

Characterization of the Spectral Distribution of Hyperspectral
Imagery for Improved Exploitation

by

Ariel Schlamm

B.S. Rochester Institute of Technology, 2006

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

April 14, 2010

Signature of the Author _____

Accepted by _____
Coordinator, Ph.D. Degree Program Date

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

CERTIFICATE OF APPROVAL

Ph.D. DEGREE DISSERTATION

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

Dr. David Messinger, Dissertation Advisor

Dr. William Basener

Dr. Emmett Ientilucci

Dr. David Ross

Dr. Carl Salvaggio

Date

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

Title of Dissertation:

**Characterization of the Spectral Distribution of Hyperspectral Imagery
for Improved Exploitation**

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

Signature _____ Date _____

Characterization of the Spectral Distribution of Hyperspectral Imagery for Improved Exploitation

by

Ariel Schlamm

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

Abstract

Widely used methods of target, anomaly, and change detection when applied to spectral imagery provide less than desirable results due to the complex nature of the data. In the case of hyperspectral data, dimension reduction techniques are employed to reduce the amount of data used in the detection algorithms in order to produce “better” results and/or decreased computation time. This essentially ignores a significant amount of the data collected in k unique spectral bands. Methods presented in this work explore using the distribution of the collected data in the full k dimensions in order to identify regions of interest contained in spatial tiles of the scene. Here, interest is defined as small and large scale manmade activity. The algorithms developed in this research are primarily data driven with a limited number of assumptions. These algorithms will individually be applied to spatial subsets or tiles of the full scene to indicate the amount of interest contained. Each tile is put through a series of tests using the algorithms based on the full distribution of the data in the hyperspace. The scores from each test will be combined in such a way that each tile is labeled as either “interesting” or “not interesting.” This provides a cueing mechanism for image analysts to visually inspect locations within a hyperspectral scene with a high likelihood of containing manmade activity.

Contents

1	INTRODUCTION	21
1.1	Spectral Image Analysis	21
1.2	A New Approach to Spectral Image Analysis	30
1.3	A New Approach to Change Detection	34
1.4	Synopsis	36
2	DEVELOPMENT OF SPECTRAL IMAGE FEATURES	37
2.1	Methodology	37
2.2	Data	39
2.3	Inherent Dimension of Spectral Imagery	43
2.4	Characteristics of High Dimensional Distributions	61
2.5	Anomaly Detection	65
2.6	Classification and Spectral Clustering	67
2.7	Feature Combination for Interest Segmentation	72
3	RESULTS OF SPECTRAL FEATURE ESTIMATION	89
3.1	Dimension Estimation	89
3.2	Gram Matrix Complexity Metrics	105
3.3	Anomaly Detection	112
3.4	Spectral Clustering	116
3.5	Interest Segmentation	119
4	CHANGE DETECTION USING PDDL	139
4.1	Point Density Change Detection	139
5	SUMMARY	153
5.1	Future Work	153

Appendices

Appendix A Additional PDP Feature Maps	161
A.1 Dimension	162
A.2 PDTL	167
A.3 TEP	172
Appendix B Additional Gram Matrix Feature Maps	177
B.1 Number of Endmembers	177
B.2 Peak Volume	182
B.3 Area Under Volume Curve	187
Appendix C Additional Anomaly Feature Maps	193
C.1 Hyperspectral	193
C.2 Multispectral	198
Appendix D Additional Clustering Feature Maps	203
D.1 Hyperspectral	204
D.2 Multispectral	206
References	212

List of Figures

1.1	Visible image of Cooke City showing large and small scale regions of interest.	22
1.2	Notional example of statistical and geometric assumptions placed on hyperspectral data.	23
1.3	Example three-band distributions of manmade and natural materials from real data.	30
1.4	Percentage of information or variability contained in an eigenband as a function of the eigenvalue for four different regions in real data. . .	32
2.1	Overall process model for locating regions of interest in a large area scene.	38
2.2	Visible image of HyMap data flown over Cooke City, MT.	39
2.3	Visible image of AVIRIS hyperspectral data of (a) Rochester, NY, (b) Cuprite, NV, (c) Indian Pines, IN, (d-e) Lunar Lake, NV, and (f) Moffett Field, CA.	40
2.4	RGB images of the Hyperion data of Oaxaca Valley, Mexico.	41
2.5	Visible display of the HYDICE (a) forest radiance hyperspectral and (b) desert radiance data sets.	41
2.6	Visible display of the HYDICE (a) forest radiance hyperspectral and (b) desert radiance data sets.	42
2.7	Quickbird data over the site of the Esperanza forest fire (a) and Indonesia (b).	43
2.8	Notional result from PCA on two different distributions, (a) and (b). Manifold dimension estimation in (c).	44
2.9	Graphical description of box-counting algorithm to determine the capacity dimension.	46
2.10	Illustration of volume-dimension relationship.	47
2.11	Illustration of counting (a) and linear regression of results (b).	50

2.12	Example distributions of synthetic data.	51
2.13	Dimension estimation plots for one-dimensional data with evenly spaced, uniformly distributed, and normally distributed data rotated into 15 dimensions.	52
2.14	(a) Dimension estimation plot for two nonintersecting lines, starting the distance calculation at the center of one line. (b) Dimension estimation plot for two intersecting lines, starting the distance calculation from the intersection point. Also shown are notional two-dimensional projections of the data.	52
2.15	Dimension estimation plots for two-dimensional data with evenly spaced, uniformly distributed, and normally distributed data rotated into 15 dimensions.	53
2.16	Dimension estimation plots for a two-dimensional square (a) and circle (b) rotated into 20 dimensions illustrating the plateau effect. Also shown are notional two-dimensional projections of the data.	54
2.17	Dimension estimation plots of (a) a large circular distribution with three smaller circles surrounding the large circle, (b) a large low point density circle with a small high point density in the center, and (c) a large low point density circle with a small high point density overlapping. Also shown are notional two-dimensional projections of the data.	55
2.18	Dimension estimation plot for seven-dimensional uniformly distributed data with 7000 points rotated into 20 dimensions.	56
2.19	Dimension estimation plots of normally distributed data on a parabolic curve in 15 dimensions.	56
2.20	Notional point density plot illustrating the incline and tail regions. . .	57
2.21	Example plot for box-counting estimation.	58
2.22	Example plot for box-counting estimation of a one-dimensional data set.	58
2.23	Example plot for box-counting estimation of a two-dimensional data set.	59
2.24	Estimated dimensions of full scene hyperspectral data as a function of sensor, image content, and GSD.	60
2.25	Estimated dimension as a function of PDP dimension for HyMap data.	61
2.26	Process used to estimate the dimension, PDDL, and TEP from the PDP.	62

2.27	Examples of incline and tail on a point density plot created from real hyperspectral data of road (a) and forest (b).	62
2.28	Example Gram Matrix volume estimation as a function of number of endmembers plot.	64
2.29	Notional result of TAD on a simple data set.	65
2.30	GUI output displaying the relation between the PDTL and the anomalous pixels in a scene.	66
2.31	Notional description of the Gradient Flow algorithm. The points in (a) are connected to their nearest neighbors in (b). The direction of highest density change are shown in (c).	68
2.32	Notional example of individual communities in a two-dimensional distributions.	69
2.33	Example of the network theory approach to spectral clustering on AVIRIS data of Braddock Bay near Lake Ontario in Western NY.	69
2.34	(a)The scatter plot showing the classes from the gradient flow algorithm and corresponding gradient flow result on an urban HYDICE scene. (b) Corresponding results for the k -means algorithm.(NOTE: The colors/gray scales do not distinguish all 13 major classes. In particular, there is one tree class and one long grass class and they are both displayed as the same color.)	70
2.35	Number of classes in a tile as a function of subjective amount of interest in the tile for (a) Gradient Flow and (b) k -Means.	71
2.36	Number of anomalies contained in a tile as a function of the subjective amount of interest contained in the tile.	73
2.37	Three natural material tiles from HyMap data near Cooke City, MT used for pixel-injection analysis.	74
2.38	Average reflectance spectra of forest (green), grass (cyan), foothills (magenta), and city (black) regions used in this analysis.	74
2.39	Dimension and error of linear piece-wise fit as a function of percentage of manmade pixels included in a forested region.	75
2.40	Example PDP without a clear incline portion illustrating that the estimated slope is too high.	76
2.41	Estimated dimension as a function of percentage of manmade pixels for three natural regions.	76
2.42	Dimension estimation plot tail length as a function of percentage of manmade pixels for three natural regions.	77

2.43	Error of linear piece-wise fit to dimension estimation plot (TEP) as a function of percentage of manmade pixels included in three natural regions.	78
2.44	Five RGB tiles from the HyMap Cooke City, MT collect. Each tile contains approximately 4000 pixels. (a) trees (b) grass (c) grass, road, & trees, (d) a construction area, and (e) a small city.	78
2.45	Complexity plots for the five hyperspectral tiles. (a) trees (b) grass (c) grass, road, & trees, (d) a construction area, and (e) a small city.	79
2.46	RGB image of the Cooke City, MT hyperspectral data cube and the same image notionally divided into 64 individual tiles for processing with the general region of interest highlighted.	80
2.47	Example feature maps (a-e) and an RGB image (f) of the Cooke City, MT hyperspectral data cube and divided into 64 individual tiles for processing. Areas in blue contain known interest.	81
2.48	Combination images from traditional boolean combination in IDRISI.	82
2.49	Addition and multiplication images from scaled metric images.	83
2.50	Results from using supervised classification methods as a way to combine the metric images.	84
2.51	Results from using unsupervised classification methods as a way to combine the metric images.	85
2.52	Results from using the Dempster-Shafer theory classification method as a way to combine the metric images.	86
3.1	RGB display of a HyMAP flight line over Cooke City, MT with regions of interest used in analysis highlighted.	90
3.2	Visible image of example forested region (a) and dimension estimation plots for a 500 pixel (b) and a 30000 pixel (c) forested region.	90
3.3	Bar chart of point density dimension estimates as a function of forest region size.	91
3.4	Dimension estimation plots of two grass samples from the Cooke City data, where (a) has 500 pixels and (b) has 3000 pixels.	92
3.5	Bar chart of point density dimension estimates as a function of grass region size.	92
3.6	Dimension estimation plot for a road sample of approximately 700 pixels.	93
3.7	Visible image of region of secluded houses in Cooke City data (a) and dimension estimation plot (b) for a 500-pixel region.	94

3.8	Bar chart of point density dimension estimates as a function of region size containing secluded houses, trees, grass, and roads.	94
3.9	Visible image of region of Cooke City, MT (a) and dimension estimation plots for a 500 (b) and a 2000 (c) pixel regions.	95
3.10	Visible image of construction site (a) and corresponding dimension estimation plot of the region (b).	96
3.11	Average correlation and PCA dimension estimates for each region of interest.	97
3.12	HyMap data of Cooke City, MT (a) and corresponding PDD feature map for the full hyperspectral scene (b) and resampled multispectral scene (c).	99
3.13	High resolution image of downtown Cooke City, MT.	100
3.14	High resolution image of a construction site outside of Cooke City, MT.	100
3.15	High resolution image of a clearing in the woods outside of Cooke City, MT.	100
3.16	Hyperion data of San Francisco (a) and corresponding PDD feature map for the full hyperspectral scene (b).	101
3.17	HyMap data of Cooke City, MT (a) and corresponding TEP feature map for the full hyperspectral scene (b) and resampled multispectral scene (c)	102
3.18	High resolution image of the area with secluded houses and buildings outside Cooke City, MT.	102
3.19	Hyperion data of San Francisco (a) and corresponding TEP feature map for the full hyperspectral scene (b).	103
3.20	HyMap data of Cooke City, MT (a) and corresponding PDDL feature map for the full hyperspectral scene (b) and resampled multispectral scene (c).	104
3.21	Hyperion data of San Francisco (a) and corresponding PDDL feature map for the full hyperspectral scene (b).	105
3.22	HyMap data of Cooke City, MT (a) and corresponding number of endmembers feature map for the full hyperspectral scene (b) and resampled multispectral scene (c).	106
3.23	RGB image of Hyperion data of San Francisco (a) and number of endmembers where the volume equals zero feature map (b).	107
3.24	HyMap data of Cooke City, MT (a) and corresponding peak volume estimate feature map for the full hyperspectral scene (b) and resampled multispectral scene (c).	108

3.25	RGB image of Hyperion data of San Francisco (a) and peak volume feature map (b).	109
3.26	HyMap data of Cooke City, MT (a) and corresponding area under the volume curve feature map for the full hyperspectral scene (b) and resampled multispectral scene (c).	110
3.27	Peak volume as a function of area under the Gram Matrix showing the faster increase of the area metric.	111
3.28	RGB image of Hyperion data of San Francisco (a) and peak volume feature map (b).	111
3.29	Anomalies detected using TAD (right) on the construction site in Cooke City, MT (left) showing the two-dimensional projection of the anomaly map in the hyperspace (center).	112
3.30	Anomalies detected using TAD (right) on the forest radiance scene (left) showing the two dimensional projection of the anomaly map in the hyperspace (center).	113
3.31	(a) TAD anomaly image of HyMap data at full spectral resolution, (b) number of anomalous pixels feature map, (c) TAD anomaly image of HyMap data at WV2 spectral resolution, and (d) number of anomalous pixels feature map.	114
3.32	TAD result on Hyperion data of San Francisco (a) and number of anomalous pixels (b).	115
3.33	(a) Gradient Flow classification image, (b) ISODATA classification image, and (c) number of spectral clusters feature map for HyMap data of Cooke City, MT.	117
3.34	Gradient Flow output of HyMap data at WV2 spectral resolution (a) and number of spectral clusters feature map (c).	118
3.35	TAD result on Hyperion data of San Francisco (a) and number of anomalous pixels (b).	118
3.36	Interest segmentation map from five spectral features (b) with the interest class turned on in (a) and off in (c).	120
3.37	Interest segmentation map from five spectral features (a), RGB image (b) and 8 spectral features (c).	121
3.38	Interest segmentation map from 5 spectral features for HyMap data resampled to WV2 spectral resolution(a) and RGB image (b).	122
3.39	Interest segmentation maps from five spectral features in two interest segments (a) and three interest segments (c) and visual maps (b) and (d) where true color (RGB) regions are “interesting” and false color (CIR) are “not interesting.”	123

3.40	RGB image with exposed (a) and hidden (b) targets and corresponding interest segmentation maps of HYDICE forest imagery.	124
3.41	RGB image of AVIRIS data of Rochester, NY near Lake Ontario (a) and corresponding interest segmentation map (b).	125
3.42	RGB image of AVIRIS data of Moffett Field, CA (a), corresponding interest segmentation map from one iteration (b), and two iterations (c).	125
3.43	Interest segmentation map (b) of Hyperion data of San Francisco, CA (a).	127
3.44	RGB images (a-c-e) and interest segmentation maps (b-d-f) for three regions of Hyperion imagery of Oaxaca Valley, Mexico.	128
3.45	Interest segmentation map for Hyperion data of Oaxaca Valley, Mexico and correctly identified regions of “interesting” manmade activity.	129
3.46	Interest segmentation maps (a,c,e) from five spectral features in two interest segments and visual maps (b,d,f) where true color (RGB) regions are “interesting” and false color (CIR) are “not interesting.”	130
3.47	RGB image (a) and Interest segmentation map for Quickbird imagery of an near the Esperanza forest fire containing a large agricultural and manmade region.	131
3.48	High resolution tiles corresponding to “interesting” tiles in the interest segmentation map in Figure 3.47.	132
3.49	RGB image (a) and Interest segmentation map for Quickbird imagery of an near the Esperanza forest fire containing a small suburban area and many large areas with no activity.	134
3.50	High resolution tiles corresponding to “interesting” tiles in the interest segmentation map in Figure 3.49.	135
3.51	RGB image (a) and Interest segmentation map for Quickbird imagery of an near the origin of the Esperanza forest fire containing multiple large and small scale regions of interest.	136
3.52	High resolution tiles corresponding to “interesting” tiles in the interest segmentation map in Figure 3.51.	137
4.1	Change detection via PDDL estimation process.	140
4.2	PDPs of two Hydice tiles illustrating the difference in PDDL between the target absent (a) and target present (b) cases.	141
4.3	Change detection result for HyMap image pair with implanted manmade material changes.	143

4.4	Change detection results for high altitude HYDICE data at various threshold values.	143
4.5	Metric value in a tile as a function of the subjective amount of change contained in the tile for Hydice forest data at full spectral resolution and a detection map where the tile brightness is scaled by the metric value.	144
4.6	Thresholded detection maps of HYDICE forest data at full spectral resolution (170 bands).	145
4.7	Thresholded detection maps of Hydice forest data at WorldView2 spectral resolution (8 bands).	146
4.8	Thresholded detection maps of HyMap data at full spectral resolution (126 bands).	147
4.9	Indonesia coastline before (a) and after (b) the tsunami of 2004. Raw PDTL map (c) and scaled map showing only the top 1% of detected changes (d).	149
4.10	Change detection results for a region off the coast on Indonesia with two boats.	150
4.11	Change detection results for a coastal region of Indonesia showing a boat and coast damage.	150
4.12	Change detection results for an area of significant coastal damage in Indonesia.	151
4.13	Change detection results for a region in the ocean off the coast of Indonesia containing a line of buoys suspended in the water in one image missing from the later image.	151
4.14	Change detection results on an inland region of Indonesia showing tsunami damage to a crop region.	152
4.15	Change detection results on an inland region of Indonesia showing a rescue camp (a) and a close up of the blue tents (b).	152
5.1	Estimated dimension for different months during the year.	154
5.2	(a) RBG image of an AVIRIS hyperspectral flightline over Rochester, NY and the particular region of interest used in processing. (b) Dimension estimation plot of the scene and a greyscale map showing that the land pixels are those contained in the inclining portion of the graph. (c) Dimension estimation plot and greyscale map showing that the water pixels are contained in the plateau of the graph. Only those points between the vertical lines in the graph are used to estimate the dimension in (b) and (c).	155

5.3	Fourier transforms of city (left) and forest (right) hyperspectral data.	159
A.1	RGB image (a,c) and PDD feature maps (b,d) of Hydice Forest imagery with exposed and hidden targets.	162
A.2	RGB image of AVIRIS data of Rochester, NY near Lake Ontario (a) and corresponding PDD feature map (b).	162
A.3	RGB image of AVIRIS data of Moffett Field, CA (a) and corresponding PDD feature map (b).	163
A.4	RGB image of sample WV2 data of a forested region (a) and corresponding PDD feature map (b).	163
A.5	RGB image (a) and PDD map (b) for Quickbird imagery near the Esperanza forest fire containing a large agricultural and manmade region.	164
A.6	RGB image (a) and PDD map (b) for Quickbird imagery near the Esperanza forest fire containing a suburban region and multiple small scale regions of interest.	165
A.7	RGB image (a) and PDD map (b) for Quickbird imagery near the origin Esperanza forest fire containing many small and large scale regions of interest.	166
A.8	RGB image (a,c) and PDDL feature maps (b,d) of Hydice Forest imagery with exposed and hidden targets.	167
A.9	RGB image of AVIRIS data of Rochester, NY near Lake Ontario (a) and corresponding PDDL feature map (b).	167
A.10	RGB image of AVIRIS data of Moffett Field, CA (a) and corresponding PDDL feature map (b).	168
A.11	RGB image of sample WV2 data of a forested region (a) and corresponding PDDL feature map (b).	168
A.12	RGB image (a) and PDDL map (b) for Quickbird imagery near the Esperanza forest fire containing a large agricultural and manmade region.	169
A.13	RGB image (a) and PDDL map (b) for Quickbird imagery near the Esperanza forest fire containing a suburban region and multiple small scale regions of interest.	170
A.14	RGB image (a) and PDDL map (b) for Quickbird imagery near the origin Esperanza forest fire containing many small and large scale regions of interest.	171
A.15	RGB image (a,c) and TEP feature maps (b,d) of Hydice Forest imagery with exposed and hidden targets.	172

A.16 RGB image of AVIRIS data of Rochester, NY near Lake Ontario (a) and corresponding TEP feature map (b).	172
A.17 RGB image of AVIRIS data of Moffett Field, CA (a) and corresponding TEP feature map (b).	173
A.18 RGB image of sample WV2 data of a forested region (a) and corresponding TEP feature map (b).	173
A.19 RGB image (a) and TEP map (b) for Quickbird imagery near the Esperanza forest fire containing a large agricultural and manmade region.	174
A.20 RGB image (a) and TEP map (b) for Quickbird imagery near the Esperanza forest fire containing a suburban region and multiple small scale regions of interest.	175
A.21 RGB image (a) and TEP map (b) for Quickbird imagery near the origin Esperanza forest fire containing many small and large scale regions of interest.	176
B.1 RGB image (a,c) and NEM feature maps (b,d) of Hydice Forest imagery with exposed and hidden targets.	178
B.2 RGB image of AVIRIS data of Rochester, NY near Lake Ontario (a) and corresponding NEM feature map (b).	178
B.3 RGB image of sample WV2 data of a forested region (a) and corresponding NEM feature map (b).	179
B.4 RGB image (a) and NEM map (b) for Quickbird imagery near the Esperanza forest fire containing a suburban region and multiple small scale regions of interest.	180
B.5 RGB image (a) and NEM map (b) for Quickbird imagery near the origin Esperanza forest fire containing many small and large scale regions of interest.	181
B.6 RGB image (a,c) and PV feature maps (b,d) of Hydice Forest imagery with exposed and hidden targets.	182
B.7 RGB image of AVIRIS data of Rochester, NY near Lake Ontario (a) and corresponding PV feature map (b).	182
B.8 RGB image of AVIRIS data of Moffett Field, CA (a) and corresponding PV feature map (b).	183
B.9 RGB image of sample WV2 data of a forested region (a) and corresponding PV feature map (b).	183

B.10	RGB image (a) and PV map (b) for Quickbird imagery near the Esperanza forest fire containing a large agricultural and manmade region.	184
B.11	RGB image (a) and PV map (b) for Quickbird imagery near the Esperanza forest fire containing a suburban region and multiple small scale regions of interest.	185
B.12	RGB image (a) and PV map (b) for Quickbird imagery near the origin Esperanza forest fire containing many small and large scale regions of interest.	186
B.13	RGB image (a,c) and AREA feature maps (b,d) of Hydice Forest imagery with exposed and hidden targets.	187
B.14	RGB image of AVIRIS data of Rochester, NY near Lake Ontario (a) and corresponding AREA feature map (b).	187
B.15	RGB image of AVIRIS data of Moffett Field, CA (a) and corresponding AREA feature map (b).	188
B.16	RGB image of sample WV2 data of a forested region (a) and corresponding AREA feature map (b).	188
B.17	RGB image (a) and AREA map (b) for Quickbird imagery near the Esperanza forest fire containing a large agricultural and manmade region.	189
B.18	RGB image (a) and AREA map (b) for Quickbird imagery near the Esperanza forest fire containing a suburban region and multiple small scale regions of interest.	190
B.19	RGB image (a) and AREA map (b) for Quickbird imagery near the origin Esperanza forest fire containing many small and large scale regions of interest.	191
C.1	Anomalies detected using TAD (left), RX (center), on the Cooke City, MT scene (right).	194
C.2	Anomalies detected on the forest radiance scene with targets out in the open (top) using TAD (center) and RX (bottom).	194
C.3	Anomalies detected on the forest radiance scene with targets in shadow (left) using RX (center) and TAD (right).	195
C.4	RGB image of HYDICE data with exposed targets (a), TAD result (b), and corresponding NANOM feature map (c).	195
C.5	RGB image of HYDICE data with hidden targets (a), TAD result(b), and corresponding NANOM feature map (c).	196

C.6	RGB image of AVIRIS data of Rochester, NY near Lake Ontario (a), TAD result (b), and corresponding NANOM feature map (c).	197
C.7	RGB image of AVIRIS data of Moffett Field, CA (a), TAD result (b), and corresponding NANOM feature map (c).	197
C.8	RGB image of sample WV2 data of a forested region (a), TAD result (b) and corresponding PDTL feature map (c).	198
C.9	RGB image (a), TAD result (b), and NANOM map (c) for Quickbird imagery near the Esperanza forest fire containing a large agricultural and manmade region.	199
C.10	RGB image (a), TAD result (b), and NANOM map (b) for Quickbird imagery near the Esperanza forest fire containing a suburban region and multiple small scale regions of interest.	200
C.11	RGB image (a), TAD result (b), and NANOM map (b) for Quickbird imagery near the origin Esperanza forest fire containing many small and large scale regions of interest.	201
D.1	RGB image of HYDICE data with exposed targets (a), Gradient Flow output (b), and corresponding NCLUST feature map (c).	204
D.2	RGB image of HYDICE data with hidden targets (a), Gradient Flow output (b), and corresponding NCLUST feature map (c).	204
D.3	RGB image of AVIRIS data of Rochester, NY near Lake Ontario (a), Gradient Flow output (b), and corresponding NCLUST feature map (c).	205
D.4	RGB image of AVIRIS data of Moffett Field, CA (a), Gradient Flow output (b), and corresponding NCLUST feature map (c).	205
D.5	RGB image of sample WV2 data of a forested region (a), Gradient Flow result (b) and corresponding NCLUST feature map (c).	206
D.6	RGB image (a), Gradient Flow output (b), and NCLUST map (c) for Quickbird imagery near the Esperanza forest fire containing a large agricultural and manmade region.	207
D.7	RGB image (a), ISODATA result (b), and NCLUST map (b) for Quickbird imagery near the Esperanza forest fire containing a suburban region and multiple small scale regions of interest.	208
D.8	RGB image (a), ISODATA result (b), and NCLUST map (b) for Quickbird imagery near the origin Esperanza forest fire containing many small and large scale regions of interest.	209

List of Tables

2.1	Dimension and Estimation	45
2.2	Proposed List of Features.	72

For my mom, who has been there since day one.

