resolve optically thin cirrus well. The ACAM mask outperforms truth data when optically thin clouds and small clouds are detected. The second image evaluated for February (Figure 107b) contains widespread, optically thin clouds. Pixels are mutually identified by both masks 50.9 percent of the time, with false positive and false negative pixels occurring 1.1 percent and 47.9 percent of the time, respectively. When optically thin clouds allow underlying surface details to be distinguished, the ACAM's inability to detect them does not necessarily indicate a weakness. If not identified as cloud cover, the details below the cloud may add value to an ISR analysis of the scene. There is little snow cover in either of the evaluated images to test ACAM's ability to resolve clouds over seasonal snow cover at this site, but the ACAM delivers acceptable performance for these winter season examples.

Eight of the ACAM's masks are compared to operational masks over Indiana during July. Correctly classified pixels range from 60.6 percent to 96.1 percent for the comparisons. The highest amount of false positive pixels from a scene is 12.1 percent. The highest amount of false negative pixels among the comparisons is 39.4 percent. Manual inspection reveals the main contributor to mask quality differences is optically thin cloud cover. This can be observed in Figures 108 and 109. Due to issues with widespread snowfall mentioned previously in subsection 1.2 of this chapter, images exhibiting large areas of surface snowfall are not included with Indiana's February results. The ACAM outperforms operational masks with small cumulus clouds. While the ACAM mask detects some areas of thin cloud cover, it underperforms with this cloud type when compared to operational truth masks. The ACAM performs best with cumulus clouds of any size above the size threshold and frequently outperforms operational masks with small clouds near the MCAT. Despite issues with optically thin clouds, the ACAM delivers acceptable masks for Indiana during July.

1.5 North Korea Results

The ACAM is evaluated with North Korea using three synthetic truth masks over North Korea during May. Individual results can be seen in Table 10. Pixels are mutually identified by both masks at least 90.2 percent of the time for all images. No more than 9.4 percent of any comparison is associated with false positive pixels. No more than 0.4 percent of any comparison is associated with false negative pixels. Qualitative analysis of each mask comparison reveals truth data outperforms ACAM masks with optically thin clouds, but ACAM masks consistently resolve more small clouds. This is indicated by the results displayed in Figure 110. The ACAM performs well against operational truth data and marginally well against synthetic truth data. Qualitatively and quantitatively, the ACAM delivers acceptable results for this site.

2. Nighttime Results: Florida

The ACAM's performance is evaluated at night over Florida using four truth data masks. Three comparisons use operational mask data and one uses synthetic mask data as truth. The results for each comparison are displayed in Tables 9 and 10. The ACAM performs the weakest against the September 11, 2017 synthetic truth mask. This comparison is displayed in Figure 112. The ACAM performs well against each operational masks comparisons. The lowest proportion of mutually classified pixels comparing the masks is 92.0 percent. The highest proportion of false positive pixels is 7.8 percent, while the highest proportion of false negative pixels is 6.3 percent. A representative comparison is displayed in Figure 111. Between the September 10 and 11 images, Hurricane Irma (2017) made landfall on Florida. In both mask comparisons with this hurricane, truth data consistently outperforms the ACAM with detection of optically thin clouds. Qualitative inspections of these mask comparisons reveal that the ACAM struggles to identify optically thin clouds more frequently at night than during the day. The ACAM performs best against operational truth data, with mask quality decreasing only slightly against the synthetic truth mask. Qualitatively and quantitatively, the ACAM delivers acceptable performance for sufficiently illuminated nights over Florida.




Figure 102: Mask comparison image overlaid atop the observed image from May 9, 2017 in (a). The observed image is given for reference in (b). Mask comparison is translucent. Clouds identified by both masks are colored gray or white. Clouds identified only by the ACAM's mask are green. These are referred to as false positive pixels. Clouds identified only by the truth data mask are red. These are referred to as false negative pixels.




Figure 103: Comparison of cloud masks on April 26, 2017 over Fairbanks, Alaska in (a) with colors as in Figure 102. The observed image is given for reference in (b).




Figure 104: Mask comparison over California on May 30, 2017 in (a) with colors as in Figure 102. The observed image is given for reference in (b).




Figure 105: Subset of mask comparison over the Sierra Nevada Mountains in eastern California on May 30, 2017 in (a) with colors as in Figure 102. The observed image is given for reference in (b).




Figure 106: Subset of mask comparison over Los Angeles, California on July 24, 2017 in (a) with colors as in Figure 102. The observed image is given for reference in (b).