Chapter 1

INTRODUCTION

1.1 Spectral Image Analysis

The collection and use of spectral imagery has grown over the years as the modality inherently measures spatial and spectral information over the full scene. Existing hyperspectral processing algorithms are not consistently reliable due to the complexity of changing phenomenology from scene to scene or within a single wide area scene.[22, 67] For this reason, much of the burden lies on the analyst to visually inspect the imagery to locate true targets, anomalies, changes, or material classes and identify false detections within the scene. However, it is not feasible for an analyst to look at every band of every scene in any reasonable amount of time. Typically used methods for any detection scheme depend on placing statistical or geometric assumptions on the distribution of the data. These are generally not robust to complex or cluttered scenes and the quality of the results changes drastically between scenes. There is a need for semi-autonomous exploitation methodologies to use the information captured by a hyperspectral sensor that are robust to these changing scene contents and take advantage of the entire collected data set. This dissertation presents methods that utilize the full distribution of hyperspectral data in the measured k spectral dimensions to identify regions of manmade activity within a wide area scene. The methods developed will not be based on known or measured spectral signatures of materials. This is a new application as the methods developed will be applied to spatial tiles of the scene in order to aid an analyst with identifying and inspecting areas of interest. Results of the processing will be used to alert an analyst to areas that require visual inspection or further data collection.

For the purpose of this research, regions of interests are defined as regions containing small and large scale manmade activity. Figure 1.1 shows the visual map of

a hyperspectral collect over Cooke City, MT taken during the summer of 2006 with the HyMap (VIS/NIR/SWIR) sensor which were under a program sponsored by AFRL. The ground sample distance (GSD) of this image is approximately 2m. This scene contains large portions of natural materials, including forest, grass, and rocky terrain, that should be classified as “uninteresting.” It also contains a small urban area, farms, and scattered individual houses. Preliminary investigations have shown that there are examples of regions of large and small manmade activity contained within the scene. A region of large scale activity is an area consisting of several “interesting” pixels that may or may not be obvious from broad visual inspection of the scene. In the Cooke City image, there is a region next to the road that has been cleared of natural materials to make room for trucks and dumpsters, shown in Figure 1.1(right). A region of small scale activity is just one or a few “interesting” pixels. Preliminary anomaly detection methods were tested on the Cooke City data. One method, Topological Anomaly Detection [11] (Section 2.5), found a very anomalous pixel in the middle of a field, shown in Figure 1.1(left). This was visually confirmed with high resolution data to be a storage tank or tower in the middle of the field.




Figure 1.1: Visible image of Cooke City showing large and small scale regions of interest.

The four basic tasks for spectral image analysis are target detection, anomaly detection, change detection, and spectral clustering/classification. In order to perform any of these tasks in an automated fashion on hyperspectral data, assumptions about what the data “look like” in the hyperspace must be made. These assumptions are typically either statistical or geometric assumptions about the data and are shown in Figure 1.2. Statistical methods describe the distribution of the data in the hyperspace and are often multivariate Gaussian and T-Elliptical assumptions.[35] These can be used in classification, target, or anomaly detection schemes. Geometric assumptions placed on the data rely on defining basis vectors, which may or may not be orthogonal, to describe how the data lie in the space. They are commonly divided into two categories: vector subspace and linear mixture models. If the data lie in a vector subspace, then a set of orthogonal basis vectors can be found to describe the entire data set where the number of basis vectors is less than the original number of bands measured. Linear mixture models (LMM) are based on the assumption that the data lie in a convex hull.[55] The data can be described by the corners of the simplex or endmembers (or non-orthogonal basis vectors) and each point contained in the convex hull can be expressed as a linear combination of the endmembers. The number of endmembers is typically less than the number of dimensions measured. Linear unmixing uses the LMM to determine the fractional amount of each endmember in each pixel.[60]

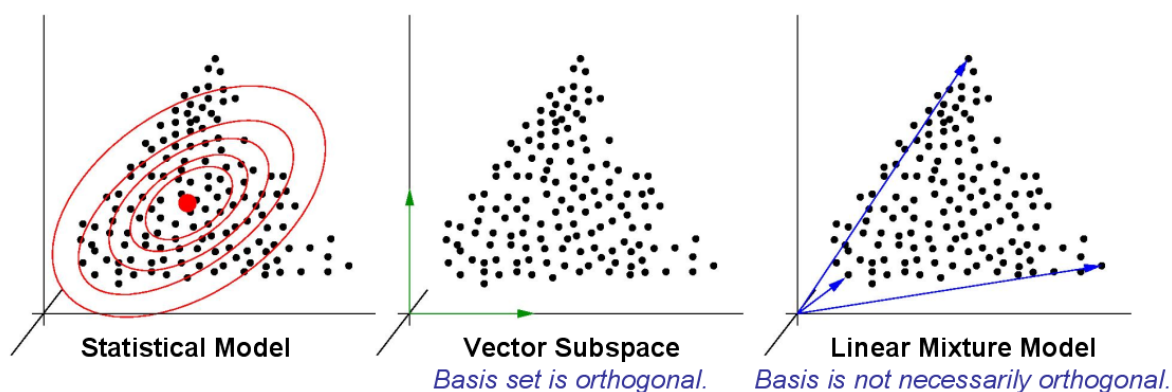


Figure 1.2: Notional example of statistical and geometric assumptions placed on hyperspectral data.

In the broadest sense, the term dimension describes the number of degrees of freedom within a space. With reference to hyperspectral imagery, dimension can refer to the number of unique spectral bands, k , collected. Hyperspectral imagery typically has hundreds of bands of spectral data; therefore, in some sense, each image has hundreds of dimensions. The inherent dimensionality of a spectral image is also a function of the subject being measured. It is commonly found that the dimension of the vector space in which the data were collected is higher than the inherent dimension of the data. [28] Therefore, it is assumed that the inherent dimension of the subject(s) within hyperspectral images is less than the number of bands measured.

Reducing the collected dimension, n , of the spectral image to the inherent dimension, d , where $d < n$, of a spectral image is commonly used as a means to reduce processing time and improve results of further analysis by transforming the data into d unique and uncorrelated dimensions. This relies on the assumption that the entire spectral image can lie in the same d -dimensional hyperspace. Dimension reduction of hyperspectral imagery is commonly done with Principal Component Analysis (PCA), Minimum Noise Fraction (MNF), Singular Value Decomposition (SVD), Virtual Dimensionality, or similar methods. [55, 17] PCA performs eigenvalue decomposition on the spectral covariance matrix in order to orthogonalize and decorrelate the hyperspectral data. The eigenvalues are related to the variance of the data in the corresponding eigenvector's dimension. Typically, to reduce the dimensionality of the data, a threshold is set on the total percentage of variability in the eigenvalues.[55] The data corresponding to the largest eigenvalues are retained and the less varying data are discarded from further processing. This results in a smaller data set to pass on to a classification or detection algorithm, which will run faster and in some cases provide better results. Each band in the new, reduced data set is a linear combination of all the originally measured spectral bands. MNF is also used to reduce the spectral dimension of hyperspectral data. In MNF, the sensor noise is adjusted for in the transform. PCA and MNF both use the spectral covariance matrix to transform the data into a decorrelated space. This process is based on the assumption that the data are normally distributed in the space. If the data are not normally distributed, the covariance between two bands is not an accurate measure. However, going through the process of rotating the data into an uncorrelated space may lead to useful, though not probabilistic, information about the data. SVD is also similar to PCA as it is used for dimension reduction. However, SVD can be done on the hyperspectral data itself, rather than the spectral covariance matrix. The singular values and basis vectors, instead of eigenvalues and eigenvectors, are calculated. For this reason, SVD is a process based on geometric

assumptions of the data.

Anomaly detection in hyperspectral imagery is the process of locating pixels that are spectrally different from the background. In order to find an anomalous pixel, there are two requirements. First, the background must be well defined; without knowing what is a typical pixel in the scene, there cannot be a pixel that is different from the background. Second, a distance measure from a pixel to the background or background model must be established. There are a variety of ways anomaly detection is performed in the hyperspace, but they primarily assume either a geometric or statistical description of the background and distance metric. For example, if LMM is used, the background is assumed to be in a convex hull defined by a set of endmembers. A point outside the convex hull is an anomaly.

A distance metric is often used to measure how far a certain point is from the background in order to decide if it is anomalous. The simplest metric is

$$d_e(\mathbf{x}, \mathbf{y}) = \sqrt{\sum_{i=0}^n (x_i - y_i)^2} \quad (1.1)$$

where d_e is the Euclidean distance between the two points in the space; it can be used in either a statistical or geometric method. The Euclidean distance is an ambiguous measure in many circumstances as it cannot differentiate between spectral and brightness differences. The spectral angle,

$$\theta = \cos^{-1} [\mathbf{x}_u \bullet \mathbf{y}_u], \quad (1.2)$$

where θ is the multidimensional angle between two unit vectors \mathbf{x}_u and \mathbf{y}_u in the hyperspace, can be used as a geometric distance metric. It measures the difference in spectral character of a pixel and is therefore independent of the magnitude or brightness of a pixel.[55] The brightness difference can be measured by subtracting the lengths of the pixel vectors; however, this is typically less useful in hyperspectral analysis. In both of these cases, the point \mathbf{y} can be a single pixel representative of the background distribution, for example the mean, median, or mode. If the background is assumed to be a Gaussian distribution, the Mahalanobis distance, sometimes called the statistical distance, can be used as a metric. This distance is defined as

$$d_M = (\mathbf{x} - \mathbf{m})^T \Sigma^{-1} (\mathbf{x} - \mathbf{m}), \quad (1.3)$$

where \mathbf{x} is a pixel vector and \mathbf{m} and Σ are the mean and covariance of a Gaussian distribution. This is the statistical distance from a point to a Gaussian distribution. There are also metrics that attempt to combine the angular distance, euclidean

distance, and correlation into a single measure termed spectral similarity, which may improve algorithm results.[64]

The RX algorithm[42] is a statistical detector with a constant false alarm rate (CFAR), which is based on the assumption that the background is normally distributed and that the anomalies have some additive spectral feature relative to the background. [55] The RX algorithm computes the Mahalanobis distance from each pixel to the the spectral distribution of the local spatial neighborhood. It is typically calculated using the local mean and covariance around the pixel of interest; to reduce computation time, sometimes the full scene covariance or both mean and covariance are used in the calculations. The RX algorithm is a generalized likelihood ratio test (GLRT). A threshold is chosen on the statistic to have a particular false alarm rate given some assumptions about the test statistic distribution.

There have been many modifications of the RX algorithm presented because it is a CFAR detector. One such modification is the Kernel RX algorithm.[30] In this version, the data is transformed into a high dimensional feature space through a nonlinear algorithm; in this new feature space, the background data should be more accurately modeled by a Gaussian distribution and therefore produce better results. The nonlinear transformation is not explicitly done; instead, the RX algorithm is applied to kernelized data where both the data and the detector are “mapped” into the higher dimensional space.

Data driven methods that do not rely on assumptions about the data are desirable in anomaly detection. Banerjee, *et al.* (2006), use the support vector data description (SVDD), which is a particular kerneling method to describe the support of the data. Unlike the Kernel RX algorithm, the SVDD method does not compute statistics on the data, but a decision rule is determined to identify anomalies.[30] Similar to the RX algorithm, a detection threshold is chosen for the normalized SVDD decision rule in order to assign pixels as background or anomalous. This is not exactly a CFAR detector; the threshold chosen is associated with an acceptable range of false alarms.

If the material of interest is known, then instead of using anomaly detection routines to locate all anomalous pixels, target detection methods can be employed to locate pixels of the particular material in a hyperspectral image. Orthogonal subspace projection (OSP) is a target detector that minimizes the background contribution through a projection onto the endmembers, and then projects the pixel onto the target to measure the similarity.[23] Statistical target detection algorithms such as the spectral matched filter (SMF) [48] use the mean and covariance of the distribution to determine whether the target is present in the data. Other statistical target detectors are variations of SMF and include constrained energy minimization

(CEM) [55], and adaptive coherence estimator (ACE) [46]. This, and other target detection methods or distance metrics, can be used to measure the similarity of a pixel to a particular target. Many target detection methods can be modified in order to measure the distance away from the background instead of the similarity to a particular spectrum for use as an anomaly detector.

Classification of hyperspectral images is the process used to segment the spatial data based on the spectral signature into a variety of material classes. There are two main methods for spectral clustering or classification: supervised and unsupervised. [55] Supervised classification requires training: the user must choose sample regions of the desired material classes from the original scene. Once the sample regions are chosen and characterized, a distance metric is used in order to assign the rest of the pixels into a particular class. The spectral angle can be calculated from the pixel to each of the class means and assigned to be in the class with the smallest spectral angle. This measure is useful because two pixels of the same material, but with different brightnesses, will likely be put into the same class. Another metric, minimum distance to the mean (MDM), calculates the Euclidean distance from a point to the mean of each class, and assigns it to the class with the smallest distance. Gaussian maximum likelihood classification uses the Mahalanobis distance to the mean as a discriminant metric. In this method, the training data are used to describe each class as a Gaussian distribution with a mean and covariance. Supervised classification is limited by the user. The user must choose the correct number of classes and select enough sample data so each class is well defined. If the training data do not represent every possible pixel for a particular class, that type of pixel will be misclassified. The user typically chooses classes based on visual interpretation of the scene and materials in the scene. For example, a person might want a class for buildings, but a building is typically made of a variety of materials that are spectrally different. This type of class may be more accurately classified into individual material clusters. The training data, and thus the classification results, are subjective and dependent on the user.

Unsupervised classification does not require the user to select training data or the exact number of classes to be used. Typically, the user will give a “best guess” at the number of classes for the algorithm to use as a starting point, but an unsupervised algorithm can add and/or subtract from that number. K -means and Isodata are the two most common algorithms for unsupervised classification. The k -means algorithm chooses k random pixels in the hyperspace to be class means. Every other pixel is assigned to one of these k classes based on a distance metric, usually MDM. The means of these classes are calculated and the process repeats until the mean converges. Isodata is a modified version of k -means that can reduce

or increase the number of classes with each iteration. If too few pixels are assigned to a class, the class is removed and the process continues. Two classes whose means are very close can be combined into one class. When the variance within a class is too large, this class can be split into two classes. These algorithms may not produce a good classification map to the human eye when the classes do not correspond to the material classes of which we think. Unsupervised classification is usually used as a method to determine the best number of classes and where to choose the training data for supervised classification.

If the data are believed to lie in a convex hull, linear unmixing can be used to find the proportions of each endmember in the pixels. It can be modified to be a classifier by thresholding the proportion, or assigning a pixel to the endmember of higher proportion. This method can be either supervised or unsupervised, depending on whether the endmembers are chosen by a user or automatically found with an algorithm such as MaxD or N-FINDR.[55] The drawbacks of this method all come from the assumption that every pixel is a linear combination of the endmembers when this is not explicitly true.

Wide area search methods are commonly employed when the hyperspectral data are recorded over a large physical area. These scenes are often used for the detection of unknown targets, activities, and events. Within the physical area of the scene, there can be a wide variety of land cover environments (urban, suburban, rural) and materials. Search requires finding all possible regions of interest at the expense of finding more false alarms.

The Wide Area Reconnaissance Hyperspectral Overhead Real-time Surveillance Experiment (WAR HORSE) was built by the Naval Research Laboratory in order to function as a real time detection and cueing operation. The imaging system consists of two line scanners, a visible hyperspectral sensor, and a high resolution visible sensor. A subspace RX detector is used as the data are collected. When an anomaly is found in the hyperspectral data, the high resolution imager is cued and a high resolution image is taken over the identified anomaly. [63] Other than the WAR HORSE system, hyperspectral imagery is not typically used for wide area search.

The types of “interest” possible within a wide area scene are diverse and difficult to model.[15] As Carlotto points out, targets and regions of interest do not occur as single moments in time. [16] The location of a building, for example, begins as a clearing of the vegetation and arrival of vehicles. Construction of the building begins, and after time passes, the building is finally finished. All of these stages would need to be accurately taken into account for an accurate target model. Automated target detection routines depend on accurate models to be successful. Carlotto uses

multitemporal images in conjunction with pattern recognition and change detection in order to locate changes that match a particular pattern, like the growth of vegetation or construction of a building, in wide area imagery. For this reason, algorithms that do not depend on prior knowledge of what is “interesting” are useful in wide area search applications.

Users of PCA and other dimension reduction techniques tend to throw away potentially useful information for anomaly detection and classification. Wide area search methods are required to locate a variety of target sizes and classes. Removing the data that hold “less information” may result in more false positives and false negatives in anomaly detection and poor classification. An accurate wide area search methodology must take advantage of all the data collected, aside from low SNR channels, in order to be robust and utilize the full distribution of hyperspectral data.

Previous spectral image complexity has not been explicitly measured. Spectral image quality and utility metrics have historically been more desired, but are related to the concept of complexity.[38] It is likely more difficult to extract useful information from a complex image but a complex image also more likely contains information that is useful to the task of spectral image analysis. A higher quality image is also generally more useful, but quality is task specific. A high resolution image, such as those from WorldView2 or HyMap, may not be as useful for locating large regions of crop expansion, just as a low resolution image, from Landsat or Hyperion, is not useful for locating small targets.

The NIIRS scale is a functional expression developed for rating the quality and interpretability of an image.[27] The concept is expanded in Shen (2003), which describes a spectral quality metric for object and anomaly detection in reflective hyperspectral imagery based on the spectral resolution, spatial resolution, and signal-to-noise ratio. Kerekes, *et al.*, (2005) compare a selection of spectral image quality and utility metrics to the General Spectral Utility Metric for use in target detection.[59] Target implantation is used to determine the utility of an image for target detection with respect to specific targets by Stefanou & Kerekes (2009). The results of all of these methods are dependent on the complexity of a spectral image, but until recently the complexity of an image has not been directly measured from the data. The characteristics of spectral distributions in the hyperspace described above indicate that the geometry of the distributions of spectral imagery containing manmade materials is different than imagery only containing natural materials. Dimension, volume, and other geometric measures of the spectral distribution can indicate the degree of scene complexity.

1.2 A New Approach to Spectral Image Analysis

Why is the full distribution of data being used? A spectral image rarely consists of a single type of landcover; often the subject matter is a combination of natural materials, including forest, grass, water, and manmade materials, including buildings and roads. The individual spectra of these materials are not the same. It follows that the distributions of these materials do not lay in the same dimensional subset of the hyperspace. Consider the plots in Figure 1.3 showing three-band projections of forest and city pixel data in the visible and IR wavelengths. The forest data are tightly clustered and appear to fall in a fairly uniform shape in the hyperspace. Clustered near the origin are the shadow pixels. In practice, it is assumed that natural materials are Gaussian or elliptical T-distributed[35] throughout the entire hyperspace. The manmade materials do not show this tendency. Consider the

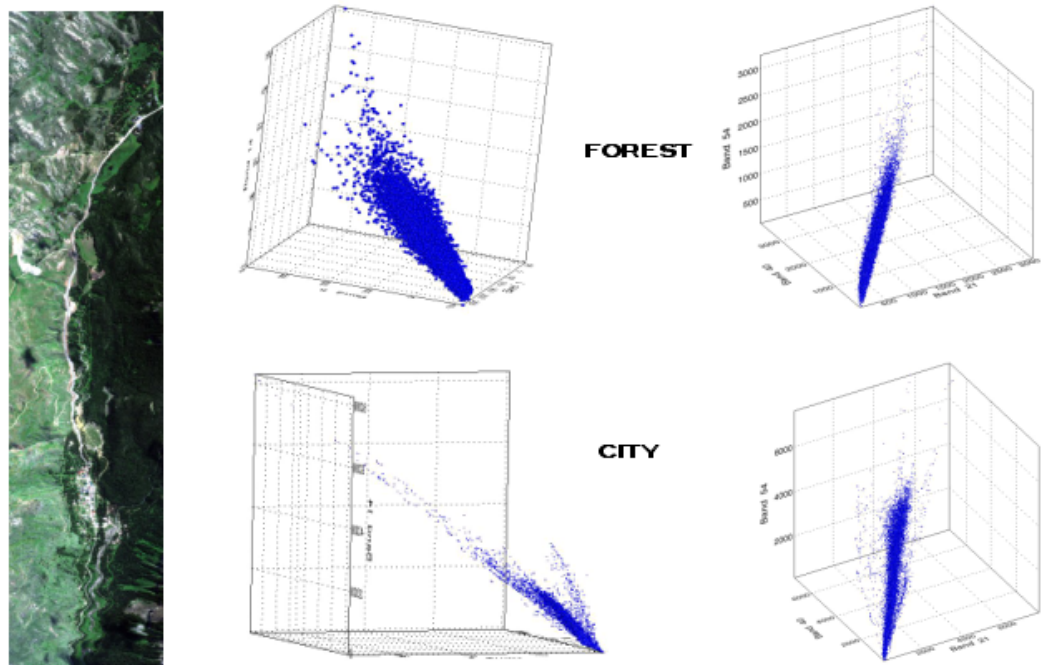


Figure 1.3: Example three-band distributions of manmade and natural materials from real data.

construction of manmade materials; the process is controlled in such a way that each time a material is manufactured, it has similar material properties. Of course, there is some variability to these processes and illumination during collection. Collectively, manmade material distributions do not look like natural material distributions in the hyperspace. Instead, particular areas of the distribution are more populated with pixels than others; each area may be a different material. Visualization of this distribution in 3D would show the center to be virtually empty. Other features seen in manmade material distributions are the long, somewhat linear features protruding from the bulk of the distribution in many directions. Reducing the dimension of the image can result in the loss of some of this information. It is clear from Figure 1.3 that manmade distributions are not Gaussian and are not contained in a well-defined convex hull or linear subspace of the hyperspace.

The eigenvalues from a PCA of four different regions of the Cooke City scene (forest, grass, city, construction) were calculated and used to discriminate features of the different region distributions in the hyperspace. Figure 1.4 shows a measure of the relative amount of information in the orthogonal set by plotting the cumulative sum of the first i th eigenvalues over the total sum *vs.* the eigenvalue number. In all cases, the majority of information is contained in the first 20 bands. This is a typical result as most hyperspectral images are reduced to approximately 20 dimensions in practice. After this point, the curves diverge. The natural materials fall off slowly because they are more well distributed throughout the entire data cloud, and therefore better represented by a multivariate normal distribution as they have variability in every direction defined by the eigenvectors. It is for this reason that there are methods that try to distinguish manmade materials from natural by assuming the natural materials are the background that can be well modeled by normal distributions. The city falls off quickly, and the construction site falls off between the two. This indicates that the distributions of manmade and natural materials are different in the bands corresponding to the lower eigenvalues and that an area with both material types (construction) falls in between. Even though PCA is not commonly used to explore the distribution of the data in the hyperspace, it is clear from this example that the distributions of manmade and natural materials manifest differently in the hyperspace.

As the full d -dimensional hyperspace of recorded data is used in the variety of tests, the algorithms may be applied to hyperspectral image data in either the reflectance or radiance domain. As many methods of calibrating recorded radiance data to true reflectance data often involve only a linear transform, for example the Empirical Line Method [55], the actual distribution of the data will remain unaltered in shape. Only the units and scale of the distribution are affected. For this reason,

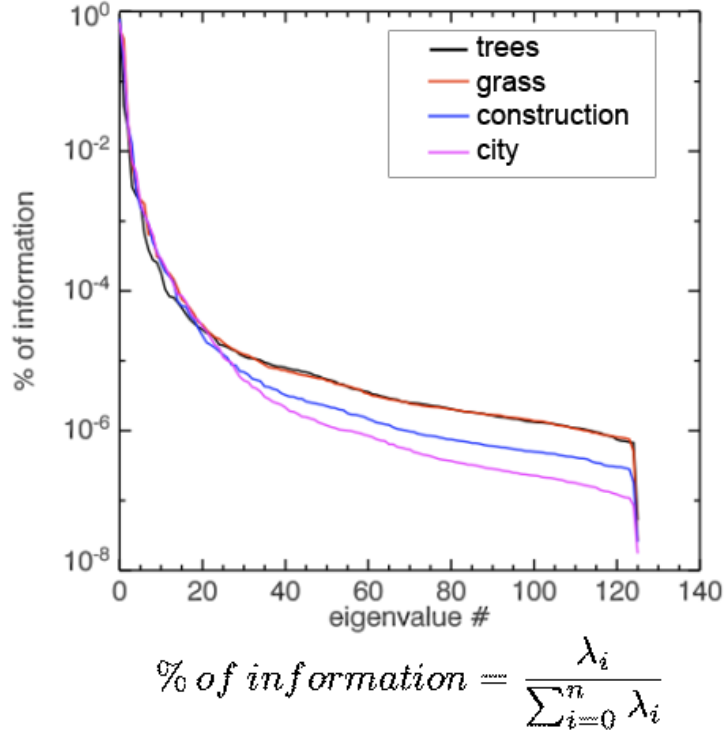


Figure 1.4: Percentage of information or variability contained in an eigenband as a function of the eigenvalue for four different regions in real data.

radiance or reflectance data may be used for this research.

Using the full distribution of data can lead to a large number of computations and therefore long processing times. As we do not need to spatially locate each interesting pixel but general regions likely containing interest and large uninteresting regions, spatially tiling the hyperspectral image is one way to avoid the computational drawbacks to processing the full scene at once. A variety of tests indicating the likely presence of “interest” will be applied to hyperspectral spatial tiles of the original scene. The results of these tests will be combined so that each spatial tile is labeled as containing or not containing “interest.” A map of the original hyperspectral scene will be created using these labels to indicate areas that are “interesting” and should be given priority for further investigation and also to indicate areas that are not interesting and should not be looked at. In this way, it will be a cueing mechanism to analysts: a way to prioritize where their time should be spent looking or if higher resolution data should be taken over smaller regions of the scene.

Tiling the image spatially provides many advantages, beyond simpler and faster processing. Applying tests and algorithms to individual tiles may make the algorithms more robust. Small spatial tiles should have less complex distributions than the full scene as well as fewer materials and less material variability. This translates to fewer false alarms both in each tile and overall. Large portions of the image will likely be declared “uninteresting,” eliminating all possibilities of false alarms in these areas. Each tile is independent from the neighboring tiles during processing, so parallel computing can be used to further decrease processing time.

As previously stated, a large scale region of interest is considered to be made of several “interesting” pixels. This corresponds to a larger structure that likely has some spatial influence. Because these should have a substantial number of pixels, a large region of interest will alter the distribution of the data in the full measured hyperspace. A small scale region of interest consists of just a few, a single, or possibly subpixel areas of manmade activity. These regions are spatially small enough that visibility from viewing the hyperspectral scene would be nearly impossible, depending on spatial resolution and object size. They also consist of so few pixels that their presence in the spectral distribution would not significantly alter the distribution of the full image data. Locally, however, a small number of pixels should be enough to alter the distribution in some way.

Existing spectral anomaly detection methods should be more robust to smaller, less complex, and less cluttered distributions. For example, the RX algorithm [55, 30] is more robust when computed with local statistics rather than image-wide statistics. By spatially tiling the image, less complicated statistics will be computed as there is less variability within a small image tile. As it has been shown that statistical models are not very accurate for complex hyperspectral data, new anomaly detection methods based on the graph and topology of the data will be developed and introduced.

The primary issue being faced in the location of small scale regions of interest is as follows. Imagine a grassy field with a few individual subpixels of boulders in one spatial tile. The large scale manmade activity tests will likely not mark the tiles as “interesting” because the distribution of the tile would not be greatly affected by the presence of subpixel boulders. A thorough anomaly detection routine would mark these boulders as anomalies (e.g. in a field of grass, they are anomalous). A few boulders are naturally occurring and not “interesting” by our definition. However, rock and spectrally rock-like materials may also be used in building new structures, which are “interesting” by the provided definition. This issue will likely always require some human input to decide what is actually occurring.

In the context of this research, manmade activity is considered “interesting,”

however, not all manmade activity is equally interesting. For example, a house or small neighborhood may appear in our described tests as regions of manmade activity, even if nothing in the region has changed for many years, and the region is therefore uninteresting. However, if construction was underway around a single house, this might be considered interesting. In one scene, what is considered a small scale region of interest may not be temporally limited to a few pixels in size. Before a large construction site exists, a single truck or dumpster is brought in.

1.3 A New Approach to Change Detection

Many of the metrics that are developed for this research attempt to quantify the “complexity” of a scene by directly measuring features in the data that are quantifiably different between purely natural distributions and those including manmade distributions. Because of this fact, some of these metrics are also useful in the field of large area change detection, specifically with multispectral data. The Digital Globe company currently has three multispectral satellites, including WorldView2, launched in 2009, which can revisit the same location nearly once per day. These sensors are expanding the fields of large area change detection.

Change detection in remote sensing is the process of identifying the differences between two or more images. In the simplest case, two perfectly registered, single band images can be subtracted pixel by pixel. High values in the result correspond to changes and low values to static regions. In remote sensing, this method is not practical in most situations, but has laid the basis for change detection in spectral imagery. In general, two or more images are registered and the difference between the two images is estimated in some fashion in order to detect change between spectral images. All change detection metrics require a method to combine the spectral information from all bands and are sensitive to some amount of misregistration error.

A typical approach to spectral image change detection involves some type of prediction used to bring images collected at different times and under different conditions into the same domain. For this to work, one must assume that the reference image (or images) contains nothing of interest while the subsequent images may contain the change. From this one can predict what the “no interest” image looks like at a later time, but the analyst must have knowledge of all the contents of the original image, including any targets. Prediction is done to bring the two images into the same domain and therefore suppress the background, or constant areas, between the two images before some type of subtraction is done to accentuate the

differences. A thorough review of statistical anomalous change detection algorithms commonly applied to spectral imagery can be found in Theiler (2008).

Schaum and Stocker (2004) predict what one image should look like at the later time by using methods based on covariance equalization followed by the chronochrome algorithm or matched change detection to determine whether a change occurred. Chronochrome and matched change detection are very sensitive to misregistration errors, but using covariance equalization as the predictor reduces the number of false alarms caused by misregistration error.[49, 47] Meola and Eismann (2009) extend standard hyperspectral change detection routines to include multiple reference scenes when predicting what the detection image may look like. The reference images used were taken under varying illumination conditions in order to better describe what the detection image could look like if an illumination change took place. This is important because it may help distinguish illumination changes from interesting changes in the detection process.[36]

Carlotto (1997) approaches change detection as if it were pattern recognition in multitemporal imagery, but with single spectral band images.[16] In this approach, a nonlinear predictor is used to measure the changes between all available pairs of images. If a particular pattern of change is known to be interesting, for example a continuous increase in brightness or reflectance over time, one can look for that pattern of change in all available patterns. Otherwise, one can explore all occurring change patterns and decide which are interesting.

The methods developed are primarily applied to hyperspectral imagery in the reflective region of the spectrum. However, once a working methodology is created, the methodology is also applied to multispectral and thermal data. The phenomenology of the scene will change, but the governing assumption that regions of manmade activity will manifest differently from natural materials in the hyperspace is still valid. Natural materials are often at the scene ambient temperature. If calibrated data are used, vegetation can easily be sorted out as they have high emissivities due to high water content (likewise, water can easily be found). The manmade materials will have temperature contrast with the background, as they are made of a variety of materials: a central air controlled building, a hot running car, a metal dumpster. The materials will also all have different emissivity spectra, and thermal data typically produce a more complex distribution in the hyperspace.

1.4 Synopsis

The remainder of this thesis describes the methodology for approaching this problem and evaluating the results when applied to hyperspectral images in order to distinguish large and small scale manmade regions of “interest” from the natural materials in the image. Methods for determining the dimension of the distribution of the data and classification schemes are used to locate large scale regions of interest. For these tests, the dimension or number of classes for the entire scene may be computed for comparison. These methods should work as manmade materials should influence both the dimension, shape, and complexity of the distribution as well as the number of spectral clusters, especially in a cluttered or large region. Anomaly detection schemes will be developed in order to locate both the large and small scale regions of interest and count the number of anomalous pixels within a tile. However, locating the interesting areas of a spectral image is not the only goal of this research. In order to develop data driven algorithms that successfully exploit the spectral data, the nature of spectral data must be better understood. Improving the current understanding of the distribution of spectral imagery in the native hyperspace is a main objective of this research.

Chapter 2

DEVELOPMENT OF SPECTRAL IMAGE FEATURES

2.1 Methodology

Established methods used on hyperspectral imagery to perform data reduction, classification, and target and anomaly detection are used on the condition of simplifying assumptions applied to hyperspectral imagery. The assumptions of multivariate normality or linear subspace do not necessarily hold when applied to cluttered or wide area scenes. Methods outside of statistical and linear mathematics, including topology, graph theory, and fractal mathematics, are investigated for use in hyperspectral image analysis. Data driven algorithms that do not depend on as many strict (and possibly inaccurate) assumptions are developed for spectral image analysis.

New methods, based on the k dimensional “graph” of the data in the hyperspace, the topology of the data, the density, and the geometry of the data, are explored in order to characterize and quantify the differences between natural and manmade distributions and to locate regions of interest in a spectral image. These methods do not assume that the data lie in a convex hull, the distribution is multivariate normal, or there are unique endmembers or pure spectral pixels in the scene. These methods have limited assumptions, primarily that the manmade activity manifests differently in the hyperspace in a quantifiable manner.

The goal of this project is to develop a methodology for a spectral analysis system to alert an analyst to regions of interest within a wide area scene. The hyperspectral data are spatially tiled for individual processing. A series of tests based on alternative, data driven methods are developed. Each of these tests provide a spectral feature measure that indicates the presence of interest in a tile. The scores

from each test are combined so each tile is assigned an identifier corresponding to the likely presence of interest. The overall process is outlined in Figure 2.1. The analyst can then prioritize how to visually inspect the scene in the areas with high interest or request higher resolution data in a particular area and ignore large regions containing only natural and uninteresting materials.

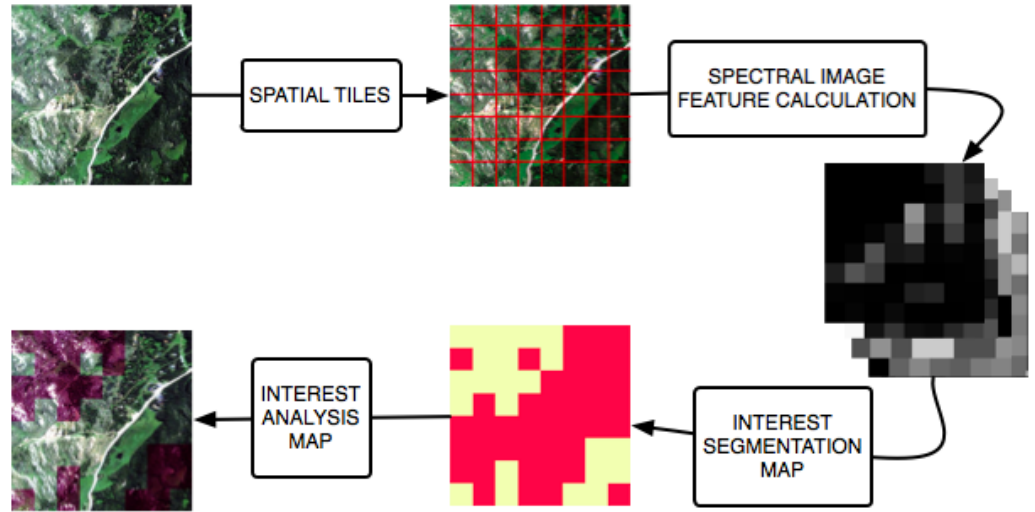


Figure 2.1: Overall process model for locating regions of interest in a large area scene.

The tests are developed for finding both large and small regions of interest or activity. Large region tests include measures based on the point density of the tile and number of material classes in the tile. The local volume of the distribution in the hyperspace is also investigated as a way to estimate the size and complexity of the distribution. Anomaly detection is the primary test for locating small regions of interest. Some of these tests prove useful in the task of change detection within temporal data. The performance and robustness of each individual test and the entire algorithmic process is verified by application to a large collection of imagery and in some cases using the target implantation method.[13] As each test is developed, it is applied to many tiles containing a variety of materials and land cover combinations.

2.2 Data

A variety of spectral imagery will be used in different parts of this analysis. Approximately 45 images are assembled for this purpose. One collection of data used for this research was taken by the HyMap sensor flown over the Cooke City, MT area on July 4, 2006.[61, 2] A visible image of the data is shown in Figure 2.2. HyMap is a commercially flown airborne hyperspectral sensor. These data have 126 usable bands (after the atmospheric absorption bands are removed) between 450nm and 2500nm with 2.5m ground sample distances(GSD). There are seven different flightlines of Cooke City available in both radiance and reflectance units. The scene is predominately natural material, including forest, farmland, grass, foothills, rock, and dirt roads. The manmade material in the scene includes downtown Cooke City, a residential neighborhood, and a construction site. These sections contain buildings, vehicles, and paved roads. Each flightline contains several known fabric targets and automobiles in different locations. These scenes are useful for the task of large area search for regions of interest. These data are used for dimension estimation, interest segmentation, and change detection.



Figure 2.2: Visible image of HyMap data flown over Cooke City, MT.

AVIRIS is the NASA operated airborne imaging spectrometer that operates in the visible and near infrared portion of the spectrum.[24] This sensor provides 224 spectral bands with a ground resolution of approximately 17m. A variety of calibrated radiance AVIRIS scenes are available, all shown in Figure 2.3. These data sets are primarily used for investigating hyperspectral distributions and dimension estimation.

Hyperion is a spaceborne hyperspectral imager that operates in the visible to near infrared spectral range and provides 30m ground pixels.[1] Data of the Oaxaca Valley in Mexico is shown in Figure 2.4. Nearly 50 data sets have been collected over Oaxaca since August 2008. Many of these have areas of significant cloud cover, so four hyperspectral cubes taken at different times of year will be used for parts of this analysis. These flightlines have been divided into arid, mountainous, and tropical




Figure 2.3: Visible image of AVIRIS hyperspectral data of (a) Rochester, NY, (b) Cuprite, NV, (c) Indian Pines, IN, (d-e) Lunar Lake, NV, and (f) Moffett Field, CA.

regions, each of which have different environments and phenomenology. Each region does contain regions of manmade activity, including farmland, large, and small cities. These data are also used for dimension estimation and interest segmentation.

There are 21 available scenes taken by the HYDICE sensor of a variety of regions and land cover types.[39, 9, 4] There are 15 data sets at three different GSDs of a

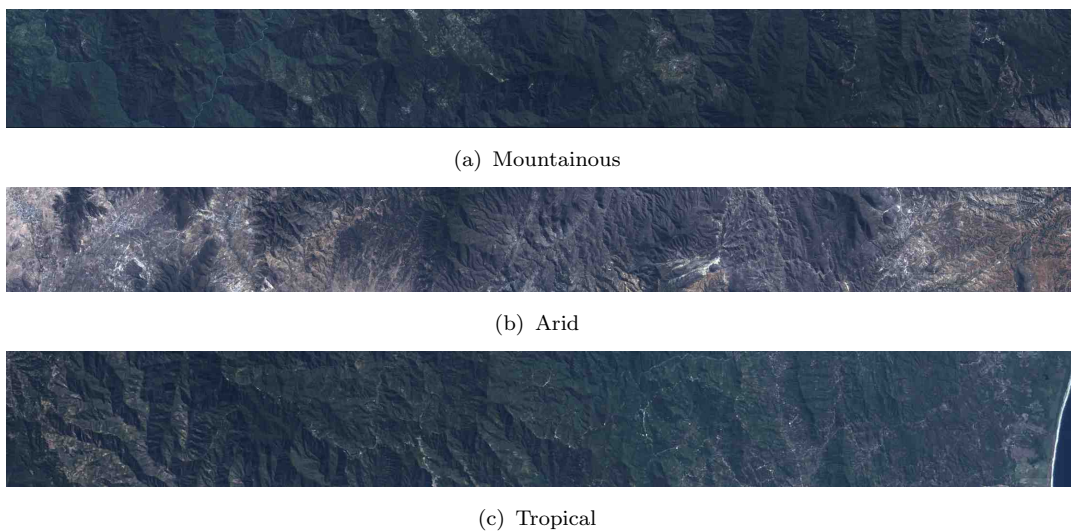


Figure 2.4: RGB images of the Hyperion data of Oaxaca Valley, Mexico.




Figure 2.5: Visible display of the HYDICE (a) forest radiance hyperspectral and (b) desert radiance data sets.


forest and grass scene, shown in Figure 2.5(a), including forest, grass, and road. In each scene, camouflaged vehicles, tarps, and targets are in the exposed grass, the treeline shadow, or in the forest. There are three data sets of a desert scene, shown in Figure 2.5(b), including a variety of ground cover types and roads at three different GSDs. These scenes contain the same vehicles, tarps, and targets as the forest scene. A scene of the entire DC Mall, including the White House and Smithsonian buildings is shown in Figure 2.6(c). Figure 2.6(a) shows an urban scene including a large shopping center, parking lot, fields, and a residential area. A natural scene of a variety of terrain, including trees, grass, exposed rocks, and hills, is shown in Figure 2.6(b). All of these data sets are used for dimension estimation. Some of the forest flightlines are used for change detection. The urban scene is used for interest segmentation.



(a)



(b)



(c)

Figure 2.6: Visible display of the HYDICE (a) forest radiance hyperspectral and (b) desert radiance data sets.




Figure 2.7: Quickbird data over the site of the Esperanza forest fire (a) and Indonesia (b).

Many of the approximately 2m GSD hyperspectral scenes will be spectrally re-sampled to the spectral resolution of the Quickbird sensor and the newly launched WorldView2 sensor in order to see how the algorithm performance changes with fewer spectral bands. WorldView2 is an 8-band VNIR/SWIR multispectral sensor.[3] QuickBird is a 4-band VNIR sensor.[3] Additional multispectral data are available from these sensors. QuickBird data of the Esperanza forest fire in California and Indonesia before and after a tsunami will also be used for some processing and are shown in Figure 2.7. Digital Globe has also provided sample imagery from WorldView2 of a forested region. These data are used for change detection and interest segmentation.

2.3 Inherent Dimension of Spectral Imagery

The distributions of different material types like those in Figure 1.3 do not take up the same amount of space in the hyperspace. Dimension estimation is one way to distinguish this characteristic. Estimating the dimension of the distribution within a spatial tile of the image can indicate whether the tile contains a significant amount of manmade, natural, or a mixture of materials if these features manifest themselves differently in the space. For this reason, estimating the dimension of each spatial tile will be developed as a test to find the presence of large scale manmade activity.

There are a variety of ways to approach estimating the inherent dimensionality of a collection of points. The spanning (or basis or encapsulating)[28] dimension is the minimum number of basis vectors required to span the space. In the hyperspectral community, it is this dimension of the data that is used in the vector subspace and

linear mixture models. The intrinsic[28] dimension refers to the smallest number of parameters needed to contain all of the variability in the data through a mapping function. PCA dimension reduction techniques estimate this value by orthogonalizing the data and choosing how much of the variability to include in the reduced image. The problem with this method is illustrated in Figure 2.8. PCA rotates the




Figure 2.8: Notional result from PCA on two different distributions, (a) and (b). Manifold dimension estimation in (c).

two different distributions in Figure 2.8 (a) and (b) into the same two orthogonal dimensions. Though the magnitude of the eigenvalues will be different for the two distributions, both directions will correspond to nonzero (or significant) eigenvalues. They will both be labeled as two-dimensional distributions. However, these two distributions are not the same and they do not take up space in the plane in the same manner; they do not lie in the same dimensional subspace of the full plane. For this reason, PCA will not accurately provide a measure of dimension unique to the data set. The spanning and intrinsic measures of dimension can be useful to reduce the dimensionality of the data if the data are well behaved and represented within the smaller hyperspace. This is often done in hyperspectral image analysis; however, the full distribution of the data usually does not conform to the underlying assumptions of PCA, the linear mixture model, or the vector subspace model. The topological dimension[28] is the dimension of the manifold or hypersurface that the data lie on or in. Figure 2.8(c) notionally demonstrates the manifold the data in (a) lie on and that it has a dimension of one.[5] The fractal dimension is an estimate of the topological dimension and a measure of how well the data fill up the space. As the goal of this research is to utilize and evaluate the distribution of hyperspectral imagery in the full k dimensions without explicitly reducing the dimensionality of the data, the fractal dimension is estimated.