Figure 107: Mask comparison for Indiana on February 23, 2017 in (a) with colors as in Figure 102. The observed image is given for reference in (b).




Figure 108: Mask comparison for July 23, 2017 over Indiana in (a) with colors as in Figure 102. The observed image is given for reference in (b).




Figure 109: Mask comparison for July 27, 2017 over Indiana in (a) with colors as in Figure 102. The observed image is given for reference in (b).




Figure 110: Mask comparison for May 5, 2017 over North Korea in (a) with colors as in Figure 102. The observed image is given for reference in (b).




Figure 111: Mask comparison for September 10, 2017 over Florida in (a) with colors as in Figure 102. The observed image is given for reference in (b).




Figure 112: Mask comparison for September 11, 2017 over Florida in (a) with colors as in Figure 102. The observed image is given for reference in (b).

C. Analysis of MSI Cloud Mask Results

MSI imagery is evaluated to determine the ACAM’s ability to detect clouds at multiple resolutions. The spatial resolution is 20 *m* for all MSI images examined in this paper. Overall, the ACAM masks compare well with truth data. The lowest proportion of correctly, or mutually classified, pixels is 85.8 percent. All but two of the images are classified correctly less than 90.4 percent of the time. Part of the ACAM’s increased performance is attributed to the smaller difference between compared mask resolutions relative to the compared VIIRS masks. Qualitative and quantitative evaluations of each MSI mask are performed. Examples from each site are also reported. The results from each test are displayed individually in Table 11. These evaluations indicate that the ACAM delivers acceptable performance at both levels of spatial resolution evaluated in this paper.

Table 11: Results from comparisons between the ACAM and MSI masks.

SITE	DATE	% CORRECT	% FALSE POS	% FALSE NEG
BAG	6/16/16	96.5	1.9	1.5
BAG	5/12/17	90.4	1.0	8.5
DJB	7/20/17	86.8	1.2	12.0
DJB	7/25/17	98.8	1.1	0.1
LA	6/8/16	92.3	1.9	5.9
LA	6/28/16	85.8	11.6	2.6

1. Bagram, Afghanistan Results

The images available for Bagram, Afghanistan are limited for the months of May and June. Those available with any cloud cover are either completely obscured or contain only small areas of cloud cover. The images selected for evaluation have limited areas of cloud cover to better test the ACAM's ability to differentiate clusters of small clouds from their surroundings. The lowest proportion of correctly identified pixels from both Bagram images is 90.4 percent. The highest false positive and false negative proportions are 1.9 percent and 8.5 percent, respectively. Individual results from these tests are displayed in Table 11. As with VIIRS images, optically thin clouds are better resolved by the operational mask than the ACAM's mask. Conversely, the operational mask struggles more with snow cover and detects cloud cover that is not distinguished with manual analyses. This is displayed in Figures 113 and 114. Based on the findings from the test images, the ACAM's masks outperforms the operational masks qualitatively and compares with them well quantitatively.

2. Ambouli, Djibouti Results

The masks compared over Ambouli, Djibouti are derived from data collected during July of 2017. The observed images contain various cloud types, including optically thin cirrus. The lowest proportion of correctly classified pixels for this site is 86.8 percent. The highest proportion of pixels determined to be false positives or false negatives are 1.2 percent and 12.0 percent, respectively. Individual results are displayed in Table 11. Manual inspection reveals that the ACAM struggles to identify thin clouds with both images. This is displayed in Figure 115. Manual inspections also reveal the ACAM resolves small cloud clusters better than the operational mask. Based on findings from the test images, the ACAM masks perform marginally well against operational masks qualitatively and well against them quantitatively.

3. Los Angeles, California Results

The masks evaluated over Los Angeles, California are derived from data collected during June 2016. The lowest proportion of correctly classified pixels for this site is 85.8 percent. The highest proportions of pixels determined to be false positives or false negatives are 11.6 percent and 5.9 percent, respectively. Individual results are given in Table 11. As with the previously discussed MSI images, the ACAM struggles most with optically thin clouds and correctly resolves small clouds more frequently than the operational masks. The operational mask continues to flag cloud-free locations as cloudy both near and far from areas of cloud cover as well. This could be attributed to spectral changes in the cityscape not resolved well by the operational cloud mask, but the exact cause could not be determined. Most of the bright city features fall below the ACAM's MCAT and are not flagged as cloud cover. The mask comparison is displayed in Figure 116. The ACAM's masks outperform the operational masks qualitatively and compare well with them quantitatively. As with the other the MSI sites, the ACAM performs well with this location.