Estimating the fractal dimension is possible under the assumption that the data

set is a finite sample of an infinite set. There are three main estimates of the fractal dimension, each providing a different result. The box-counting algorithm estimates the capacity dimension. The information dimension is calculated through the estimation of entropy. This is related to the capacity dimension as the probability used in the entropy calculation is the probability of point n falling in the i^{th} box. The correlation dimension is typically a less computationally intense method of calculating the fractal dimension.[6] These methods are briefly shown in Table 2.1 and further explained in the following subsections.

Table 2.1: Dimension and Estimation

Dimension	Method
Hausdorff	theoretical dimension
Capacity	box counting
Information	entropy estimation
Correlation	correlation integral

2.3.1 Capacity and Information Dimensions

Box-counting algorithms estimate the capacity dimension of a data set.[65, 6] These algorithms determine the dimension of the space by counting how many boxes (in the limit obtained by successively decreasing box size) are required to cover the entire space. The number of boxes, N , needed to cover a line segment, or a one-dimensional distribution, is proportional to the length, L , of the line segment divided by the size, ϵ , of the boxes used to cover it. N for a two-dimensional distribution, such as a square, is $\left(\frac{L}{\epsilon}\right)^2$; for a three-dimensional distribution, such as a cube, it is $\left(\frac{L}{\epsilon}\right)^3$. This relationship is shown in Figure 2.9. This is extended to a d -dimensional space as

$$N(\epsilon) = \left(\frac{L}{\epsilon}\right)^{D_{box}}. \quad (2.1)$$

Approximating the solution for the capacity dimension, D_{box} , results in the expression

$$D_{box} = \lim_{\epsilon \rightarrow 0} \frac{\log N(\epsilon)}{\log \frac{1}{\epsilon}}. \quad (2.2)$$




Figure 2.9: Graphical description of box-counting algorithm to determine the capacity dimension.

Notice the limit is as ϵ approaches zero. In implementation this means reducing the box size (and increasing the number of boxes) from a single box that covers the entire data set to a sufficiently small box size such that each box contains no more than a single point. To practically solve for D_{box} , the log number of boxes required is obtained by regression against the log reciprocal length of the box for a range of box sizes.

Entropy is a measure of the complexity within a system and is defined as

$$I(\epsilon) = - \sum_{i=1}^N p_i \log_e p_i, \quad (2.3)$$

when the data set is covered by N boxes of size ϵ and the probability of a point being in the i^{th} box is p_j . This probability used in calculations is typically a normalized count of the number of points in a particular box. The information dimension is defined as

$$D_{inf} = \lim_{\epsilon \rightarrow 0} - \frac{I(\epsilon)}{\log \epsilon} \quad (2.4)$$

and is also practically calculated through linear regression.

2.3.2 Estimating the Correlation Dimension

Estimating the correlation dimension of a data set is a different approach to fractal dimension estimation. Box-counting algorithms count the number of small elements or boxes needed to cover a space. This method does not consider whether a neighboring box is full or empty. The information dimension calculation brings the information about neighboring boxes into the calculation through the inclusion of probabilities. However, the nature of the distribution of hyperspectral imagery is not well defined in a statistical sense and thus this method is not ideal. The point density algorithm presented here does consider the density of points within small elements, in this case spheres.[6]. Grassberger and Procaccia (1983) state that the convergence of the box-counting algorithm needs a minimum of 200,000 points before the capacity dimension estimate is accurate. When applied to the same data set, correlation dimension estimation converges to an accurate result with only a few thousand data points. Because of these facts, correlation dimension estimation through the point density algorithm was implemented to estimate the inherent dimension of the distribution of hyperspectral imagery.



Figure 2.10: Illustration of volume-dimension relationship.

The correlation dimension is calculated using a correlation function. The correlation function,

$$C(N, r) = \frac{1}{N(N-1)} \sum_{i \neq j}^N H(r - \|x_i - x_j\|), \quad (2.5)$$

is essentially an averaging function dependent on the total number of points, N , and the radius, r , of the sphere where H is the Heavyside-step function that counts the number of points in the sphere. The $\|\cdot\|$ denotes the l_2 norm. The dimension is

then

$$D = \lim_{r \rightarrow 0} \frac{\log C(N, r)}{\log r}, \quad (2.6)$$

which can be solved for by estimating the slope of a line fit to the log of the correlation function *vs.* the log of the estimating sphere radii.

The point density algorithm is a different approach to estimating the correlation dimension.[20] As the name implies, the point density implementation considers the density of points within small elements, in this case spheres. These algorithms determine the dimension by relating it to the volume, V , of the space. This is illustrated in Figure 2.10. The “volume” of a line or of points on a line is proportional to the length or radius r of that line. The “volume” of a filled circle is the area of the circle, or $\pi r^2 \propto r^2$. The volume of a filled sphere is $\frac{4}{3}\pi r^3 \propto r^3$. It can be shown that the volume of a d -dimensional sphere is proportional to r^d for any dimension, d . [65] A similar relationship between the volume of a hypercube and the dimension of the hypercube is discussed in Landgrebe.[32] In a uniformly distributed space, the number of points, N_V , in the data set is proportional to the volume. It follows from this assumption that

$$V \propto N_V \propto r^d \leftrightarrow \log N_V \propto d \log r. \quad (2.7)$$

The dimension can then be approximated by

$$D \approx \frac{\log N_v}{\log r} \quad (2.8)$$

and practically solved for by estimating the slope of the line between the log of the number of points within the volumes and the log of the radii of the volumes.

2.3.3 Relating Dimension Estimations

The capacity, information, and correlation dimensions are all related to each other and to the true dimension of the data through the inequalities

$$D_{cor} \leq D_{inf} \leq D_{box} \leq D, \quad (2.9)$$

where D is the true dimension of the data.[21] The capacity dimension is an estimate of the actual (Hausdorff) dimension of the data. The information dimension is an estimator of the lower bound of the capacity dimension. The correlation dimension is the lower bound of the information dimension.[21] The three estimates are all sensitive to the distribution of the data in different ways; however, it can be shown

that if the points are uniformly distributed in the space, all estimations will equal the actual dimension. For this reason, when the estimate of the dimension does not equal the true dimension, it provides information about the changing density of pixels within the hyperspace.[20] It is possible for the inequality in Equation (2.9) to fail in practice, particularly when the data are very irregularly distributed or do not fit the number of data point criteria, or when a wide variety of spheres and/or boxes are not used in the computation. When the different dimension estimates do not conform to each other, the nonconformity provides information about the nature of the data.

2.3.4 Original Implementation of Point Density Algorithm

Implementation of this algorithm does not require calculating a correlation function when using the assumption made in Equation (2.7). Figure 2.11 shows points uniformly distributed on a line. The Euclidean distance from the center point of the line to all other points is calculated. The points within radii of increasing size are counted as a function of radius. To even further simplify the computation, the number of points within a particular radius does not need to be explicitly counted. The distance vector,

$$D = \{d_i\}_{i=0}^n, \quad (2.10)$$

is sorted such that

$$N(r) = i \quad (2.11)$$

where n is the total number of points, $N(r)$ is the number of points contained within radius r of the center point, and i is the vector index. A log-log plot of these values is produced and initially the slope of the line is found through basic linear regression and recorded as the dimension. As this method was applied to more and more data, it was found that a more robust algorithm for estimating the slope of the plot was needed. The improvements made to this algorithm are described in Section 2.3.7.

2.3.5 Data Issues and Synthetic Data Creation

As stated earlier, the fractal methods that estimate the intrinsic dimension are dependent on the distribution of the data and the assumption that the data are a finite sample of an infinite set. For example, the point density algorithm is based on the assumption that the volume enclosing the data is proportional to the number of points within the data set. This is true when the data are evenly spaced (*i.e.*,




Figure 2.11: Illustration of counting (a) and linear regression of results (b).

on a grid) in the hyperspace and approximately true when the data are uniformly distributed. However, real data are not uniformly distributed in the hyperspace.

According to Camastra,[14] there exists a requirement on the number of data points needed to accurately estimate a dimension D . The number of data points must satisfy

$$D < 2 \log N \quad (2.12)$$

where N is the number of data points in the set. For a 100 by 100 pixel image, this implies accurate estimation of a dimension up to 8. PCA typically reduces an image to 20 or 30 bands. In order to estimate a dimension of 30, approximately 600 billion points or pixels are required. But PCA is generally applied image-wide, not on a single class. We expect an image-wide estimation of dimension to be higher than any individual material class. Grassberger and Procaccia found that for the same data set, more points are required in order to reliably estimate the capacity dimension than the correlation dimension.

In order to determine algorithm performance with varying distributions in varying dimensions, synthetic data with known dimensionality were created. The number of data points in the synthetic data sets was also varied to explore the necessity of the requirement in Equation (2.12). One and two-dimensional data sets were created and embedded in a 10-dimensional hyperspace. A parabolically curved test set was created to compare to the example in Figure 2.8. These were all generated in three distribution types: evenly spaced (*i.e.*, on a grid), uniform, and normal, each of which is shown in Figure 2.12.




Figure 2.12: Example distributions of synthetic data.

2.3.6 Verification of on Point Density Estimation Synthetic Data

The results from applying the correlation dimension estimation algorithm to synthetic data sets are presented in this section. Each synthetic data set presented contained enough data points to sufficiently estimate the known dimension according to the Camastra criteria. None of the data sets generated have larger dimensions than would fit this criteria.

In each case (described below) the number of “enclosed” points as a function of radius is calculated to estimate the dimension as in Equation (2.8). The $\log N$ vs. $\log r$ point density plots (PDP) are generated for estimation of the linear slope representing D . In Figures 2.13 through 2.19, the plotted red lines represent the least squares linear regression to the plot. The red dimension is the corresponding slope of the line. The blue lines represent a visual best fit to the meaningful section of the plot described below with the corresponding slope also in blue.

One-Dimensional Data

Figure 2.13 shows the evenly spaced, uniform, and normal data results for a sample size of 500. As speculated in the example in Section 2.3, the estimated dimension of the linear data evenly spaced is exactly 1. The uniform case also shows a very straight line; however, the slope of 0.935 slightly underestimates the true dimension. In the normal case, the data begin to plateau at higher radii. This is due to the nature of the normal distribution; there are fewer points further away from the mean, so the count does not increase as quickly. If the plateau is ignored during linear fitting, the slope of the inclined portion of the line is approximately one.

Figure 2.14(a) shows the result when two linear distributions are crossed at the mean; the dimension estimate is still 1, as if the second line was not there. This




Figure 2.13: Dimension estimation plots for one-dimensional data with evenly spaced, uniformly distributed, and normally distributed data rotated into 15 dimensions.

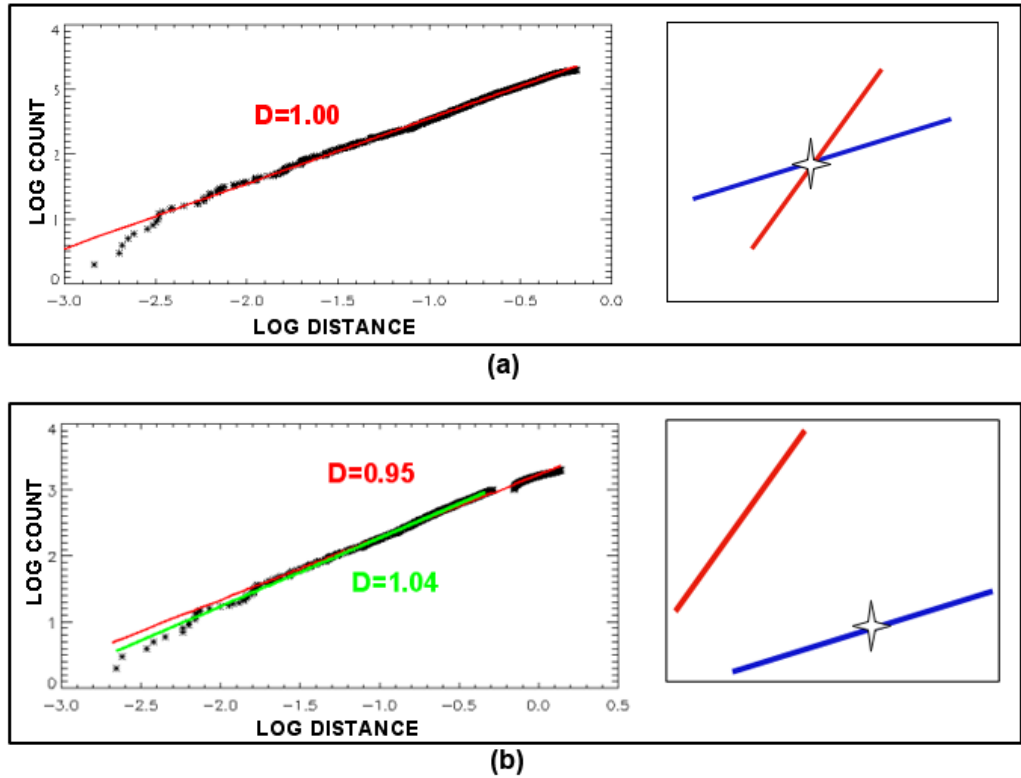


Figure 2.14: (a) Dimension estimation plot for two nonintersecting lines, starting the distance calculation at the center of one line. (b) Dimension estimation plot for two intersecting lines, starting the distance calculation from the intersection point. Also shown are notional two-dimensional projections of the data.

same phenomenon is seen when up to 7 lines in 20 dimensions were forced to cross in the center. Figure 2.14(b) shows the resulting plot where there are two non-crossing lines in a 20 dimensional space, and the calculation begins on the center of one line. The slope is still estimated to be about 1; however, there is a gap in the plot that corresponds to the empty space between the two lines.

Two-Dimensional Data




Figure 2.15: Dimension estimation plots for two-dimensional data with evenly spaced, uniformly distributed, and normally distributed data rotated into 15 dimensions.

Similarly, the results for two-dimensional data are shown in Figure 2.15. The evenly spaced and uniform data result in dimension estimates that are very close to 2. Notice that even though the estimate is accurate, there is a small plateau in the plot at high radii. These two-dimensional data, visualized in Figure 2.12 are constrained to lie within a square. Larger radii are required to count the corner points of the square, but there are fewer points at this distance to count.

The same plots for a two-dimensional square and a circle of data contained in that square are shown in Figure 2.16; the plateau is not present in the circular data. The plateau is, again, more severe in the normal distribution case, for the same reason as stated above. By excluding that portion of the plot during linear fitting, there is a best fit line with the slope equal to 2.

Multiple-Dimensional Data

A remotely sensed scene typically contains many materials; it is unlikely these will always lie in the same dimensional subspace. Data sets combining uniformly distributed data within circles of different radii and density of points were created and embedded independently into the same 20 dimensions. Figure 2.17(a) shows the

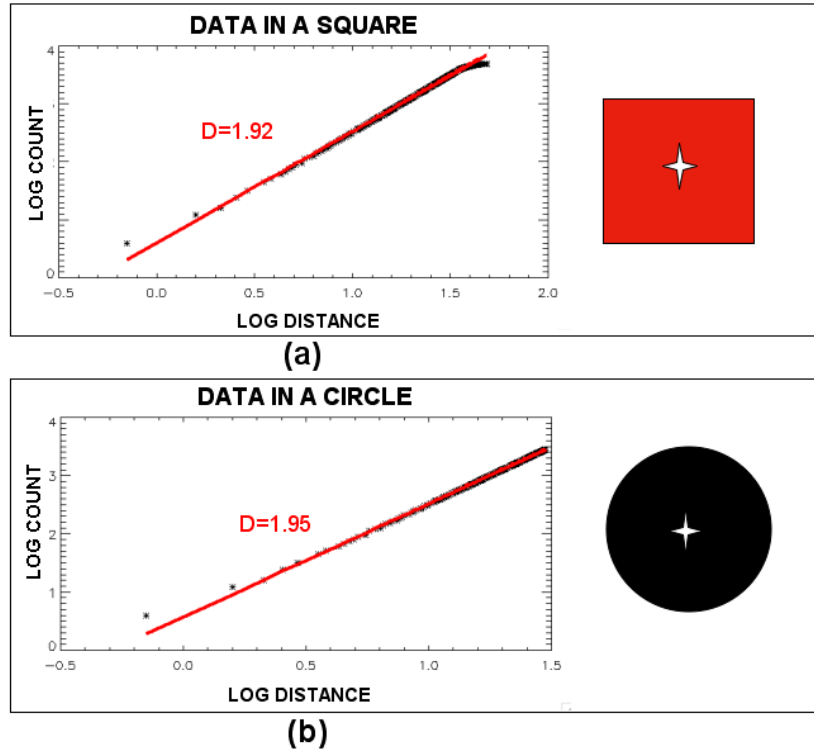


Figure 2.16: Dimension estimation plots for a two-dimensional square (a) and circle (b) rotated into 20 dimensions illustrating the plateau effect. Also shown are notional two-dimensional projections of the data.

plot for a large circle surrounded by three lower point density circles. There is a sharp plateau when the distance exceeds the radius of the primary circle. The best fit slope to the entire line is now much lower than 2. A high point density circle surrounded by a low density circle shows a similar plateau in Figure 2.17(b), and the best fit slope is also significantly reduced, to 1.6. When a high density circle is partially overlapping a low density circle, instead of a plateau, Figure 2.17(c) indicates that the slope increases at large distances from the center.

A uniformly distributed, seven-dimensional circle with 5000 data points was created and rotated into 20 dimensions. The resulting plot in Figure 2.18 shows the best fit line to the entire graph has a slope of approximately 7.

Parabolically Curved Data

Section 1 described the methodology for using PCA and other, similar statistical methods to find two independent dimensions in a curved data set. Other techniques,

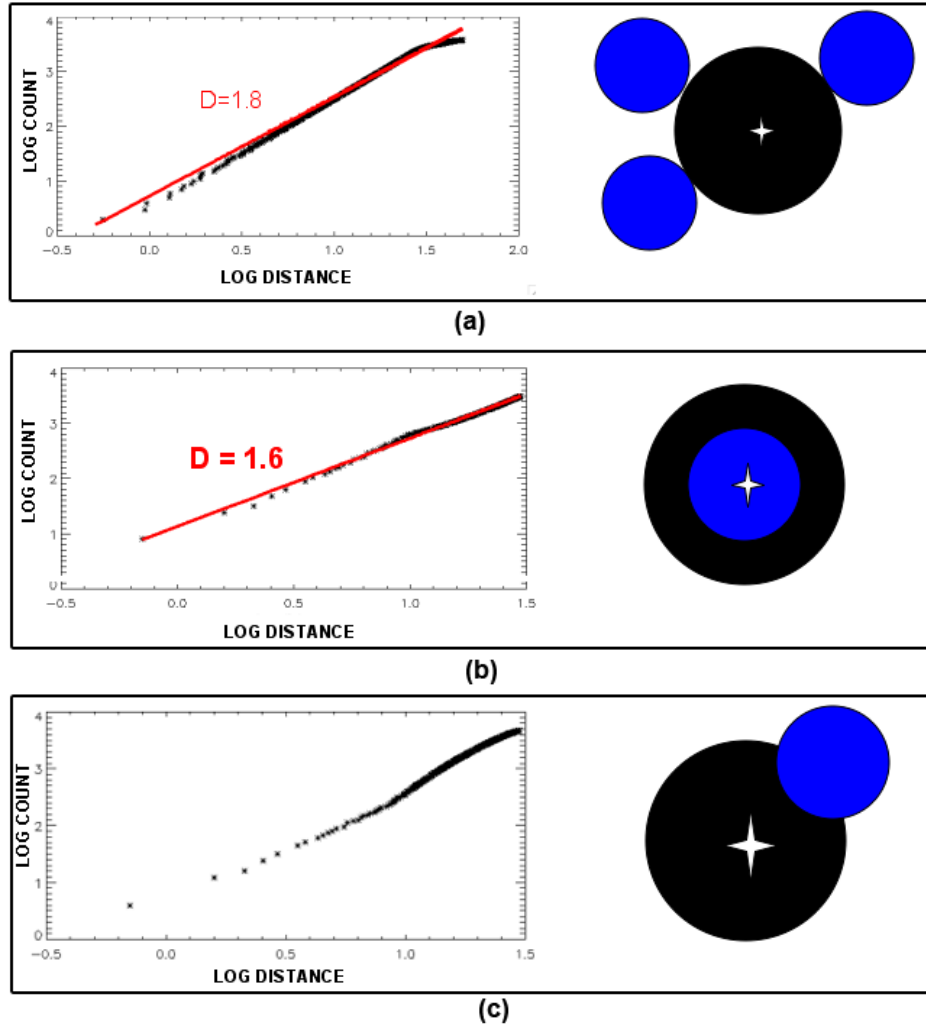


Figure 2.17: Dimension estimation plots of (a) a large circular distribution with three smaller circles surrounding the large circle, (b) a large low point density circle with a small high point density in the center, and (c) a large low point density circle with a small high point density overlapping. Also shown are notional two-dimensional projections of the data.

like manifold dimension, estimate the dimension of any curved set to be one. The correlation dimension algorithm was applied to a parabolically curved set (Figure 2.12). The resultant dimension estimation plot is shown in Figure 2.19. The estimated dimension of the generated parabolically curved data set is found to be 0.9. This approximately agrees with the manifold dimension of one.

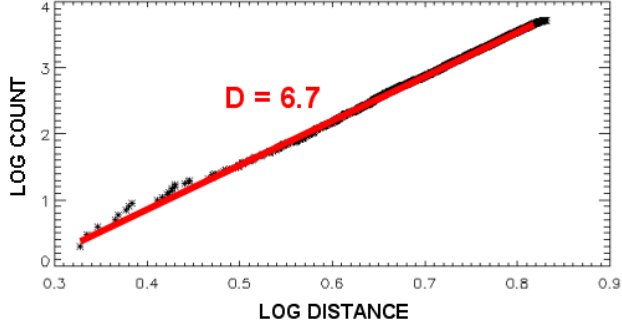


Figure 2.18: Dimension estimation plot for seven-dimensional uniformly distributed data with 7000 points rotated into 20 dimensions.

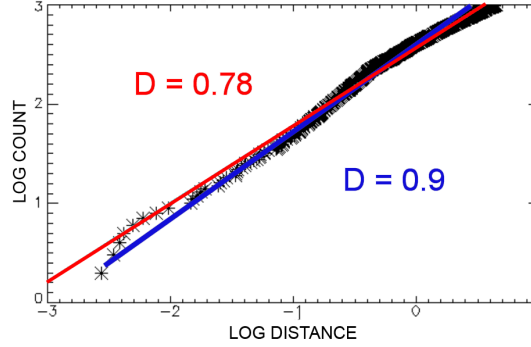


Figure 2.19: Dimension estimation plots of normally distributed data on a parabolic curve in 15 dimensions.