Figure 113: Mask comparison for Bagram on May 12, 2017. Red-colored pixels at the bottom of the image are not all associated with cloud cover in (a) with colors as in Figure 102. The observed image is given for reference in (b).




Figure 114: Subset of the May 12, 2017 mask comparison near Bagram, Afghanistan in (a) with colors as in Figure 102. While both masks flag parts of the mountain snow cover as clouds, the ACAM’s mask outperforms the operational mask. The operational mask struggles most with the optically thin clouds surrounding the cumulus clouds near the center of the image. The observed image is given for reference in (b).




Figure 115: Mask comparison for Ambouli, Djibouti on July 20, 2017 in (a) with colors as in Figure 102. The observed image is given for reference in (b).




Figure 116: Mask comparison for Los Angeles, California on June 8, 2016 in (a) with colors as in Figure 102. The observed image is given for reference in (b). Issues with the operational mask flagging non-cloudy pixels as cloudy can be seen in the upper half of the image.

D. Cross-Platform Background Test

The ability to detect cloud cover in high-resolution imagery is desirable, but applying the technique presented in this paper requires background files that increase in size as spatial resolution becomes greater. Small synthetic background images are preferred for operational applications of the ACAM. One method for reducing file size is presented in this paper. At the cost of degrading background image quality, a cross-platform background image for Los Angeles, California is developed and tested for use with MSI images using VIIRS background images. The compared channels have the same central wavelength, but have different spatial resolutions. Differences between the background images are evaluated using the chi-square test. Masks created using the original and cross-platform backgrounds are also compared.




Figure 117: The mean image of Los Angeles determined by averaging pixel values from California's June and July VIIRS background images.

The original MSI image for Los Angeles, California is composed of parent images from June and July. The backgrounds from VIIRS images of California for these months are averaged together using ArcMap's Raster Calculator to compensate for the time period difference. The resulting image is clipped to the spatial extent of the MSI-based background. ArcMap's Rescale by Function Tool is used to recode radiance values from the VIIRS-based background to approximate those in the MSI-based background. This is accomplished by supplying the tool with spectral characteristics from both images. The tool requires the upper and lower bounds of each image to perform the transformation. The VIIRS background image has a minimum radiance value of $5.0 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$ and a maximum radiance value of $124.5 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$. The MSI background has a minimum radiance value of $5 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$ and maximum radiance value of $8043.1 \text{ Wm}^{-2}\text{sr}^{-1}\mu\text{m}^{-1}$.

A transformation function must be selected to recode the pixel values from the input image. The tool's logistic growth function is chosen. This function changes high-value pixels from the scene most and small-value pixels least (ESRI 2007b). This setting is selected after manually comparing the output from each of the function options to the original MSI background image. The tool also resamples the recoded image to the resolution of the MSI image. It uses the nearest neighbor method to resample the recoded image. This technique determines new pixel values using the value of the original pixel located closest to the pixel being created to generate an image with higher spatial resolution (McGrew et al. 2014). The original VIIRS image is displayed in Figure 117. The cross-platform and original MSI image are displayed in 118.




Figure 118: The final cross-platform synthetic background image for Los Angeles, California is displayed in (a). It is comprised of VIIRS backgrounds from June and July. The MSI background for this time period is displayed in (b).

Sample points are randomly generated across the domain, but are not allowed to occur over water. Simulated water values are known from the VIIRS imagery and observed to differ little from the MSI background image. Including them in the chi-square test will introduce positive bias into the results of the chi-square analysis.

After removing sample points outside state borders, 89 sample points remain with an associated 88 degrees of freedom. The magnitude of difference between the two images is displayed in Figure 119. The calculated test statistic is $13.1 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. The critical value is determined to be $64.8 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. The null hypothesis that the images are not significantly different fails to be rejected. The size of the chi-square test statistic indicates that the goodness of fit between the images is strong. The magnitude of difference between pixel values at sample points range from 35.6 to $3455.2 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$, with a mean value of $1578.3 \text{ } Wm^{-2}sr^{-1}\mu m^{-1}$. Qualitatively, higher contrast is apparent in the cross-platform image when it is compared to the MSI background. The cross-platform image also exhibits a blurred appearance. This is expected given the difference in spatial resolution between the two radiometers. Despite the differences between from the original synthetic background for Los Angeles, the cross-platform image is determined not to be significantly different and is acceptable for use with the paper's ACAM.

The cross-platform background is tested against an observed MSI image from June 8, 2016. The MSI ACAM settings are used to generate masks with both background images. Each mask is manually inspected before quantitative tests are performed. The mask associated with the cross-platform background detects clouds over the sea near Los Angeles well, but mask quality declines slightly inland. The mask also flags clear-sky areas of the cityscape as cloud cover. This is likely due to the presence of bright pixels in the MSI background that are not resolved with VIIRS due to the sensors coarser resolution. The lower input values are not recoded high enough by the rescaling tool as a result. The change detection image resulting from the sum of both masks and the original image can be seen in Figure 120. Manual examination reveals that mutually identified pixels compose 95.3% of the change detection image. False positive and false negative pixels comprise 0.4 and 4.3% of the scene, respectively. All




Figure 119: Magnitude of difference between the original and cross-platform background images for Los Angeles, California. Sample points are marked with green rhombi. The largest differences are associated with shades of white and mostly occur over the cityscape.

pixels not mutually identified by the masks are flagged incorrectly. No cloud mask pixels associated with the cross-platform image are observed to outperform pixels from the other mask. This degradation of quality is expected, as low-resolution input image is expected to underperform high-resolution input. The cross-platform mask is determined to perform exceptionally considering the difference the large difference in spatial resolution between the radiometers. This provides evidence that a cross-platform image can act as an acceptable surrogate in some instances for sites where cloud-free images are not available for cloud detection.

The cross-platform file is also smaller in size compared than the original background file. The new background is 395.0 KB before resampling. Compared to the

MSI background's 115.0 MB file size, the cross-platform image is 99.7% smaller. This represents a massive decrease in file size with relatively little loss in quality of the resulting cloud mask. If cross-platform background images perform well in future research, given the Earth's surface area is $5.1 \times 10^{11} \text{ km}^2$ and the background tested covers $1.5 \times 10^4 \text{ km}^2$ according to NASA (2017), a global mosaic of background images could be created that occupies 13.5 GB of file space, rather than the 3.9 TB its high-resolution counterpart would consume. The decrease in file size would allow a global mosaic created by this method to be built and saved on the hard drive of an off-the-shelf computer system.




Figure 120: Mask comparison between ACAM runs using the cross-platform background and original MSI synthetic background from June 6, 2016 for Los Angeles, CA is displayed in (a). The mask comparison is translucent and overlaid atop the background image. Cloud-free pixels mutually identified by both masks are not assigned colors. Clouds identified only by the original background's mask are colored red. Clouds identified only by the cross-platform mask are colored green. The original image is displayed in (b).

V. Conclusion

A. Summary of Results

Clouds can cover 70 percent of the planet at any given time (Stubenrauch et al. 2013). Identifying cloud cover in satellite imagery is possible using manual techniques, but is time-consuming and subjective. Automated methods are faster than manual analyses, but can be less reliable. The varying shapes and spectral signatures of clouds have made them difficult to detect by remote sensing since the earliest automated techniques were devised. Validation of detection methods is also complicated by the difficulty of defining cloud cover quantitatively and obtaining reliable truth data (Girolamo and Davies 1997). Several detection methods and their constraints are given in this paper. Imperfections in cloud detection methods are acceptable and expected, but negatively influence the quality of cloud masks and derivative products.

The ACAM presented in this paper relies on threshold-based change detection and synthetic background images to identify areas of cloud cover. It uses a meteorological definition of cloud cover to define clouds qualitatively, but applies an ISR approach to quantifying them. Detecting any cloud thin enough where details from the scene below can be interpreted is desirable, but does not necessarily detract from the value of a product. The inclusion of both spectral and geometric thresholds are used to minimize the issues found with simple threshold methods, without adding the complexity of multispectral or radiative transfer techniques. The cloud detection method in this paper balances simplicity with accuracy to create a cloud mask product that can be derived from multiple platforms at multiple resolutions. With improvements, it could also be applied to meteorological products where detection of

optically thin clouds has higher priority. Better cloud detection is desirable for both audiences because it improves product quality and enables better decision making.

Three methods are presented in this paper for generating synthetic background images that use a first-guess image. Backgrounds can be built either manually or algorithmically. The automated methods allow for large datasets to be incorporated into the final synthetic images. Backgrounds generated during this paper were evaluated qualitatively (by manual inspection) and quantitatively (by chi-square tests). Test sites were chosen at multiple latitudes throughout different seasons. Backgrounds were constructed at multiple resolutions for sites with both limited and robust data archives (see Table 6). A method for generating backgrounds for use between platforms was also presented. It delivered a mask with acceptable quality. The coarser resolution of the cross-platform background image reduced the file's size 99.7 percent with only a 4.7 percent decrease in the ACAM's ability to detect clouds. The methodology used to generate the final background images produced reliable emulations of all sites examined in the paper. In addition to serving as input for the ACAM, these cloud-free images can be used as stand-alone products.