2.3.7 Improvements to Point Density Dimension Estimation

The results from non-ideal synthetic data presented in Section 2.3.6 show the log-log plots created with the point density algorithm to have two distinct features shown in Figure 2.20: a steep incline and a plateau or tail. The steep incline is used to provide the estimate of the dimension of this primary cluster in the data. The tail indicates the presence of points at far distances from the primary cluster center. The tail may be the result of a sudden or gradual reduction in the point density of the cluster or the presence of lower density clusters outside the primary cluster. In the synthetic data, the transition between these two sections is sudden because there are no mixed pixels. In real data, the transition may be sudden or gradual, depending on the pixel contents in the region taken. Mixed pixels are far enough from the primary cluster that they will not generally affect the slope estimate. An exception to this case is when there is no primary cluster, in which case the plot

will not have a clear incline at all.




Figure 2.20: Notional point density plot illustrating the incline and tail regions.

Due to the presence of the transition and tail regions of the plot in real data, an automated method to calculate the slope of the primary incline was developed. This routine chooses the appropriate slope based on minimizing the residual error between the points that fall on the incline and a straight line drawn through those points. To estimate the best fit slope of the incline, the tail points are located and removed from the slope calculation. A straight line is drawn from the first point in the graph to a random midpoint and the error between the line and points is calculated. The error function is minimized to find the best fit line from which to estimate the slope. The pixels mapped to the transition and plateau region are spectrally different from those used to estimate the dimension and are excluded from the dimension estimation. The transition and plateau regions of the point density plot have proven useful in indicating the amount of interest in a spectral tile and are further described in Section 2.4.1.

2.3.8 Implementation of the Box-Counting Algorithm for Capacity Dimension Estimation

The box-counting algorithm for capacity dimension estimation was provided in the form of C-code by Dr. Ronald G. Resmini of the Geography and Geoinformation Science Department at George Mason University but is presented here for comparison. The output of this program is the number of boxes that contain pixels, $N(\epsilon)$, for each iteration of a particular box size, ϵ . This information can be used to generate the $\log N(\epsilon)$ vs. $\log \frac{1}{\epsilon}$ plot, shown in Figure 2.21. From this plot, the dimension is estimated by fitting a straight line to the linear mid-section of the plot.[43] Similar to the PDP, other information can be gathered from the box-counting process which may be of use in anomaly detection and box-counting in future research.




Figure 2.21: Example plot for box-counting estimation.

2.3.9 Verification on Synthetic Data




Figure 2.22: Example plot for box-counting estimation of a one-dimensional data set.

Results from the box counting algorithm applied to synthetic one- and two-dimensional data sets in Figures 2.22 and 2.23. In both cases, the correct dimension can be estimated from the plot. As the dimension of the data increases, a significant number of data points are required to correctly estimate the dimension, which also dramatically increases the run time. The verification of this algorithm on the same




Figure 2.23: Example plot for box-counting estimation of a two-dimensional data set.

simple data sets as the point density algorithm is taken as sufficient evidence that the two algorithms can be compared with application to real data, presented in Section 2.3.10.

2.3.10 Comparison of Dimension Estimates

Comparison of these dimension estimates on real spectral data is necessary to determine what phenomenology of the data each method is measuring and which is most appropriately used for a tiling scheme.[54] A total of 37 hyperspectral images and over 250 subset cubes derived from these images were used. For each image and subset, three dimension values are estimated: correlation, capacity, and principal component dimension. The PC dimension, denoted PCA in the following figures, is defined as the number of principal components needed to contain 99.5% of the variability in the scene. The capacity and correlation dimensions are denoted BOX and PDP in the resulting plots. Figure 2.24 shows the estimated dimension of full images as a function of scene content, sensor, and GSD. In general, the large GSD scenes, like those from Hyperion, have larger dimension estimates than higher spatial resolution data. Some scenes with small GSDs show increased dimension estimates in at least two of the algorithms. These scenes, mainly Moffet Field, Cooke City, Urban, and DC Mall, all contain urban regions and can be said to be more complex than the scenes containing only natural materials. An increased dimension estimate may indicate increased complexity of or material variability within the scene.

Figure 2.25 shows the estimated dimension of HyMap data as a function of the PDP dimension. The points on the left (labeled “subset”) are the average estimated dimensions of seven flightlines for five unique material subsets. The points on the




Figure 2.24: Estimated dimensions of full scene hyperspectral data as a function of sensor, image content, and GSD.

right (labeled “image”) are the estimated dimension of seven full flightlines covering the same area. The individual material subsets have no regular pattern across dimension estimates and are dependent on the content of the subsets. In general, the PCA and PDP dimension estimates are higher than the BOX estimate, except for a single point. These patterns are not observed for the image wide estimates. The PCA and BOX dimension estimates of full images are consistent across the seven flightlines. The PDP estimates of full images are consistently larger than both the PCA and BOX dimension estimates, but they are not constant and range from approximately 5.5 to 8.2. Also different is that the BOX dimension estimate is higher than the PCA dimension estimate image wide, the opposite of what is seen for image subsets. For full image dimension estimation, the BOX or PCA algorithms will provide more consistent results than PDP. However, for dimension estimation of spectral image subsets such as those to be analyzed in a tiling scheme, PDP provides a more reliable result. With only a few thousand pixels, the box-counting algorithm cannot accurately converge on the inherent dimension of the data resulting in significant variability of the results.




Figure 2.25: Estimated dimension as a function of PDP dimension for HyMap data.

2.4 Characteristics of High Dimensional Distributions

2.4.1 Additional Features from the Point Density Plot

The PDP is an alternative way to represent the changing density of the data in the hyperspace with respect to the center of the distribution. In the ideal case, where the distribution is uniformly spaced and spherical, the PDP is a straight line as in Figure 2.11 and the slope can be easily calculated. However, as the distribution differs from ideal, the slope tapers into a zero slope line at large radii. When practically using the point density approach, the PDP generally can be divided into two regions, shown in Figure 2.20: the incline and the tail. Basic linear regression was not an accurate method for estimating the slope of the PDP.[50, 53] Instead, a new method to estimate the slope of the incline was developed. In order to accurately estimate the slope of the incline, the points corresponding to the transition region and tail are identified and removed from the slope calculation. As a result, two additional metrics, the point density tail length (PDTL) and total error to the PDP (TEP) are recorded each time the dimension is estimated. The PDTL is essentially the

length, in radii units, of the plot range which includes only tail points. The TEP, a measure of the transition region, is the total error between the PDP data and the linear-piecewise fit to the calculated slope and PDTL. The process for estimating the inherent dimension, the PDTL, and the TEP is shown in Figure 2.26.

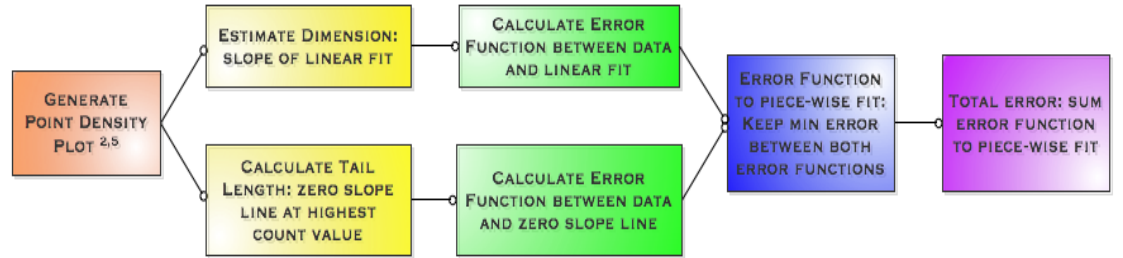


Figure 2.26: Process used to estimate the dimension, PDTL, and TEP from the PDP.




Figure 2.27: Examples of incline and tail on a point density plot created from real hyperspectral data of road (a) and forest (b).

Figure 2.27 shows two examples of the PDP for a group (a) of pure road pixels and (b) a region of forest pixels. The plots have the two distinct linear regions demonstrated in Figure 2.20. The slope of the incline portion of the plot is an estimate of the dimension around the center or highly populated region of the data cloud. This estimate of dimension excludes influence from anomalies and others points at a large distance from the cluster center. The flat tail of the plot comes from the relatively few points in the hyperspace that are at large distances from the center and violate the assumption that the points within a single material cluster

have common dimensionality. The length of this tail corresponds to the range in spectral space in which the spectrally distant points outside of the primary material cluster fall. If the length of the tail is short, there are few pixels at a large distance from the primary cluster, which may indicate a more pure material. If a long tail is present, there are many pixels at a large distance from the primary cluster. This may be evidence of a cluster of points, far from the centroid, or many small or isolated anomalies in the sample. The TEP is a measure of the transition region of the plot, located between the points that are a part of the steep incline and the tail. An image with a significant amount of spectral variability, whether due to natural or material variability, will have a large transition region and TEP. On the other hand, a collection of single material pixels, such as the road in Figure 2.27, will have a relatively small transition region. Both the PDTL and TEP can indicate whether “interest” exists in a tile.

2.4.2 Volume Estimation

The estimated inherent dimension of a spectral distribution is one way to measure the “size” of a distribution in the hyperspace. Another way to characterize this “size” is to estimate the volume of the distribution in the hyperspace, which is related to the complexity. A complex distribution has a variety of spectrally diverse materials and as a result will occupy a “larger” volume in the hyperspace.[38, 69] A distribution with few materials that are spectrally similar is neither complex nor interesting. An obvious method for calculating the volume of a distribution would be to calculate the volume of the convex hull in which the data lie. If the endmembers of the convex hull are found, the determinant of the endmember matrix is an estimate of the volume. However, for this to work, the dimensionality of the endmember matrix must be reduced to a square matrix of possible lower dimensionality than the data. As a result, information is lost and the result may not be an accurate estimate of the volume. This also assumes that the correct number of endmembers needed to describe the data is known *a priori*, which is not the case.

Messinger, *et al.*, (*in prep*-2010) present an iterative method that calculates the volume of the parallelepiped that encloses the data as a function of the number of endmembers used in the calculation. This is done through the use of the Max-D algorithm[55] and the Gram Matrix, which both have unique properties useful for volume estimation. The Max-D algorithm returns the endmembers in order of decreasing magnitude of separation. The Gram Matrix, G is defined as

$$G_{i,j} = \langle x_i, x_j \rangle, \quad (2.13)$$

an $n \times n$ matrix where n is the number of vectors (or pixels) in the test set, $x_i|_{i=1}^n$. The local Gram Matrix, computed with respect to a particular element, k , of the set, is defined as

$$G'(x_k)_{i,j} = \langle (x_k - x_i), (x_k - x_j) \rangle. \quad (2.14)$$

The determinant of the Gram Matrix, termed the Gramian, is the squared volume of the parallelepiped formed by the vectors in the test set.[29] Due to the geometry of the data, the volume of the parallelepiped that encloses the data is the same as the volume of the parallelepiped that encloses the endmembers. For the complexity method used here, the Gramian is calculated iteratively as a function of the number of endmembers used in the calculation, producing a plot shown in Figure 2.28. The Gram Matrix approach to volume estimation of a distribution may also be useful for change detection. The presence of unusual and distinct materials in the distribution causes the volume to increase. Given two registered tiles, one of which contains an additional and unique material, the two tiles will have two significantly different volumes.[69]

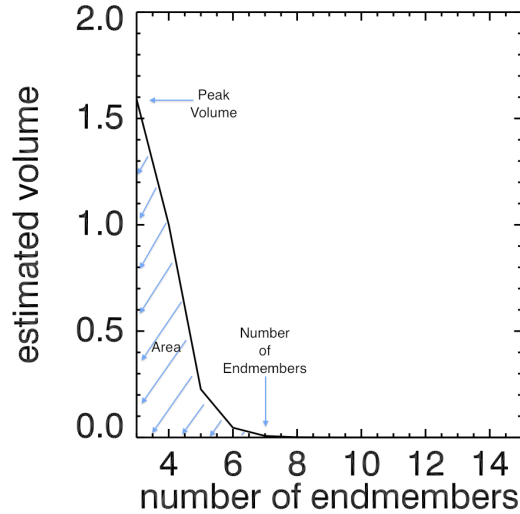


Figure 2.28: Example Gram Matrix volume estimation as a function of number of endmembers plot.

2.5 Anomaly Detection

Anomaly detection locates pixels in the scene that are unlike the background distribution. These routines do not depend on the number of anomalous pixels, so they can find both single and groups of pixels. Thus, anomaly detection can be used to locate both large and small scale regions of interest in a hyperspectral tile. The result from an anomaly detection routine can be used to develop features relating to the amount of interest in the scene ranging from the number of singular anomalies or anomalous objects and how many pixels are in the cluster.

The RX algorithm, described in Chapter 1, is the most commonly used statistical detector because it is theoretically a CFAR detector that works well in scenes consisting of mostly natural material. The scenes this methodology would be applied to may or may not consist of mostly natural material. However, even if they do, it is unlikely that every individual tile is predominately made of natural material. In the case where a static city is a spatial tile, the whole city may appear to be anomalous, although it is not interesting by our definition. In addition, a single tree stump or rock in a grass field will also be anomalous when using the RX algorithm, though it is not interesting or even manmade. In order to separate the “interesting” anomalies from “uninteresting” anomalies to which the RX algorithm is not robust, other methods must be used.

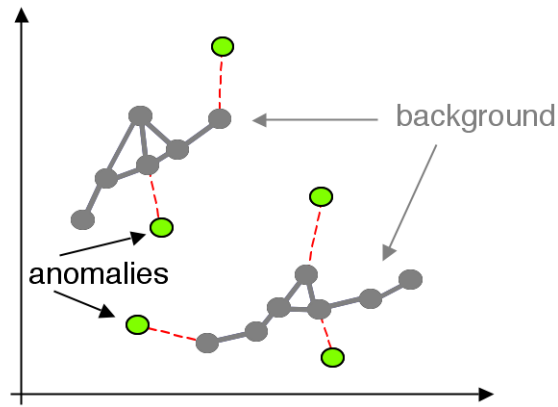


Figure 2.29: Notional result of TAD on a simple data set.

Topological analysis and graph theory provide ways to make data driven measurements on the data. The Topological Anomaly Detector (TAD) [11] uses methods from these fields in order to locate spectral anomalies. This algorithm connects points in the hyperspace that are “close” to each other in order to build a simplicial

complex, where the closeness is a free parameter of the algorithm. This process may find a single or many graph components of a variety of different sizes. As anomalies have a low frequency of occurrence, the largest components are labeled as background. The small or single point components are labeled anomalies while the large components are labelled background, shown in Figure 2.29. How anomalous these pixels are is determined by measuring the distance between the anomaly and the nearest edge of the background component.

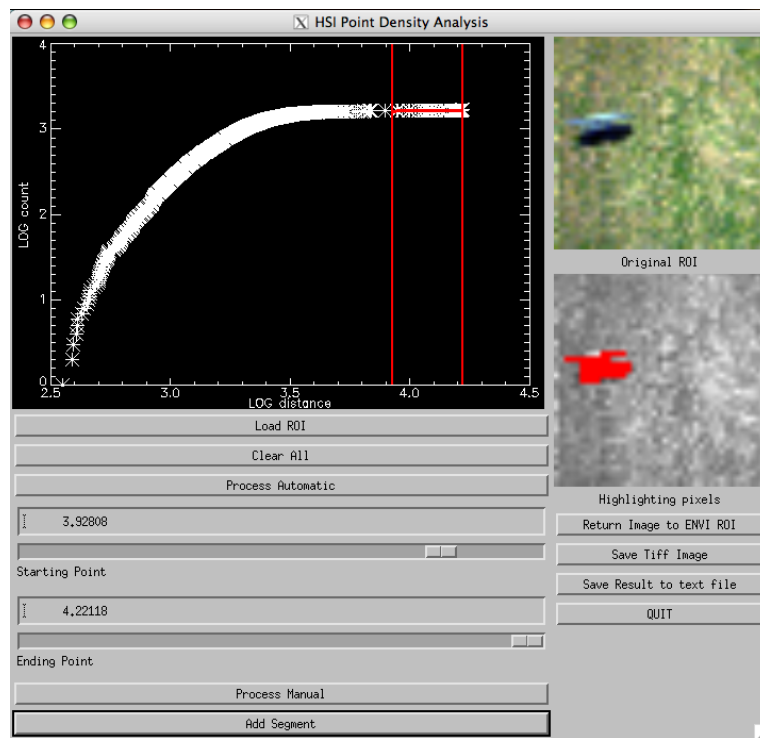


Figure 2.30: GUI output displaying the relation between the PDDL and the anomalous pixels in a scene.

Once anomaly detection is applied to the full hyperspectral scene the resulting anomaly map is tiled. For each tile, a score is assigned. Using the number of pixels or fraction of pixels in the tile is one score. However, 10 single-pixel anomalies and one 10-pixel anomaly in a tile would be assigned the same score, but they refer to two different phenomena. A multiple-pixel anomaly that is large will likely manifest itself in the full distribution in such a way that another test, perhaps dimension or classification, would find it. The primary function of anomaly detection in this methodology is to locate small scale anomalies. At this time, the number of

anomalous pixels contained in a tile determined by the TAD algorithm is measured as a spectral feature.

The number of anomalies contained in a scene tile is not the only feature measuring anomalousness used in this methodology. In some sense, the PDDL is also a measure of the amount of “anomalous” pixels contained within a scene. This is shown in Figure 2.30, where the pixels corresponding to points on the tail of a distribution are falsely colored red. The number of anomalies contained in a tile from TAD is primarily used for measuring the amount of small scale anomalies in a tile while the PDDL is more sensitive to large scale anomalies within a tile.

2.6 Classification and Spectral Clustering

Classification can provide information indicating the amount of interest in a tile. For example, if a tile can be reliably shown to contain a class of one material, it is not interesting. A tile containing multiple classes may be interesting or not depending on the size of the classes. The supervised and unsupervised classification methods described in Chapter 1 work well on low spatial resolution data where each pixel is an average of a large piece of land. The limitations of the described algorithms come from the fact that, when applied to high resolution data, the descriptions of the data in the hyperspace are no longer accurate, resulting in poor classification maps. To be applied to a wide variety of scenes and resolutions, data driven methods for classification are established in this thesis.

A recent example of these alternative, data driven approaches is the work of Bachman, *et al.*, who use manifold geometry in the hyperspace.[5] The driving force behind using manifold geometry is the nonlinear characteristics of hyperspectral imagery. ISOMAP, the isometric mapping algorithm for deriving manifold coordinates, is applied to spatial tiles of hyperspectral imagery. The computational and memory requirements of using the algorithm on a full hyperspectral scene do not enable a practical processing scheme. Instead, the manifold coordinates are determined for non-overlapping spatial tiles of the scene and then connected to create the full-scene manifold. Classification based on manifold coordinates was compared to a linear unmixing method and showed fewer false alarms on spectrally similar targets.[5]

Classification of hyperspectral images is difficult because what an analyst visually sees as a class may not be spectrally similar. Unsupervised clustering of hyperspectral data removes the user from the situation and finds spectrally similar clusters in the hyperspace. Basener, *et al.*, [10, 31] introduce a clustering algorithm, named the gradient flow algorithm, which clusters points together based on the direction

of increasing density. For each point in the hyperspace, the n nearest neighbors are found and the density of each point with respect to the n neighbors is calculated. A smoothing step locates points that are in the center of high density regions. Finally, the gradient differential is calculated; for each point, this will find the direction of the highest increasing density and “push” the point in that direction, thus assigning it to the same cluster as other points associated with this high density sink. This step is iterated until a steady state is reached and a cluster map can be displayed in image form. This algorithm is illustrated in Figure 2.31. This algorithm requires no prior estimation of the number of classes in the scene and is not restricted to geometric or statistical assumptions. The number of nearest neighbors and smoothing iterations is controlled by the user and can affect the class resolution of the output. A high class resolution coincides with individual clusters for each general class. For example, if vegetation is the general class, a high class resolution would result in each type of tree and grass in the scene becoming a separate cluster; low class resolution would put all vegetation into one cluster.




Figure 2.31: Notional description of the Gradient Flow algorithm. The points in (a) are connected to their nearest neighbors in (b). The direction of highest density change are shown in (c).

A novel clustering method, inspired by TAD, takes a network theory approach to spectral clustering. Similar to TAD, the first step is to treat the points in the hyperspace as a complete graph, where every point is connected. Maximum modularity is a concept used to describe whether the full distribution can be described in terms of communities of points that are “like” each other, illustrated in Figure 2.32.[33] This algorithm works in iterations, creating a “tree” shown in Figure 2.33. The first iteration divides the image into 2 spectral clusters, in this case land and water. Each of these spectral clusters can then be divided into two smaller clusters. This process is self-terminating and continues until no cluster can be further divided in two. More research into this approach to spectral clustering is underway.




Figure 2.32: Notional example of individual communities in a two-dimensional distributions.

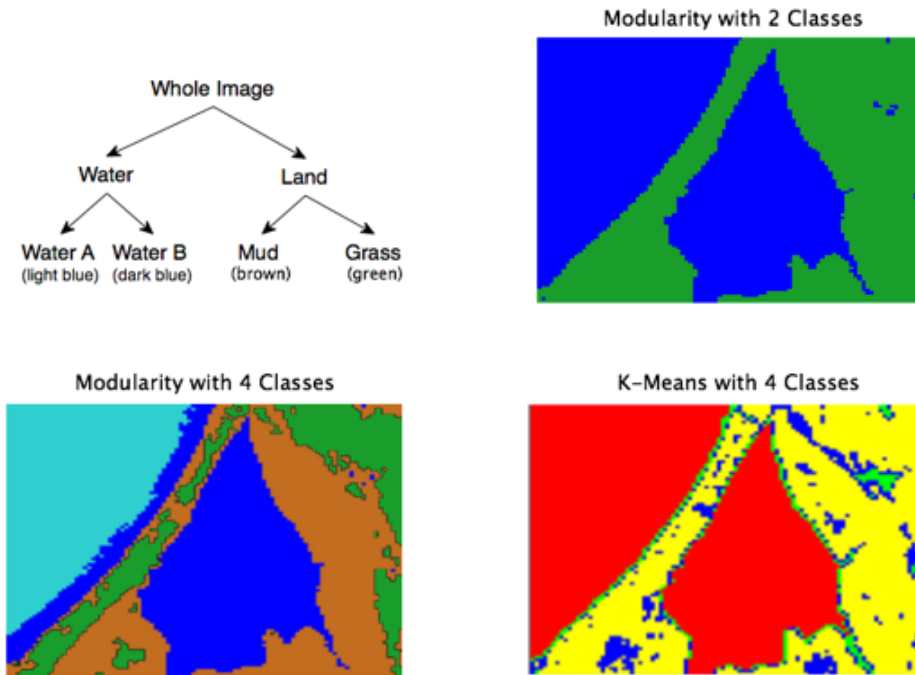


Figure 2.33: Example of the network theory approach to spectral clustering on AVIRIS data of Braddock Bay near Lake Ontario in Western NY.

The gradient flow clustering algorithm is still in the early stages of development. In order to illustrate the concept, a version of the algorithm was applied to the




Figure 2.34: (a) The scatter plot showing the classes from the gradient flow algorithm and corresponding gradient flow result on an urban HYDICE scene. (b) Corresponding results for the k -means algorithm. (NOTE: The colors/gray scales do not distinguish all 13 major classes. In particular, there is one tree class and one long grass class and they are both displayed as the same color.)