Quality of synthetic background images was constrained by two primary factors: navigation errors and cloud shadows. These issues impacted the generation of Afghanistan's VIIRS-based synthetic background image most. Manual inspection revealed numerous observed images where pixels had been incorrectly georeferenced. Some surface features shifted slightly between images as a result. The high frequency of cloud cover over large areas of the country also led to quality issues with the background by introducing cloud shadows into many of the cloud-free images. The shadows left artifacts resembling hills and valleys in the background imagery that were not present in either observed images or DEM topography. This was overcome by devising the alternate automated background generation method which used a

spectral smoothing technique to reduce the influence of navigation error and minimize the number of cloud shadows present in synthetic parent images. Frequent cloud cover also degraded the quality of the Fairbanks, Alaska background image. A background image was created that did not differ substantially from observed imagery, but widespread cloud cover and frequent shifts in snow cover patterns in the background's parent images degraded its performance. Manual inspection was used to conduct quality checks of each background using observed imagery of sites. DEMs were referenced as necessary. Both daytime and nighttime background images frequently appeared less hazy than observed images. The nighttime background of Florida also revealed further details about the moonlit surface at night than what could be gathered from a single observed image. For each site, the spatial extent of the input imagery also allowed the domain evaluated to be expanded. This is most apparent in Figure 74, where the cloud-free scene includes portions of Arizona, New Mexico, Oregon, Utah, and Nevada. These parts of the background were not subjected to formal evaluations, but manual inspection indicates that larger background images than those presented in this paper can be generated.

The cloud detection technique developed for this paper was tested on both daytime and nighttime imagery. Static threshold settings were tailored to suit each platform and channel. Masks were evaluated using both qualitative and quantitative analyses. Manual inspections were performed by evaluating each mask generated by the ACAM individually and by comparing them to truth data masks. Truth data incorporated operational cloud masks from each platform and synthetic truth masks generated using Erdas Imagine. Contingency tables were used to assess mask performance quantitatively. The technique's performance was measured by determining the proportions of pixels in each image classified correctly or incorrectly with respect to truth data masks. The technique classified pixels correctly 84.3 percent on average

for VIIRS imagery and 91.8 percent of the time for MSI imagery (see Table 8). The evaluations did not include widespread, fresh snow.

Daytime masks created for this paper generally outperformed truth data masks with seasonal snow cover, such as that associated with mountain tops. Changes in dendritic patterns between months could be distinguished between the three synthetic backgrounds for California. Manual inspection revealed that optically thin clouds presented the greatest challenge to the ACAM. Nighttime masks performed marginally well, but optically thin clouds were resolved less often than they were with daytime masks. Individual results are given in Tables 9-11. The results from both day and night tests were deemed acceptable when compared to imperfect truth data.

The current state of the cloud detection technique presented in this paper is not yet accurate enough to support operational users throughout the year and at all latitudes, but delivers promising results for applications in low latitudes or areas that remain mostly snow-free throughout the year. Further development is required before it is ready to use for ISR purposes or by the meteorological community. This paper concludes the technique presented is a valid method for detecting cloud cover. It was determined not to be platform-specific and functions at multiple resolutions. Its validity is especially true when detection of optically thin clouds is not a major concern to the end-user. Cloud masks resulting from the technique compared well on average with operational masks during both day and night. In many cases, the ACAM's masks outperformed operational cloud masks. This was especially true with clouds that were small relative to the horizontal resolution of each evaluated sensor. Simple spatial and spectral thresholds allowed this technique to perform well against operational methods without their added complexity or additional channels.

B. Recommendations for Future Research

Further development of the technique presented in this paper should focus firstly on improving the quality of cloud mask output with the ACAM's current architecture. Efforts should be focused on finding the best threshold settings that work either throughout the year or by season. Additional thresholds or bins for radiance values will increase the ACAM's ability to detect optically thin clouds. Methods for obtaining and generating truth data should also be sought to better test the strengths and limitations of the technique. Products such as daily cloud coverage charts or Lidar-based cloud masks can be used to expand the volume of truth data for validating masks. Future efforts should investigate software capable of synthesizing the component files to increase available truth data as well. Testing the technique with other sensing platforms and at other resolutions would also be beneficial. Developing a Lidar-based synthetic background image for a given area would also be possible with minimal changes to ACAM's design. This would allow the mean state of a given location to be generated in three-dimensions.