HYDICE urban radiance data set. These results are shown in Figure 2.34. The left side of the graphic shows a two-dimensional projection of the data, color coded by class. The right side shows the classification map. This method found 13 different clusters in the data. For comparison, the same data set was classified using the k -means algorithm where $k = 10$, shown in Figure 2.34. This number was chosen because there were ten unique, pure pixel materials identified in the scene. The gradient flow algorithm does a much better job of class separation than the k -means




Figure 2.35: Number of classes in a tile as a function of subjective amount of interest in the tile for (a) Gradient Flow and (b) k -Means.

algorithm, particularly in the more cluttered, residential areas, which include houses with different roof types, a variety of vegetation, and multiple road classes.

Figure 2.35 shows the number of clusters in each tile as a function of the amount of interest contained in the tile for two algorithms: k -Means and Gradient Flow. The classification algorithms were applied to full images and the classification maps were tiled. The amount of interest was determined qualitatively. From this result, it is clear that with this particular data set, the number of clusters found by Gradient flow has the linear relationship expected with amount of interest. A tile with more

clusters is more interesting than a tile with few clusters. There is no relationship between interest and number of clusters found by k -Means as this method found the maximum number of classes in nearly every tile. As a result, Gradient Flow was chosen for this research and is applied to the full image and the resulting classification map is divided into spatial tiles. A tile with only one or two classes, containing only vegetation, is a large scale region that can immediately be labeled as not interesting, while a tile with many classes is deemed more interesting.

2.7 Feature Combination for Interest Segmentation

The previous sections describe the mathematical theories from which the feature metrics are developed. Table 2.2 shows the list of feature metrics that have been developed for this research. Each feature will be calculated on a per tile basis. Dimension estimation, PDDL, TEP, and volume estimation are calculated on each hyperspectral tile. Spectral clustering and anomaly detection are applied to the full hyperspectral scene and the features are measured for each tile in the classification and anomaly map. The relationships between the features developed and amount of “interest” contained in a tile are verified in the next section.

Table 2.2: Proposed List of Features.

Number	Metric	Method
1	Tile Dimension (PDD)	PDP
2	Point Density Tail Length (PDDL)	PDP
3	Total Error to Ideal PDP (TEP)	PDP
4	Number of Spectral Clusters (NCLUST)	Gradient Flow
5	Number Anomalous Pixels (NANOM)	TAD
6	Peak Volume (PV)	Gram Matrix
6	Area of Volume Plot (AREA)	Gram Matrix
6	Number of Endmembers (NEM)	Gram Matrix




Figure 2.36: Number of anomalies contained in a tile as a function of the subjective amount of interest contained in the tile.

2.7.1 Preliminary Verification of Feature Utility

These features have been chosen because they derive information about the distribution of the data with limited assumptions and can be easily measured. Each of these metrics can also be correlated with the amount of interest contained in a spectral image tile. For example, Figure 2.35(a) verifies the intuitive relationship between the number of Gradient Flow spectral clusters in a tile and the amount of interest in a tile. Figure 2.36 shows the number of anomalies found in a tile as a function of subjective amount of interest in a tile. In general, the “interesting” tiles contain a large number of anomalies while the tiles with little or no interest contain very few or no anomalies. However, there are “interesting” tiles with few or no anomalies, so alone, the number of anomalies is not a sufficient metric.

To test the accuracy of using the estimated dimension, PDTL, and TEP as features to indicate the presence of manmade materials in a predominantly natural scene, similar natural scenes with and without manmade materials are required. Three tiles, with approximately 10,000 pure natural material pixels each, were selected from a HyMap reflectance hyperspectral data set of Cooke City, MT, shown in Figure 2.37.[61] The presence of manmade pixels is simulated in these tiles. Caefer, *et al.*(2007), present a subpixel target implantation method that takes manmade material spectra from a library and mixes them with in-scene pixels.[13] In our case, manmade pixels in downtown Cooke City, including various buildings, parking lots, and cars, were selected from the same data set as the tiles and put into a library of 2333 manmade pixels. Taking manmade pixels from the same scene guaranteed the spectra were all in the same units and had gone through the same atmosphere




Figure 2.37: Three natural material tiles from HyMap data near Cooke City, MT used for pixel-injection analysis.

and calibration. These manmade spectra were used to replace spectra in the natural tiles so there were full pixel manmade materials in the scene. Figure 2.38 shows the average spectra of the three natural tiles and the average spectrum of the manmade material spectral library. In order to observe how the PDD, PDTL, and TEP change




Figure 2.38: Average reflectance spectra of forest (green), grass (cyan), foothills (magenta), and city (black) regions used in this analysis.

as more manmade pixels are included in the scene, an increasing number of random pixels in the natural tiles were replaced with random manmade spectra taken from the library. At every increment of the number of manmade pixels, 200 random trials of the pixel injection were done, resulting in 200 estimates of each metric for each

percentage of manmade pixels. The mean estimated metric of 200 trials is displayed on all plots.




Figure 2.39: Dimension and error of linear piece-wise fit as a function of percentage of manmade pixels included in a forested region.

Figure 2.39 shows the mean estimated dimension and the mean line fitting error of 50 and 200 trials for the forest scene on the same plot. The error is relatively constant as the dimension begins to increase. At around 15% manmade materials added, both the dimension estimate and the error spike. This may be that at around 15%, the manmade materials begin to make up a significant portion of the distribution and as a result affect the center of the distribution. The algorithm assumes that the dimension is being estimated from a starting point that falls in the center of a single cluster of data. Adding a significant portion of manmade pixels to the forest increases the clutter substantially and an accurate dimension estimation cannot be made due to the multi-material clutter. This prevents the PDP from developing a clear incline portion from which to estimate the dimension, so a large number is estimated. An example is shown in Figure 2.40. The TEP may be a way of estimating the accuracy of the dimension estimate. Typically, the piece-wise function has a poor fit, or a large TEP, when the point density plot has no linear incline and only a transition region and flat tail, which presents as an increasing curve with decreasing slope. For this reason, the remaining graphs only display up to 15% manmade pixels injected. Note that this effect was best observed




Figure 2.40: Example PDP without a clear incline portion illustrating that the estimated slope is too high.

in the forest tile and the physical cause is not entirely clear. The forest is a less cluttered region in this scene than the grass or foothill regions. It may be that since the background material is more pure, it is more easily affected by the addition of manmade materials. The grass and foothill regions, though both natural, contain multiple types of natural material, adding to the amount of clutter in the tile.




Figure 2.41: Estimated dimension as a function of percentage of manmade pixels for three natural regions.

Figure 2.41 shows the mean estimated dimension as a function of percentage of manmade pixels in the tile. The dots on the y-axis show the dimension of the hyperspectral tiles with no manmade materials and the estimated dimension of the library of manmade materials. In general, the more manmade spectra added to the tile, the higher the estimated dimension. The foothills are the most cluttered region before the manmade materials are added and as a result it takes a larger number of manmade pixels until the dimension increases significantly.




Figure 2.42: Dimension estimation plot tail length as a function of percentage of manmade pixels for three natural regions.

The tail length as a function of the percentage of manmade materials is shown in Figure 2.42. In this case, the trend for each tile is identical. When just a few manmade pixels are added to the tiles, they are at a large distance from the center of the distribution, and therefore lengthen the tail of the point density plot. As more manmade pixels are added, the center of the distribution is pulled towards the manmade pixels and the tail length starts to shorten again.

Another effect of adding a significant number of manmade pixels is the lengthening of the transition region and increase in the TEP, shown at the higher percentages of manmade pixels in Figure 2.43. The TEP is initially lower for the forest region than the grass or foothills, which are both more cluttered tiles, and continues to be lower with increasing numbers of manmade pixels.

Messinger, *et al.*, use the relationship between the peak volume of the Gram




Figure 2.43: Error of linear piece-wise fit to dimension estimation plot (TEP) as a function of percentage of manmade pixels included in three natural regions.

Matrix volume and spectral image complexity. However, the number of endmembers where the volume is zero and the area under the curve also relate to complexity. Five hyperspectral tiles of increasing spectral image complexity from left to right are shown in Figure 2.44. The corresponding volume plots in Figure 2.45 show that the three metrics that can be measured from this plot all increase as the spectral complexity of the tile increases and therefore all can indicate the relative amount of complexity contained in a tile.



Figure 2.44: Five RGB tiles from the HyMap Cooke City, MT collect. Each tile contains approximately 4000 pixels. (a) trees (b) grass (c) grass, road, & trees, (d) a construction area, and (e) a small city.

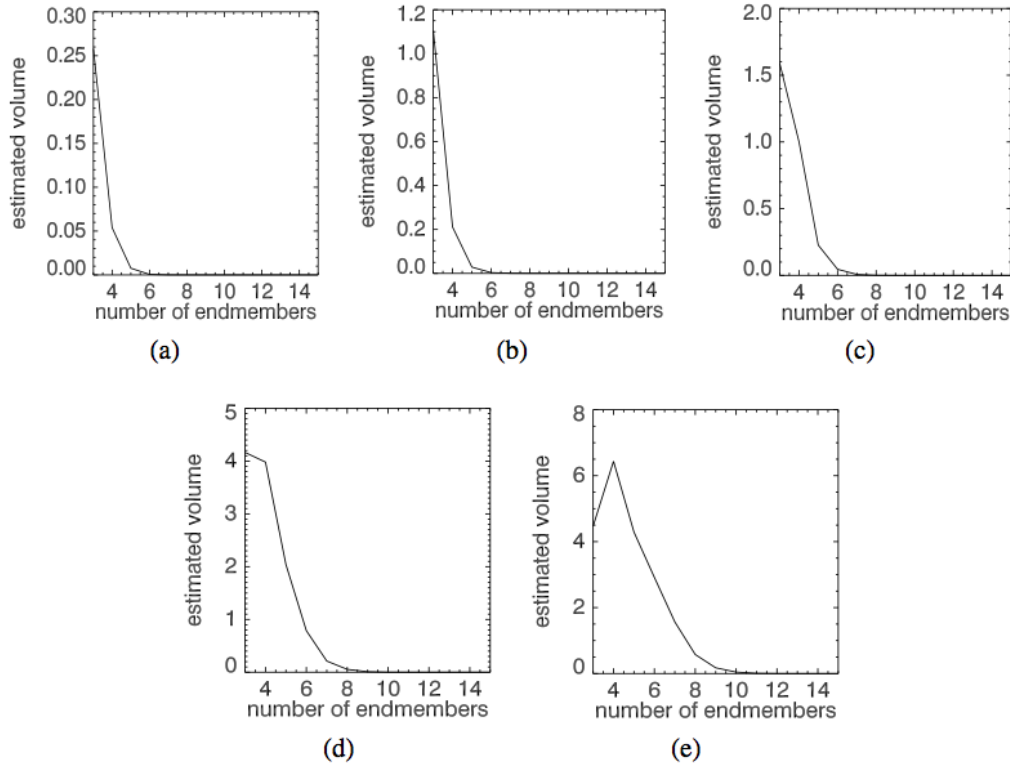


Figure 2.45: Complexity plots for the five hyperspectral tiles. (a) trees (b) grass (c) grass, road, & trees, (d) a construction area, and (e) a small city.

2.7.2 Feature Combination

The ultimate goal of this research is to use the implications of the features introduced here to determine whether a particular spatial tile in a large area scene is interesting. Hypotheses are generally made in the scientific community about what is “the truth.” The National Image Interpretability Rating Scale (NIIRS) is an example of the scientific method that is used by imagery analysts to determine the quality of an image.[55] NIIRS is used in order to rate visible images based on an analyst’s ability to perform exploitation tasks based on the visual resolution of the images. All analysts use the same standard scale, however all analysts do not always assign the same image into a single category. This is called epistemic or subjective uncertainty.[56]

The process of manual inspection of hyperspectral imagery is time consuming and unreliable due to variability in the analyst. The goal of this research is to

automatically determine areas that may or may not be interesting, with little intervention from an analyst, based on the results of a variety of tests to be developed. Traditionally, probability is used as a mathematical method to determine the likelihood of event occurrence, but it is not appropriate for use with epistemic uncertainty. Traditional probability theory has many assumptions, including knowledge of the probability events and the axiom of additivity.[56] Methods to combine evidence when little or no information on evaluating a probability is available, or the provided information is conflicting or ambiguous, are used in situations with epistemic uncertainty. These, called generalized uncertainty-based information theories or monotone measure theory, include the Dempster-Shafer theory of evidence and possibility theory.[56]

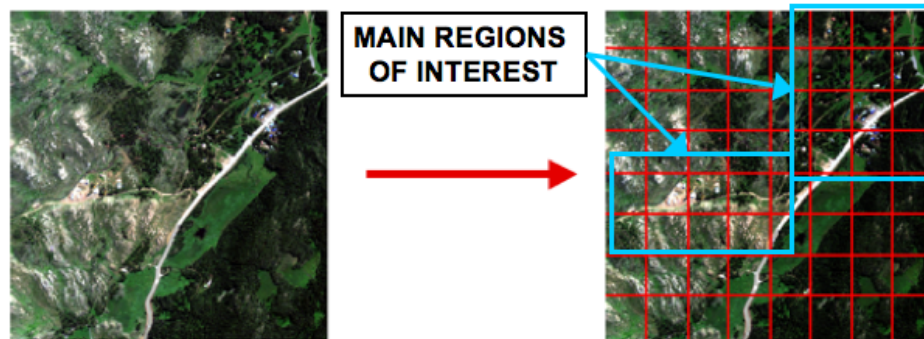


Figure 2.46: RGB image of the Cooke City, MT hyperspectral data cube and the same image notionally divided into 64 individual tiles for processing with the general region of interest highlighted.

The IDRISI software package for Geographical Information System (GIS) was used as a means to explore a variety of ways to combine the independent feature maps into a single map with two segments: “interesting” and “not interesting.” Traditional image classification involves assigning each pixel in an image to a particular class based on spectral similarity and results in a class map. Fuzzy classifiers do not assign each pixel to a single class; most methods assign the probability that the pixel belongs in each class in the image. The result is a probability map for each class in the image. For example, a pixel with a low probability in Class Map A is probably not a member of Class A. On the other hand, a pixel with a high probability likely is a member of Class A. If there is no dominant probability, the pixel may be a mixed pixel. Some fuzzy classifiers are based on hypothesis testing and instead

assign a belief and a plausibility score to each pixel, which represent the probability and the feasibility of the hypothesis being true.[56]




Figure 2.47: Example feature maps (a-e) and an RGB image (f) of the Cooke City, MT hyperspectral data cube and divided into 64 individual tiles for processing. Areas in blue contain known interest.

A HyMAP image of an area northeast of Cooke City, MT was chosen for this analysis as it is a predominantly natural (aka “uninteresting”) scene that also contains a few “interesting” areas including houses, roads, and construction. This image is tiled into 64 square tiles, each approximately 2,500-pixels. The areas of the image containing known interest are highlighted in Figure 2.46. In the IDL/ENVI environment, algorithms that estimate the number of spectral clusters, number of anomalies, point density dimension (PDD), PDDL, and TEP are applied to each tile. The raw results from each algorithm are used to make images where the brightness of the entire tile is the result from the algorithm, called feature images. For example, if the image tile is classified into four spectral clusters, then a value of 4 is assigned to the entire tile. The feature maps for anomaly detection, spectral clustering, PDD, PDDL, and TEP algorithms are shown in Figure 2.47, where the values are linearly scaled between 0 and 1. This is done to accommodate requirements for some of

IDRISI's classification routines. These feature maps show the relation between the relative feature value and the amount of interest. The anomaly (a) and cluster maps (b) correctly show high values along the road, over the houses, and over the construction site. The PDD map (c) shows the highest and lowest dimension estimates around the construction site. The PDTL (d) and TEP (e) maps also generally show higher feature scores over the known regions of interest.

The feature images are loaded into IDRISI as though they were a single image with multiple spectral bands. Simple boolean combination routines in IDRISI as well as regular and fuzzy classification routines are used to combine the metric images. The accuracy of the combination routines is determined by visual analysis and knowledge of the interesting locations in the Cooke City image.

Boolean Combination

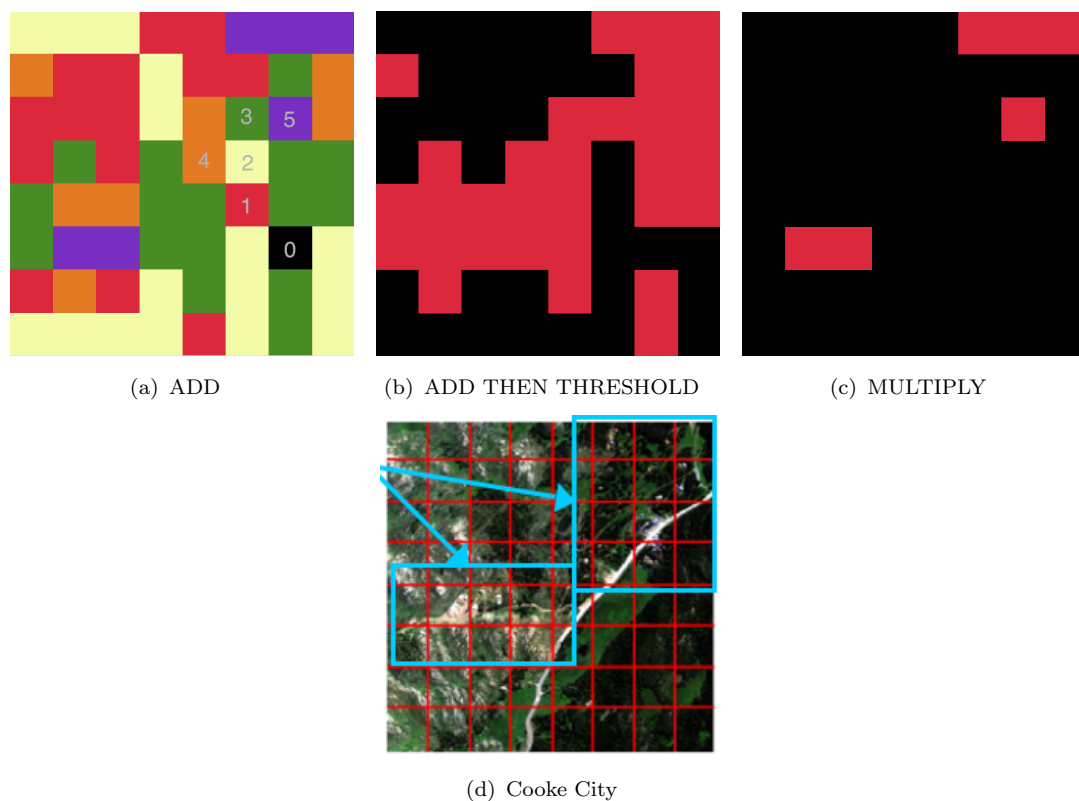


Figure 2.48: Combination images from traditional boolean combination in IDRISI.

Boolean combination of metric scores is the simplest way to achieve a single score for each image tile. This includes the AND and OR operations. These operations were used in two different ways. Typically, in order to use boolean combination, the images must be boolean. The scaled metric images from Figure 2.47 are thresholded to 1 if the value is above a certain score, and 0 otherwise. The image calculator in IDRISI was used to add and multiply the metric images together for the results shown in Figure 2.48. The multiplication image provided poor results, because if any metric value for a tile is zero, the result will be assigned a score of zero as well. The numbers in certain tiles in Figure 2.48(a) are the number of features with a score of one in the boolean images. For example, the black tile was below the threshold for all five features and the purple tiles were above the threshold for all five features. The addition image was thresholded so everything with a score of three or more, which means at least three features with a high score, is defined as interesting. When compared with the Cooke City image, the known areas of interest are all found.




Figure 2.49: Addition and multiplication images from scaled metric images.

Because the input images were scaled between 0 and 1, the scaled images could also be used as input for addition and multiplication combinations with a threshold set on the result. In both cases, a higher value would also indicate interest. Unlike true boolean combination, the resulting image will have non-integer values where a higher score should indicate more interest than a lower score. These results, shown in Figure 2.49, are similar to the boolean addition and thresholding. The multiplication image (b) was thresholded because most of the values were close to zero.

Supervised Classification

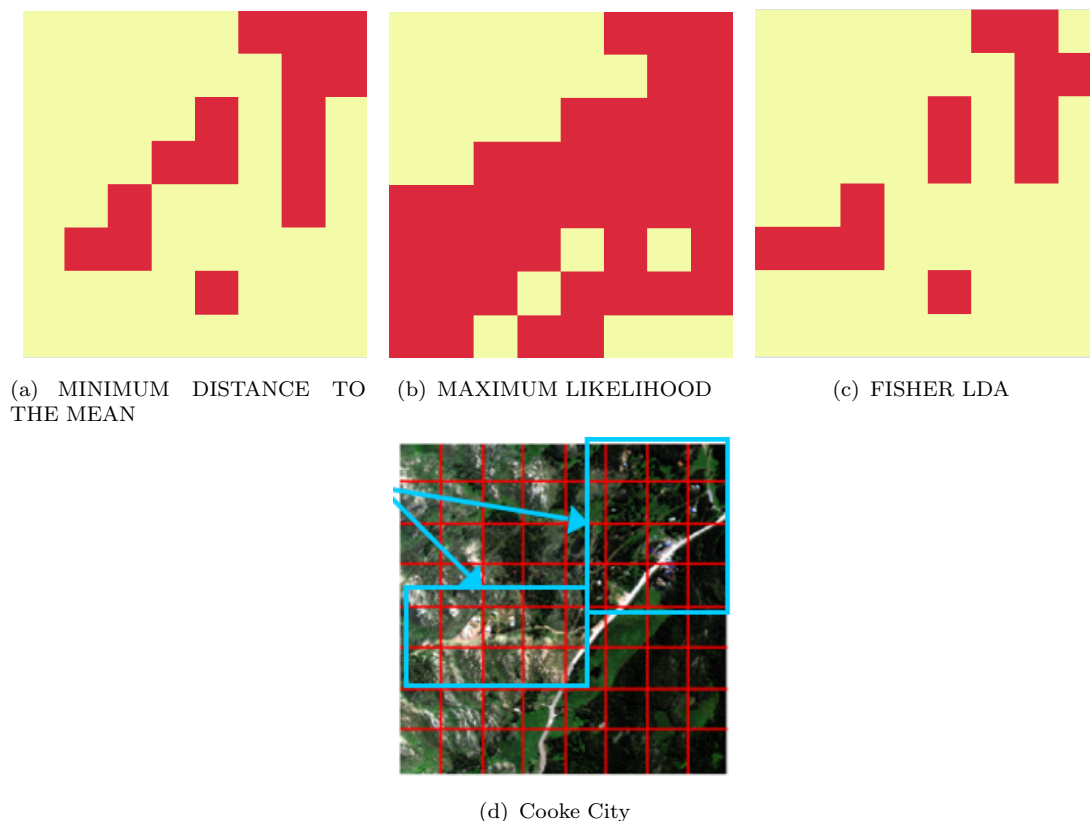


Figure 2.50: Results from using supervised classification methods as a way to combine the metric images.

Supervised classification is commonly used by, accepted by, and available to the hyperspectral community and for this reason was used as a method to combine the metric images into a single score tile. For training data, the tiles with the highest values in the scaled anomaly metric image (Figure 2.47(a)) were used for the interest class, and the areas with no anomalies were used for the no-interest class. All supervised classification methods available in IDRISI were used and the three most successful were chosen for inclusion. These basic supervised classification methods—minimum distance to the mean, maximum likelihood, and Fisher Linear Discriminant Analysis (LDA)—are shown in Figure 2.50. Maximum likelihood, the most popular of the three, performed the best. Many of the false alarms in this combination map are due to tiles that contain both trees and grass. Minimum

distance to the mean and Fisher LDA classified the majority of the known areas with interest as uninteresting. However, in realistic cases of this method, the regions that are interesting are unknown and there will not be available training data. An unsupervised method is more suited to this task.