The technique presented in this paper produced acceptable background imagery, but background synthesis methodology can be improved. Dendritic snow patterns changed regularly across observed images. Future efforts should focus on a method for smoothing pixels around certain features in background imagery, such as locations above certain elevations where seasonal snow cover is likely. Pixel smoothing can be restricted to certain portions of a scene by creating a binary mask that indicates where smoothing is or is not allowed to occur. This technique can also be applied to bodies of water where land-water boundaries change regularly due to factors such as changes to water level or to the flow paths rivers and streams. Implementation of these features can be done using conditional statements within Raster Calculator.

For daytime background images, synoptic or hemispheric-scale synthetic background mosaics should be developed to evaluate and hone the technique over large areas. Demonstrating the scalability of the technique beyond the sites presented in this paper would further demonstrate its value. An image of the continental United States can be synthesized by building on imagery from this paper. Creating a method for giving the most recently observed images greater weight in the final background would also add value to the ACAM by further refining cloud detection capabilities. Additionally, the influence of semi-permanent clouds on background imagery should be investigated. This is especially important for locations where features such fog, maritime stratus, or terrain-induced standing clouds can routinely obscure the surface. Finally, background images that cover extended time periods and allow for radiance adjustments like those created by Zhu et al. (2015) should be developed.

For nighttime backgrounds, methods to improve the surface light mask used in conjunction with the background image should be investigated. The current technique does not allow for cloud detection over pixels corresponding to surface lights. Clouds moving over these lights were observed to decrease radiance values directly over pixels identified as lights during manual inspections, while also increasing the radiance of cloudy pixels in the vicinity of the original light source during manual image inspections. Devising a method for detecting clouds that move over light sources would increase the value of the final masks. Oil rigs were observed to occupy multiple pixels with bright light as well. The binary light masks in this paper incorporated oil platforms, since their positions were static. A shapefile with points plotting positions of semi-permanent surface lights can be used to designate pixels near the structures that should be treated separately during cloud detection. Point positions can then be updated on a regular basis to account for changes in oil rig locations. Similarly, the ability of the ACAM to detect sufficiently illuminated areas after sunset and before

sunrise should be investigated. The number of moonlit nights is limited each month, but the day-night terminator presents itself daily at most locations across the planet. Investigating whether the limited illumination supplied by light just before sunrise and after sunset is sufficient for cloud detection to produce reliable cloud masks.

The ACAM presented in this paper was designed around a modular architecture to allow for future feature detection capabilities to be added in the future. These modules should first focus on honing the ACAM's ability to handle snowfall and sea ice. Daily snow mask products can be used to generate binary masks that alter cloud detection thresholds in Raster Calculator. This could either be incorporated into the current ACAM or be added to it as a separate module. Other modules should focus on identifying non-cloud elements in a scene, such as sun glint or airborne aerosols, once clouds are masked out. Mountain shadows are given as a potential source for Fairbanks cloud masks in Chapter IV. Future research should also explore a model that mitigates terrain shadow influence. ArcMap's Hillshade Tool or Skyline Barrier tool can be used to build another module that would handle topographic shadow evolution for the background image. These tools allow for solving shadow problems due to natural terrain or anthropogenic structures. Once connected to the ACAM, each additional module can rely on multivalue mask output from prior modules to simplify the detection process. The task of characterizing all features from a scene will become less complex as each module masks out more of the scene. An advanced version of the system can also run each of these modules across multiple channels to create a multispectral detection algorithm capable of functioning independently of platform and spatial resolution. This would come at the cost of the ACAM's simplicity, but opens it up to further applications.

Bibliography

- Campbell, J. B., and R. H. Wynne, 2011: *Introduction to Remote Sensing*. 5th ed., The Guilford Press, New York, NY, 667 pp.
- Cao, C., J. Xiong, S. Blonski, Q. Liu, S. Uprety, X. Shao, Y. Bai, and F. Weng, 2013: Suomi NPP VIIRS sensor data record verification, validation, and long-term performance monitoring. *J. Geophys. Res. Atmos.*, **118**, 11664–11678.
- Environmental Systems Research Institute, cited 2007a: ArcGIS. [Available at <http://www.arcgis.com>.]
- , cited 2007b: *ArcGIS desktop*. [Available at <http://desktop.arcgis.com>.]
- European Space Agency, 2015a: *Colour Vision for Copernicus. The Story of Sentinel-2*. 2 pp. [Available at <http://esamultimedia.esa.int>.]
- , 2015b: *Sentinel-2 User Handbook*. 64 pp. [Available at <https://sentinels.copernicus.eu>.]
- , cited 2017a: Copernicus open access hub. [Available at <https://scihub.copernicus.eu>.]
- , cited 2017b: Sentinel-2 MSI technical guide: Cloud masks. [Available at <https://sentinel.esa.int>.]

Evans, H., J. Lange, and J. Schmitz, 2014: *The Phenomenology of Intelligence-*

Focused Remote Sensing. 1st ed., Riverside Research, New York, NY, 528 pp.

Girolamo, L. D., and R. Davies, 1997: Cloud fraction errors caused by finite

resolution measurements. *J. Geophys. Res.*, **102**, 1739–1756.

Glickman, 2000: *Glossary of Meteorology*. 2nd ed., American Meteorological Society,

Boston, MA, 850 pp.

Hall, R. C., 2001: *A History of the Military Polar Orbiting Meteorological Satellite*

Program. Office of the Historian, National Reconnaissance Office, 56 pp.

[Available at <http://www.nro.gov>.]

Hecker, C., M. van der Meijde, and F. D. van der Meer, 2010: Thermal infrared

spectroscopy on feldspars - Successes, limitations and their implications for

remote sensing. *Earth-Sci. Rev.*, **103**, 60–70.

Hexagon Geospatial, cited 2017: Producer field guide. [Available at [https://](https://hexagongeospatial.fluidtopics.net)

hexagongeospatial.fluidtopics.net.]

- Hutchison, K. D., A. K. Heidinger, T. J. Kopp, B. D. Isager, and R. A. Frey, 2014: Comparisons between VIIRS cloud mask performance results from manually generated cloud masks of VIIRS imagery and CALIOP-VIIRS match-ups. *Int. J. Remote Sens.*, **35**, 4905–4922.
- Hutchison, K. D., J. K. Roskovensky, J. M. Jackson, A. K. Heidinger, T. J. Kopp, M. J. Pavolonis, and R. Frey, 2005: Automated cloud detection and classification of data collected by the Visible Infrared Imager Radiometer Suite (VIIRS). *Int. J. Remote Sens.*, **26**, 4681–4706.
- Kelso, N. V., T. Patterson, D. Furno, T. Buckingham, N. Springer, and L. Cross, cited 2017: Natural Earth. [Available at <http://www.naturalearthdata.com>.]
- Kidder, S. Q., and T. H. Vonder Haar, 1995: *Satellite Meteorology: an Introduction*. 1st ed., Academic Press, Inc., San Diego, CA, 466 pp.
- McGrew, J. C., A. J. Lembo, and C. B. Monroe, 2014: *An Introduction to Statistical Problem Solving in Geography*. 3rd ed., Waveland Press, Inc., Long Grove, IL, 320 pp.
- Mueller-Wilm, U., O. Devignot, and L. Pessiot, 2016: *L2A Product Definition Document*. Tech. Rep., ESA, 35 pp.

National Aeronautics and Space Administration, 2014: Joint Polar Satellite System

(JPSS) VIIRS Cloud Mask (VCM) Algorithm Theoretical Basis Document

(ATBD), Version E. 117 pp. [Available at <https://www.star.nesdis.noaa.gov/>.]

———, cited 2017: *Earth: By the numbers*. [Available at

<https://solarsystem.nasa.gov/>.]

National Institute of Standards and Technology, and Semiconductor Manufacturing

Technology, cited 2013: *NIST/SEMATECH e-handbook of statistical methods*.

[Available at <http://www.itl.nist.gov/>.]

National Oceanic and Atmospheric Administration, cited 2017a: NOAA's

comprehensive large array-data stewardship system. [Available at

<https://www.class.ngdc.noaa.gov/>.]

———, cited 2017b: VIIRS bands. [Available at <http://rammb.cira.colostate.edu/>.]

———, cited 2017c: Visible infrared imaging radiometer suite data records from

Suomi NPP. [Available at <https://www.class.ncdc.noaa.gov/>.]

Patterson, T., cited 2017: CleanTOPO2: Edited SRTM30 plus world elevation data.

[Available at <http://www.shadedrelief.com/>.]

Rogalski, A., 2010: History of infrared detectors. *Opto-Electron. Rev.*, **18**, 121–136.

Rogalski, A., and K. Chrzanowski, 2014: Infrared devices and techniques (revision).

Metrol. Meas. Sys., **21**, 565–618.

Space Science and Engineering Center, cited 2017: Polar2Grid documentation.

University of Wisconsin, [Available at <http://www.ssec.wisc.edu>.]