Unsupervised Classification




Figure 2.51: Results from using unsupervised classification methods as a way to combine the metric images.

Unsupervised classification is also commonly used by analysts because they do not have to rely on training data for good performance of the algorithm. These methods are often not as robust, but are used because they let the computer decide what is a spectral cluster in the data. The results from the two methods chosen, histogram peak clustering (cluster) and isodata (isoclust), are shown in Figure 2.51. The two results are identical and are similar to the maximum likelihood result. In both cases, the areas of known interest are put into a single class, with a few false alarms on tiles with natural material clutter.

Dempster-Shafer with BELIEF Classification

The Dempster-Shafer theory calculates the probability of an event occurring by combining separate pieces of evidence. Interest contained in a tile is the desired event in this case. The possible hypotheses are then “interesting,” “not interesting,” and “either interesting or not interesting.” The evidence provided must support one of these hypotheses. In this analysis, all of our evidence, or metric images, are scaled such that a high number indicates interest. However, some of the evidence strongly

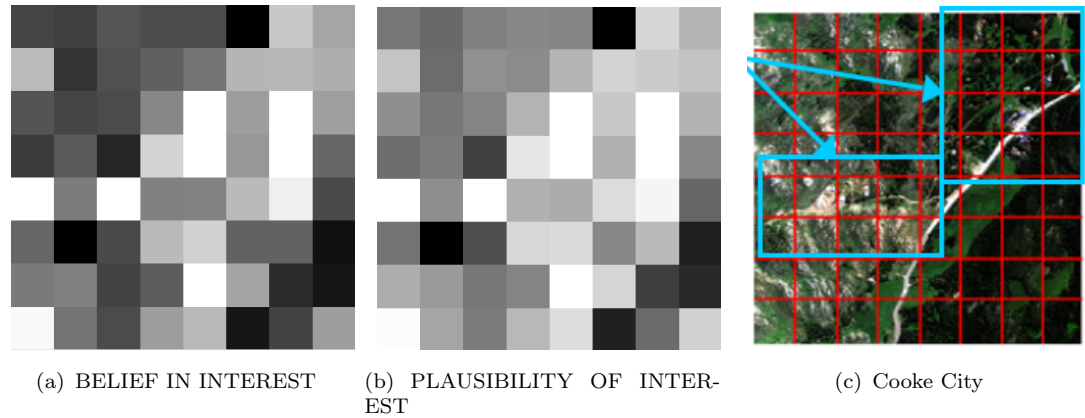


Figure 2.52: Results from using the Dempster-Shafer theory classification method as a way to combine the metric images.

supports interest, while other evidence only slightly supports interest. The output of BELIEF, the function that performs Dempster-Shafer combination in IDRISI, is a belief map and a plausibility map. The belief value indicates the amount of direct support of the event occurring and is the lower bound of the probability. The plausibility is one minus the amount of direct support against the event occurring and is the upper bound of the probability. In other words, it indicates how likely it is that the event could possibly occur if only considering the evidence against the event. In our case, if the belief and/or plausibility scores are high for a particular tile, then the tile may contain interest. These results, shown in Figure 2.52, in general, have brighter values in locations containing known interest and lower values in the regions that do not.

Combination Routine

For this analysis, CLUSTER and ISOCLUST produced an accurate map dividing the image area into “interesting” and “not interesting” segments in IDRISI with minimal user interaction. These results were reproduced in the IDL/ENVI implementation of k -Means and Isodata unsupervised spectral clustering. For these reasons, these are the algorithms that are used as a means of segmenting the multi-band feature image into interesting and uninteresting regions. A drawback of this method is that there is inherently no way to estimate the confidence of the segmentation map on a per tile basis. Even fuzzy classifiers, which do provide a probabilistic result, do not have an obvious way to accurately compute the confidence of the classification. The ideal combination method would be able to put a confidence or belief score for each

tile assignment. Instead, for this research, the number of “interest” clusters will be varied in order to divide the image into multiple levels of interest. For example, Figure 2.48(a) shows the image area divided into 6 segments where the tile score corresponds to the number of high valued spectral features. A tile that had a high value in all five features is more likely to contain interest than one that only had a high value in 4 features, and so on. This is not precisely a confidence score, but it provides a means for the analyst to prioritize based on likelihood of “interest” contained in the scene.

Chapter 3

RESULTS OF SPECTRAL FEATURE ESTIMATION

3.1 Dimension Estimation

3.1.1 PD Dimension Estimation on Hyperspectral Data

The hyperspectral data set used in this analysis was taken by the HyMap sensor flown over the Cooke City, MT area on July 4, 2006. This was done on five chosen regions of interest (ROIs). Samples of forest, grass, roads, secluded houses, urban areas, and a construction site were taken through visual segmentation. The full scene and example ROIs are shown in Figure 3.1. Multiple samples of each region were taken by varying the number of points in order to verify Camastra's data point requirement.[14] The PDD is estimated for all of these regions using the methodology described in Section 2.3.7.

Forest

The dimension was estimated for twenty forested regions of varying sample sizes based on the slope of dimension estimation plots. The mean slope was 5.8 and varies with a standard deviation of 1.19. There is no clear trend between the sample size and the dimension; however, the contents of the sample do affect the dimension. The reason for so much variation comes from the variability in tree cover across the scene. In a high resolution image that comes from the Cooke City flight line, Figure 3.2, patches of dirt, grass, and shadows can be seen between the trees. Some forest samples have significant amounts of non-tree pixels, which influence the dimension

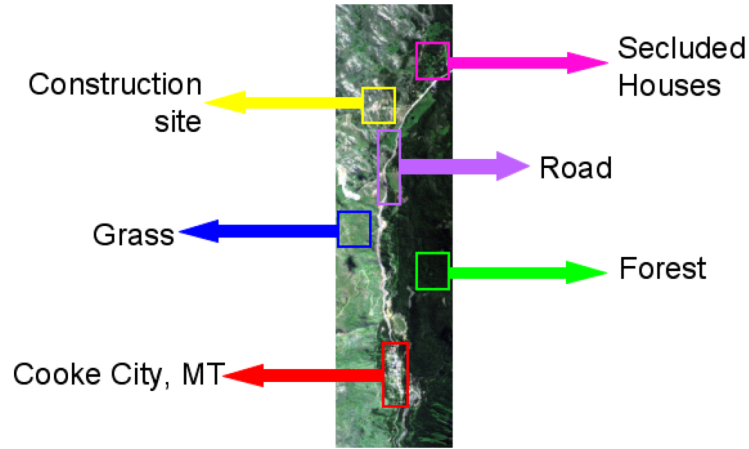


Figure 3.1: RGB display of a HyMAP flight line over Cooke City, MT with regions of interest used in analysis highlighted.




Figure 3.2: Visible image of example forested region (a) and dimension estimation plots for a 500 pixel (b) and a 30000 pixel (c) forested region.

estimation. These non-tree pixels manifest themselves in the plots primarily in small samples sizes that do not have enough tree pixels to overwhelm their presence. Figure 3.2 shows example dimension estimation plots for two independent samples of forest from the Cooke City image. The 500-pixel plot in (b) shows a short, steeply inclining portion, a secondary linear portion, and the plateau. The estimated dimension around the centroid pixel is 9.28 and corresponds to the slope of the steeply inclining portion. This region contains a significant patch of non-tree natural material. Figure 3.2 (c) shows the result from a 30000-pixel forest region with a longer and less steep primary incline, a short transition region, and a plateau. The estimated dimension from this sample is 5. This is a large sample with enough tree pixels to get a better estimation of the dimension of forest pixels. The estimated dimension as a function of sample size is shown in Figure 3.3. The sold black line




Figure 3.3: Bar chart of point density dimension estimates as a function of forest region size.

represents the average dimension estimate. With less than 1000 points in a sample, the dimension estimate is not reliable due to the number of data points required to accurately estimate the dimension. Once approximately 1000 data points are included, the dimension estimate begins to converge to 4.88. With the exception of a few anomalous very high estimates, the dimension estimate is approximately 5 regardless of sample size.

Grass

Grass and forest in some ways are very similar materials; they are both natural vegetation so they appear green to the naked eye. However, they do not have the same spectral signature or distribution in the hyperspace. Forest canopies are high above the ground and therefore cast shadows on lower levels of the trees. Grass grows very close to the ground and casts few shadows. Because it is so close to the ground, mixtures of soil and grass are common in any sample size.

The dimension estimation plots, examples shown in Figure 3.4, seem to have two general shapes, regardless of sample size. The first is similar to that of the trees—a steep incline followed by a transition to the plateau. The second plot appears as a




Figure 3.4: Dimension estimation plots of two grass samples from the Cooke City data, where (a) has 500 pixels and (b) has 3000 pixels.




Figure 3.5: Bar chart of point density dimension estimates as a function of grass region size.

curve to a plateau, with no distinct linear portion. This is likely due to the first example primarily containing grass and the second including a significant amount of soil. The estimated dimension as a function of sample size in Figure 3.5 shows that a distribution of these grass pixels lives in a lower dimensional space than the forest,

the average dimension being 3.78 with a standard deviation of 1.59. Unlike with the forest, the dimension estimate of grass continues to increase as the sample size increases even when the sample size requirement is met and exceeded. This is likely because the distribution of this grass contains a considerable amount of exposed soils and is not truly a single material cluster.

Roads

There is one main road that covers nearly the length of the entire flightline. This provides only enough pixels for one sample of pure road pixels. The dimension estimation plot for this sample is different than the forest and grass results as the presence of a transition region on the plot is visually negligible for the road while it was very apparent for the case of forest and occasionally present for grass. This indicates that pure road pixels were used in the road sample, while the grass and forest samples were often not pure samples. The magnitude of the primary slope is 3.11, significantly smaller than any estimate for the natural materials. There are not enough road pixels in this image to study how the dimension estimate changes as a function of sample size.




Figure 3.6: Dimension estimation plot for a road sample of approximately 700 pixels.

Secluded Houses

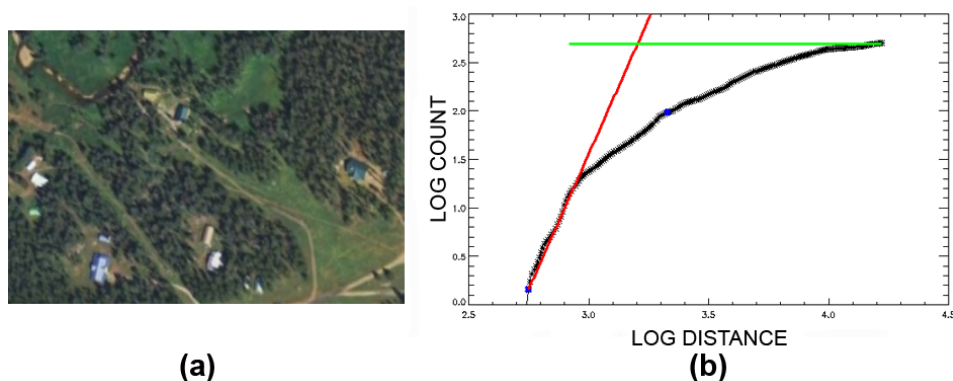


Figure 3.7: Visible image of region of secluded houses in Cooke City data (a) and dimension estimation plot (b) for a 500-pixel region.




Figure 3.8: Bar chart of point density dimension estimates as a function of region size containing secluded houses, trees, grass, and roads.

Figure 3.7 (a) shows a cropped image from the corresponding aerial image over Cooke City, MT. In one area of the scene there are a few houses scattered in a forest and grass region. In the high resolution image, it is clear that these houses are surrounded by grass, dirt roads, and forest. Because the regions including these houses must include forest and/or grass, it follows that the dimension of these regions would be similar to both trees and grass; however, because they have multiple materials in the plot, there should be a large transition region. The distribution on the plot shown in Figure 3.7 (b) confirms this. The plot appears to have the general shape of the forest plot in Figure 3.2 (b), which contained a patch of dirt in the sample. However, estimating the dimensionality of this cluster from the steeply inclining portion results in this region having a similar dimension to that of the grass samples. Figure 3.8 shows the estimated dimension as a function of region size. The dimension estimate is fairly constant across any sample size for this region. The mean dimension is 3.82 which is nearly identical to the mean dimension estimate for grass regions.

City




Figure 3.9: Visible image of region of Cooke City, MT (a) and dimension estimation plots for a 500 (b) and a 2000 (c) pixel regions.

Cooke City is the most urban area in this scene. It is apparent in the high resolution image that there is a significant amount of grass and road within the downtown area. Because this city has a significant amount of natural material in the scene, it may not be a good representation of a high clutter, fully urban area. Since the region does not contain only manmade materials, when ROI's were taken at a variety of sample sizes, some may have high amounts of grass while other are predominantly manmade material. This translates into significant variability in the log-log plots. A feature not seen in the previously introduced plots is wobbling

or steps of the incline, shown in Figure 3.9. This feature is likely an effect of the particular sample containing a variety of different material types in different amounts. Each bump may signify when a new cluster of data has been reached, causing the slope to increase again. Variability of this portion of the plot indicates different levels of clutter exist in the different samples. Even though the shape of the plot has a considerable amount of variability, the estimated dimension from both of these plots is approximately 4.0.

Construction

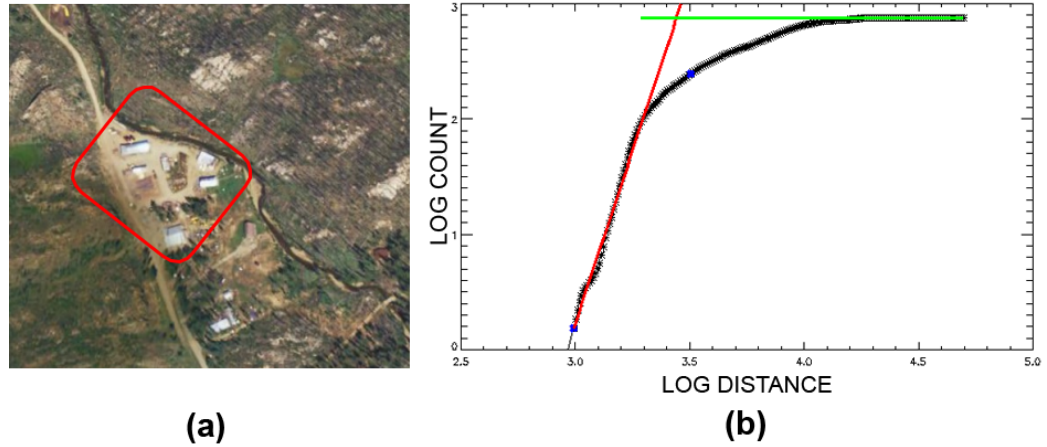


Figure 3.10: Visible image of construction site (a) and corresponding dimension estimation plot of the region (b).

In the original processing of this image, Topological Anomaly Detection[11] was executed. One area with a few hundred pixels was identified to be a very strong anomaly with TAD, indicating that a small group of pixels is far away from its closest background cluster. When looking at the entire scene, the anomaly could not be identified easily. When looking at high resolution data over the small anomalous area, it was determined to be a construction site; one can observe trucks, dumpsters, ground flattened to exposed dirt, and little to no vegetation. The high resolution image shows the construction site in Figure 3.10. Because this area is small compared to the other ROI's (≈ 700 pixels), only one sample was taken. The estimated slope of this region is 6.04. This plot has a steeper incline than the other materials and a fairly linear transition segment. It is clear that the distribution of the data in the

hyperspace is dramatically different from the previous examples in this ROI, free of vegetation and containing significant manmade phenomena and clutter.

Comparison




Figure 3.11: Average correlation and PCA dimension estimates for each region of interest.

The results for all regions are shown in Figure 3.11. The magnitude of the mean slope over all sample sizes is different for many of the materials. The road has the smallest inherent dimension, but it is also the relatively most pure material region used in this analysis. The road pixels were taken from the center of the single long road going through the image to prevent the inclusion of mixed pixels with dirt and metal often found on roadsides. The forest and grass appear to be single materials to the human eye, but in fact there is spectral variability caused by changing illumination conditions and exposed soil in the regions. The construction site has the highest estimated dimension and is one of the more complex regions in this image. It has a significant amount of manmade material and very little natural material; there are no grass or tree pixels in the region but there is some exposed dirt and rock. The city is also a very cluttered region, but because it has a significant amount of grass, the estimated dimension of the city is heavily influenced by the grass content. Grass, forest, and the secluded-houses region lie in a higher dimensional space than a pure material like road and a lower dimensional

space than the construction site due to the different degrees of material variability. These results indicate that regions of “interest” containing manmade material may have either an increased or decreased dimension estimate from the typical image tile estimate.


3.1.2 Dimension Estimation Feature Map

The feature maps in Figure 3.12 shows that the PDD for this particular image, from the HyMap sensor of Cooke City, MT, does not strongly correspond to the three regions of interest shown in the original image. The feature maps for full spectral resolution and resampled to WV2 spectral resolution at 30 square pixel tiles are shown. The previously introduced experiments did show a correlation between increased dimension and interest and this correlation is seen in other images. However, the red boxes corresponding to the city (Figure 3.13) and the construction site (Figure 3.14) show that these areas of manmade materials have a relatively constant dimension estimate. The wooded and foothill areas have a significant amount of speckle, with neighboring tiles having extremely high or low dimensions. In the top-center of the image, shown in a blue box, is an area of increased dimension estimate that corresponds to a high feature value for nearly every other feature developed. There is no ground truth of this area corresponding to the collect; however, high resolution imagery of the area is available on Google Earth. Figure 3.15 shows the contents of the blue box in the top center of the flightline containing two clearings in the woods as well as some exposed soil or rocks. For some applications, a newly created clearing in the woods may be “interesting.” In the RGB image, there are two regions that are brighter within the blue box—these may be small areas of trees that were knocked down in order to make room for a building. For the multispectral case, the PDD is not be a good indicator of the dimension due to the data being collected in a lower dimensional space.


The PDD feature map for a Hyperion scene of San Francisco is shown in Figure 3.16 for 40 pixel tiles. This result indicates the opposite of the relationship between increased dimension estimate and increased amount of interest. For large GSD data, an extremely high dimension or extremely low estimate corresponds to the water and the large grass fields, not the urban areas. The middle dimension estimates here indicate the urban regions in the scene. A large GSD corresponds to spatial averaging of the materials within a pixel, resulting in a less complex scene spectrally. This fact may make the PDD a better indicator of amount of “interest” for low spatial resolution data than for high spectral resolution data due to the phenomenology of the data.



(a) RGB IMAGE



(b) HYPERSPECTRAL FEATURE MAP



(c) MULTISPECTRAL FEATURE MAP

Figure 3.12: HyMap data of Cooke City, MT (a) and corresponding PDD feature map for the full hyperspectral scene (b) and resampled multispectral scene (c).

3.1.3 Additional Metrics from the PDP

Figure 3.17 shows the resulting feature maps from TEP estimation for the HyMap image at full spectral resolution and spectrally resampled to WV2 resolution. In both feature maps, the region corresponding to the wooded area on the top of the image displays brighter tiles, or larger TEP, than the regions below the road that are predominantly grass. The downtown area of the small city, the construction site, and the area of secluded buildings (Figure 3.18) are predominately dark tiles with a low TEP. This follows the observations made in Section 3.1.1 that manmade materials typically have a low TEP because of the uniformity in material produc-



Figure 3.13: High resolution image of downtown Cooke City, MT.



Figure 3.14: High resolution image of a construction site outside of Cooke City, MT.



Figure 3.15: High resolution image of a clearing in the woods outside of Cooke City, MT.




Figure 3.16: Hyperion data of San Francisco (a) and corresponding PDD feature map for the full hyperspectral scene (b).

tion of manmade materials. The region of secluded houses in the top left appears indistinguishable from the wooded areas—this is because the TEP metric is not sensitive to small or single pixel anomalies but instead is a measure of the regularity of the distribution. In the multispectral case, the secluded buildings also appear dark, similar to the city and the construction site. In both the hyper- and multispectral cases, there are a significant number of tiles that would be labeled “interesting” if low TEP, or dark tiles, are defined as interesting. The known areas of interest are all found, as well as some false alarms.

The observation between low TEP and amount of interest found for high spatial resolution data is not seen for low resolution Hyperion data in Figure 3.19. The TEP relationship between amount of interest is very similar to the PDD relationship for low spatial resolution data. The darkest and brightest tiles (or smallest and largest TEP values) correspond to water and the large fields while the grey tiles correspond to the the urban areas on land. In this feature map, the brightness over the urban area may indicate the amount of clutter within the tile. The brighter urban tiles

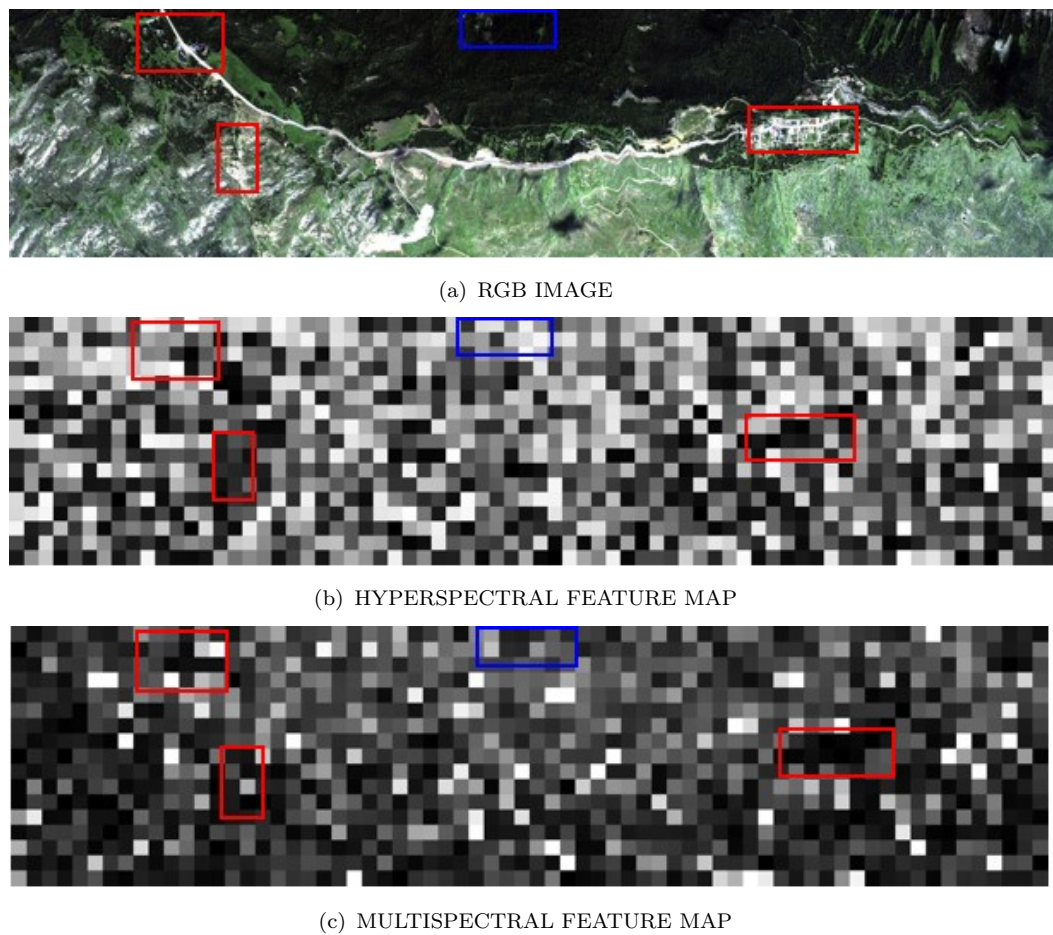


Figure 3.17: HyMap data of Cooke City, MT (a) and corresponding TEP feature map for the full hyperspectral scene (b) and resampled multispectral scene (c)



Figure 3.18: High resolution image of the area with secluded houses and buildings outside Cooke City, MT.