Stubenrauch, C. J., and Coauthors, 2013: Assessment of global cloud datasets from satellites: Project and database initiated by the GEWEX radiation panel. *Bull. Amer. Meteor. Soc.*, **94**, 1031–1049.

Tattersall, G. J., 2016: Infrared thermography: A non-invasive window into thermal physiology. *Comp. Biochem. Physiol. A Mol. Integr. Physiol.*, **202**, 78–98.

Wolfe, R. E., G. Lin, M. Nishihama, K. P. Tewari, J. C. Tilton, and A. R. Isaacman, 2013: Suomi NPP VIIRS prelaunch and on-orbit geometric calibration and characterization. *J. Geophys. Res. Atmos.*, **118**, 11508–11521.

Zhu, Z., C. E. Woodcock, C. Holden, and Z. Yang, 2015: Generating synthetic Landsat images based on all available Landsat data: Predicting Landsat surface reflectance at any given time. *Remote Sens. Environ.*, **162**, 67–83.

Vita

Captain William J. Graff graduated as salutatorian from LaVille High School in Lakeville, Indiana. He entered undergraduate studies at Valparaiso University in Valparaiso, Indiana where he graduated with a Bachelor of Science degree in Meteorology in 2011 as a Christ College Associate. He was commissioned through the Detachment 225 AFROTC at the University of Notre Dame.

His first assignment was to the 21st Operational Weather Squadron, Sembach AFB, Germany as a weather officer in October 2011. In December 2013 he was assigned to the 509th Operations Support Squadron, Whiteman AFB, Missouri where he served as a wing weather officer. While stationed at Whiteman AFB, he attended Squadron Officer School at Maxwell AFB, Alabama in August 2015. In May of 2016, he was awarded a Graduate Certificate in Geographic Information Science (GIS) from Northwest Missouri State University while pursuing a Master of Science (MS) degree. In August 2016, he entered the Graduate School of Engineering and Management, Air Force Institute of Technology. Upon graduation, he will be assigned to Keesler AFB, Mississippi as an instructor.

REPORT DOCUMENTATION PAGE					<i>Form Approved</i> OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.						
1. REPORT DATE (DD-MM-YYYY) 22-03-2018		2. REPORT TYPE Master's Thesis			3. DATES COVERED (From — To) August 2016 - March 2018	
4. TITLE AND SUBTITLE EVALUATION OF A CLOUD DETECTION TECHNIQUE USING SPATIAL AND RADIOMETRIC THRESHOLDS FOR NEAR INFRARED SATELLITE IMAGERY				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER		
				5d. PROJECT NUMBER		
6. AUTHOR(S) Graff, William J., Captain, USAF				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 Hobson Way WPAFB OH 45433-7765					8. PERFORMING ORGANIZATION REPORT NUMBER AFIT-ENP-MS-18-M-083	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) A National Program Office					10. SPONSOR/MONITOR'S ACRONYM(S)	
					11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A: APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.						
13. SUPPLEMENTARY NOTES This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.						
14. ABSTRACT Knowledge of cloud location in near infrared (NIR) imagery is of interest to the meteorological community given the wavelengths greater spatial resolution compared to longwave infrared and its potential nighttime applications. This method consists of an algorithm that can be employed by multiple instrument platforms. It analyzes changes between satellite image radiances and a seasonal synthetic background radiance image. The NIR sensor bands of the Sentinel-2 Multispectral Instrument (MSI) and Suomi National Polar-Orbiting Partnership (Suomi NPP) Visible Infrared Imaging Radiometer Suite (VIIRS) satellite instruments are used to evaluate the performance of a monochromatic change detection technique designed to locate areas of cloud cover. VIIRS imagery is examined for its higher temporal resolution compared to MSI; whereas, MSI imagery is examined for its greater spatial resolution. Background images are constructed either manually or algorithmically using a first-guess image. Observed and background images are differenced based on user-defined radiance, size, and shape thresholds. Pixels that meet these thresholds in the first algorithm are flagged as cloud cover. Output is compared to operational cloud masks that rely on multispectral techniques. Findings indicate that the developed algorithm identifies cloud cover above the specified size threshold well, but optically thin clouds and fresh snow still present limitations.						
15. SUBJECT TERMS ArcGIS, ArcMap, ArcScene, Change Detection, Cloud Detection, GIS, MSI, Sentinel-2, Suomi-NPP, Synthetic Image, POLAR2GRID, VIIRS						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT	c. THIS PAGE			Dr. Robert A. Stenger, AFIT/ENP	
U	U	U	UU	234	19b. TELEPHONE NUMBER (include area code) (937) 490-4884; robert.stenger@altamiracorp.com	