Figure 3.19: Hyperion data of San Francisco (a) and corresponding TEP feature map for the full hyperspectral scene (b).

shown in red correspond to a more cluttered region of the city than the tiles just below. However, some of the water appears as a similar shade of grey, meaning the TEP score is also similar. This water is not clean and clear—there are areas of different colors in the RGB, which relate to the depth and the constituents in the water.

The feature maps for PDDL estimation on HyMap data are shown in Figure 3.20. The PDDL is a strong indicator of the presence of manmade materials for high spatial resolution data, regardless of the spectral resolution. The wooded area is generally brighter with a longer PDDL than the rest of the image. The area following the main road is darker and has a very short PDDL. The grass and foothill regions fall in the middle. The city region contains many bright tiles with long PDDL's and one tile that is very dark—this is a tile that is mostly grass in the city. The construction site has a very long PDDL and appears as a single very bright tile surrounded by dark tiles. In this case, the construction site completely falls within in a single tile. The region with houses has significant variability of PDDL, or brightness. The dark regions near the houses are those which do not contain houses, while the lighter regions contain houses, trees, and roads.

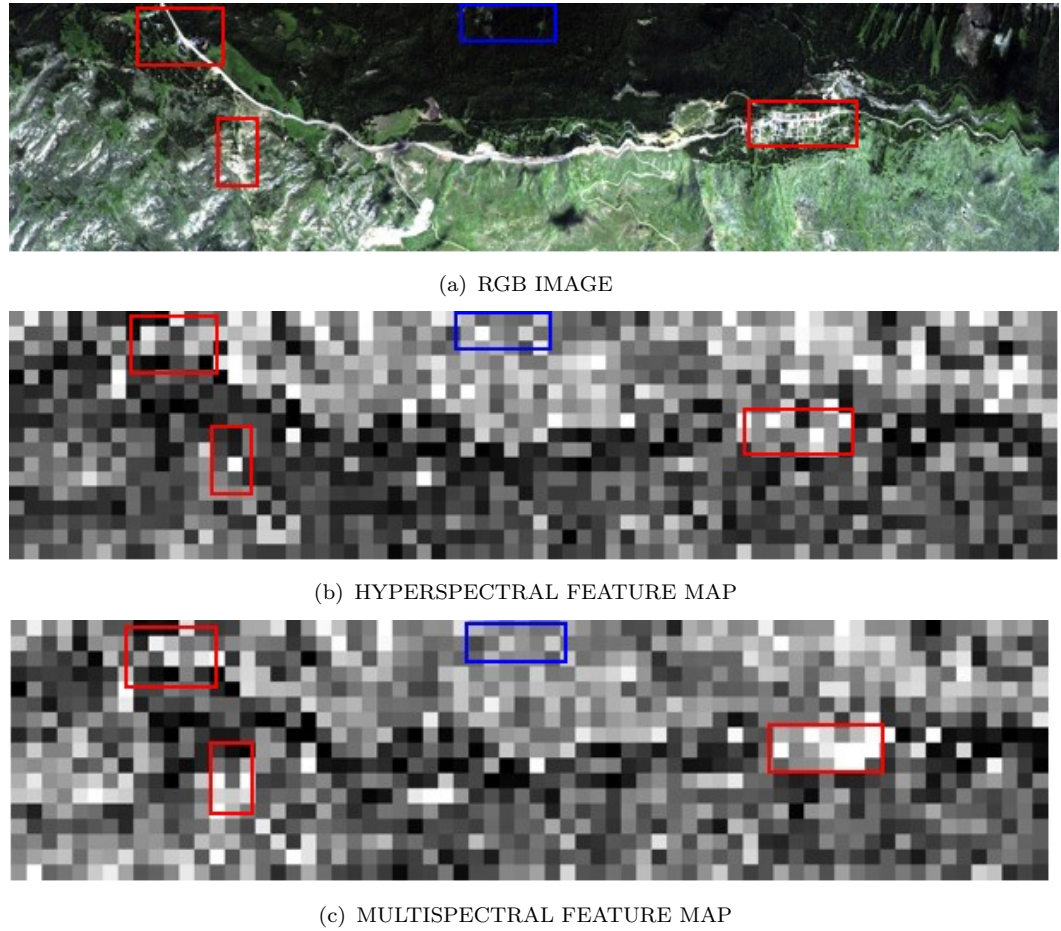


Figure 3.20: HyMap data of Cooke City, MT (a) and corresponding PDTL feature map for the full hyperspectral scene (b) and resampled multispectral scene (c).

The PDTL is not a good indicator of “interest” for low spatial resolution data as shown by the feature map for Hyperion data in Figure 3.21. Though all of the dark tiles do correspond to uninteresting areas, so do many of the bright tiles with a large PDTL value. This is due to reduced spectral complexity in low spatial resolution data. The PDTL is a measure of the very distant pixels from the center of the distribution—in some sense, anomalies. In 30m pixels, every pixel is a mixture of multiple materials and no pixel is very spectrally anomalous.




Figure 3.21: Hyperion data of San Francisco (a) and corresponding PDTL feature map for the full hyperspectral scene (b).

3.2 Gram Matrix Complexity Metrics

The Gram Matrix approach to spectral image analysis results in a plot, shown in Figure 2.28, from which three metrics can be estimated: the peak volume, the area under the volume curve, and the number of endmembers at which the volume is approximately zero. Each of these metrics correspond to the complexity of the spectral data within a particular image tile. A large value of any of these metrics corresponds to increased spectral diversity in terms of number of unique materials and spectral distance between these materials.

Figure 3.22 shows the number of endmembers feature map for Hymap data at two spectral resolutions. The number of endmembers required for the volume plot to go to zero is an estimate of the spanning or encapsulating dimension of the data because the endmembers are a geometric representation of the data. It is also an estimate of the number of unique and pure materials in the scene. The areas of the HyMap image that require the most number of endmembers all correspond to the city, construction site, and neighborhood of houses. The forest region in general needs the least number of endmembers because it is made of the least spectrally

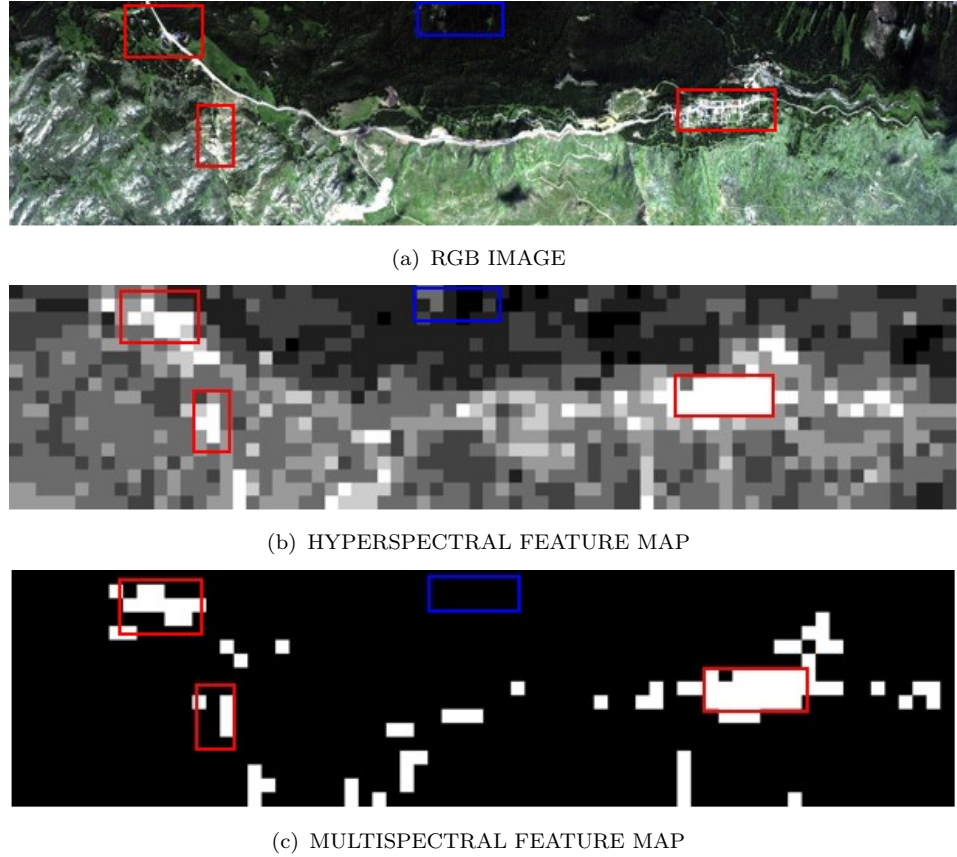


Figure 3.22: HyMap data of Cooke City, MT (a) and corresponding number of endmembers feature map for the full hyperspectral scene (b) and resampled multispectral scene (c).

diverse tiles. The grass and foothill regions are more spectrally complex than the forest and need more endmembers to span the space accordingly. The area of forest in the top center of the image; however, has an increased number of endmembers corresponding to the region likely containing interest. The number of endmembers needed to describe multispectral data accurately is significantly less than that of the hyperspectral case due to the number of spectral bands. The forest clearings in the top center of the image, shown in the blue box, do not appear strongly for the number of endmembers metric in the hyperspectral result or at all in the multispectral result. For the full hyperspectral image, tiles required between 6 and 15 endmembers for the volume plot to reach zero. In the 8-band multispectral case, the volume plot needed either 3 or 4 endmembers to reach zero.

Due to the material spectral averaging within large GSD pixels—no pixels are




Figure 3.23: RGB image of Hyperion data of San Francisco (a) and number of endmembers where the volume equals zero feature map (b).

pure pixels at 30m—endmember extraction and linear unmixing is more challenging on low spatial resolution imagery. The endmember feature map for the Hyperion data in Figure 3.23 shows that the tiles that fully consist of water need fewer endmembers than those which have any mixture of materials, urban or natural. For this image, the volume plot reach zero at either 3 or 4 endmembers for each tile.

Figure 3.24 shows the feature maps where brightness corresponds to peak volume of the Gram matrix plot for HyMap data. In this case, the city, roads, construction site, and neighborhood have the largest volumes. The city is brighter and more complex than the construction site, which is brighter and more complex than the neighborhood, and so on. If this were provided to an analyst, he/she would know that the forest region is the least spectrally diverse region. The possible region of interest in the top center of the forest does not appear significantly brighter in this metric as it did in many of the other feature maps. However, it is brighter than most of the forest region and indicates that it is likely more interesting than the surrounding tiles. The peak volume is an even stronger indicator of complexity and “interest” in a tile for the multispectral case, where the city, construction site, and

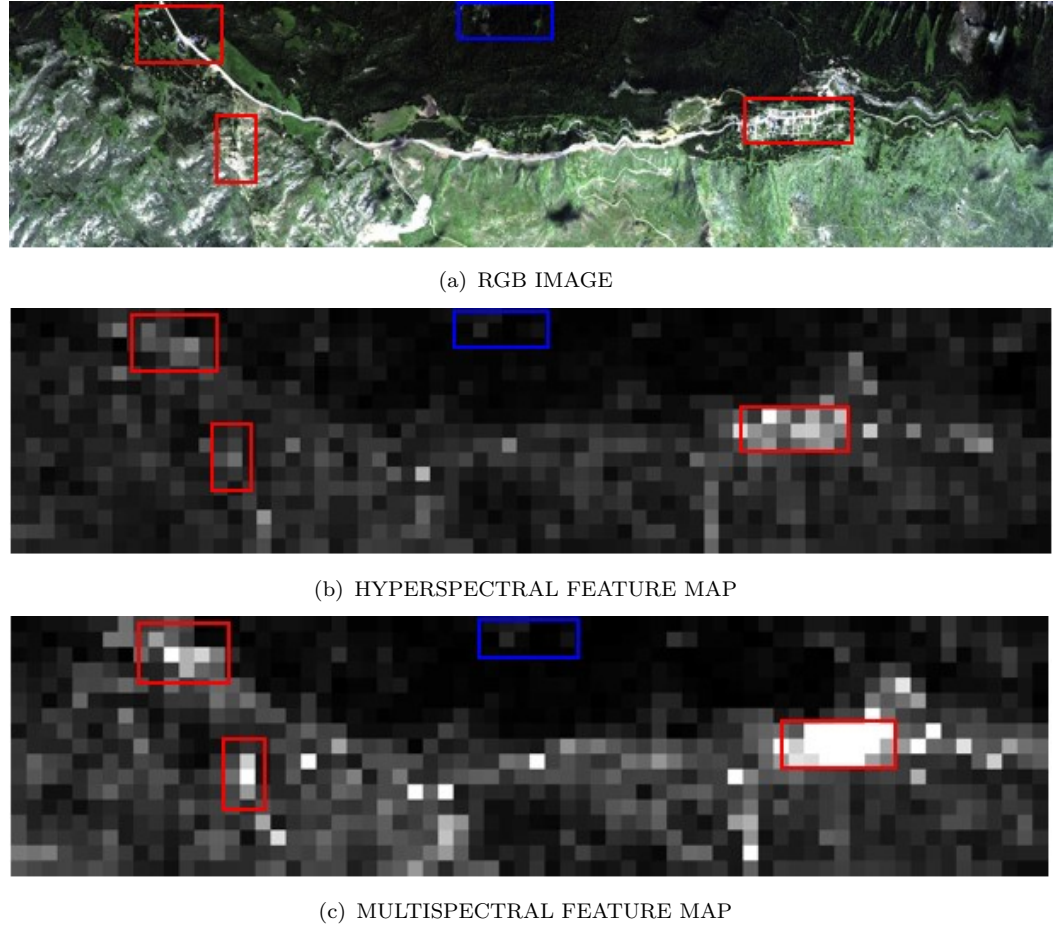


Figure 3.24: HyMap data of Cooke City, MT (a) and corresponding peak volume estimate feature map for the full hyperspectral scene (b) and resampled multispectral scene (c).

secluded buildings are all significantly brighter in the multispectral feature map.

The volume metric is also a strong indicator for low spatial resolution hyperspectral data, shown in Figure 3.25. The areas with water are dark with low peak volumes. The land regions of natural material near the coasts have slightly larger volumes and appear brighter in the feature map. The brightest tiles correspond to the urban areas with the largest peak volumes.

Figure 3.26 shows the area under the volume curve feature map for the HyMap data. In the hyperspectral case (b), this metric maps a significant portion of the image to “black” or “uninteresting”—more than the peak volume metric. In the multispectral case (c), the peak volume and area under the volume curve maps are





Figure 3.25: RGB image of Hyperion data of San Francisco (a) and peak volume feature map (b).

nearly identical. The three brightest regions once again correspond to the known areas of interest in the scene. The additional bright spots correspond to an intersection of the road and another region of the foothills that may contain something similar to the construction site. The peak volume and area under the volume curve are correlated—a tile with a large peak volume will also have a larger area—but not necessarily linearly, depending on the data.


Figure 3.27 shows the relationship between the peak volume and the area under the curve for hyperspectral data in Figures 3.24(b) and 3.26(b). The area increases much faster than the volume. One reason for this is the correlation between tile complexity and number of endmembers needed to describe the distribution, demonstrated in Figures 2.44 and 2.45. As the tile becomes more complex, the number of endmembers where the volume equals zero is larger which increases the domain of the area calculation. Additionally, the least complex tiles peak in volume around three endmembers before the volume drops significantly. As the complexity increases, the drop after the peak becomes less steep, demonstrated in Figure 2.45(d). With very complex tiles, the location of the peak moves from three to four endmembers, shown in the volume plot for the city tile in Figure 2.44(e). This relationship is not seen in



(a) RGB IMAGE



(b) HYPERSPECTRAL FEATURE MAP



(c) MULTISPECTRAL FEATURE MAP

Figure 3.26: HyMap data of Cooke City, MT (a) and corresponding area under the volume curve feature map for the full hyperspectral scene (b) and resampled multispectral scene (c).

multispectral data primarily because the number of endmembers needed to describe even complex data is lower for low spectral resolution data.

For large GSD data, the area metric map in Figure 3.28 is nearly identical to the peak volume feature map in Figure 3.25. The area metric does not increase faster for this data for the same reason that it did not for the multispectral case—there are fewer endmembers needed to describe the data and therefore physically less area in the Gram Matrix plot, so the area is driven by the peak volume estimate.




Figure 3.27: Peak volume as a function of area under the Gram Matrix showing the faster increase of the area metric.




Figure 3.28: RGB image of Hyperion data of San Francisco (a) and peak volume feature map (b).

3.3 Anomaly Detection

The TAD algorithm was applied to the region of the Cooke City image known to contain manmade activity, the construction site in Figure 3.14. This tile consists of a handful of construction vehicles of varying sizes and shapes amidst a natural material background, mostly grass and dirt. Figure 3.29 shows the RGB tile, the anomaly map, and the 2D projection of the graph built in the hyperspace. The

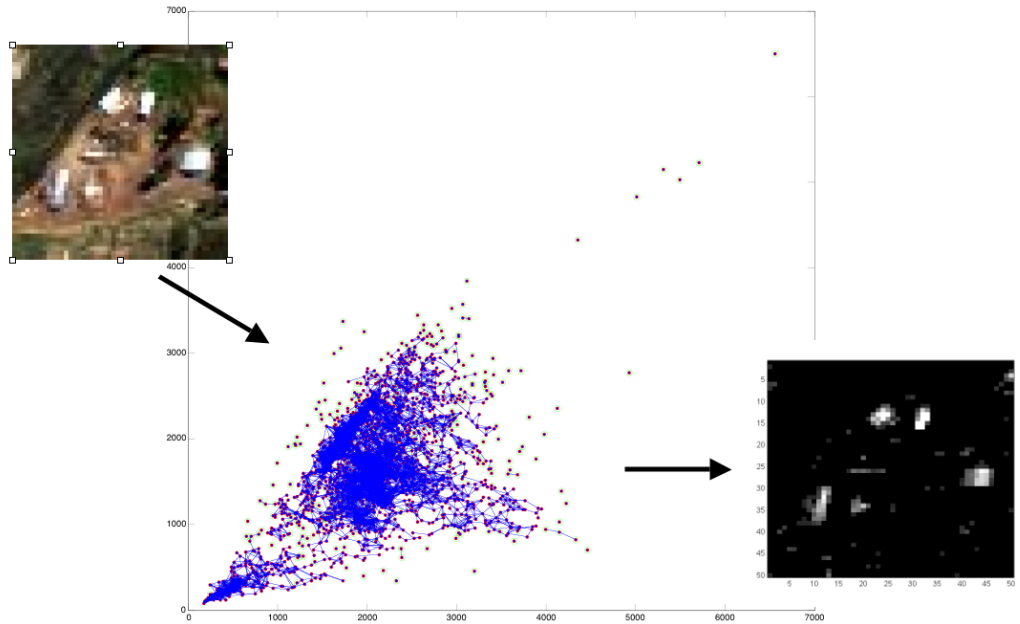


Figure 3.29: Anomalies detected using TAD (right) on the construction site in Cooke City, MT (left) showing the two-dimensional projection of the anomaly map in the hyperspace (center).

points in the graph connected by blue lines are labelled background components because they are spectrally “close” to one another. The points that are circled in red are not spectrally close to the background and are labelled anomalous. These anomalous points are displayed in varying shades of grey in the anomaly map. All of the large scale manmade objects, including trucks and trailers, are correctly found to be strong anomalies regardless of the number of pixels or orientation in the scene. There are some false alarms, mainly on regions of material transition between the soil and grass, but these are less anomalous (darker grey) than the trucks and trailers (bright grey).






Figure 3.30: Anomalies detected using TAD (right) on the forest radiance scene (left) showing the two dimensional projection of the anomaly map in the hyperspace (center).

A similar process was applied to the HYDICE forest scene shown in Figure 3.30. In the center graph, each individual background component in the hyperspace is circled in pink. In this scene, there are four disjoint background components, likely trees, grass, dirt, and road. The anomalies are shown circled in green. All of the tanks and targets are correctly labelled as strong anomalies. Even what looks to be a dead, brown-leafed tree is identified as anomalous in the forest. Material transition regions are again identified as anomalous. What defines a pixel as being “anomalous” is defined by the user with a closeness parameter; the results of TAD may be improved by changing this parameter.


3.3.1 TAD Feature Maps




(a) HSI TAD OUTPUT



(b) HSI FEATURE MAP



(c) MSI TAD OUTPUT



(d) MSI FEATURE MAP

Figure 3.31: (a) TAD anomaly image of HyMap data at full spectral resolution, (b) number of anomalous pixels feature map, (c) TAD anomaly image of HyMap data at WV2 spectral resolution, and (d) number of anomalous pixels feature map.

Figure 3.31 shows the TAD anomaly map and the number of anomalies per tile feature map on the full hyperspectral data. The tiles with roads, city, construction site, and houses all have a large number of anomalies compared to the bland grass and forested areas of the image. The area of the image in the blue box, which had a very high dimension estimate, a large TEP, and a long PDTL, also has a large number of anomalies, providing additional evidence that there is in fact something of interest in that area of the forest. The multispectral results in Figure 3.31(c-d) are very similar to the hyperspectral results. Some of the areas are less anomalous than their corresponding regions in the hyperspectral case. This degree of anomalousness is a measure of the distance from a point to a background cluster. For multispectral imagery this distance will typically be less because there are fewer bands used to calculate the distance. The water that flows along the road on the right side of the image and appears again in the top left corner are labeled anomalous by TAD because there is very little water in the scene.




Figure 3.32: TAD result on Hyperion data of San Francisco (a) and number of anomalous pixels (b).

The anomaly detection results on the San Francisco data set from the Hyperion sensor are shown in Figure 3.32. For this low spatial resolution data set, the number of anomalous pixels is not a good indicator of the amount of interest in the tile. As stated previously, because of the material variety in each 30m pixel, nothing in the image is truly anomalous. As a result, the brightest areas of the image are rated the most anomalous by TAD and drive the creation of the number of anomalous pixels map.


3.4 Spectral Clustering

In most cases, the Gradient Flow method for unsupervised spectral clustering is used to create a number of clusters feature map. This method is chosen because it is completely unsupervised, self-ending, there are few user inputs, and it is entirely data driven. No user guess at the number of classes in the image is required, but the user inputs a smoothing factor which relates to the number of classes Gradient Flow converges on. In some cases, other spectral clustering algorithms, including K-means or ISODATA, have been used both for comparison and due to memory issues with Gradient flow.


3.4.1 Feature Maps

Figure 3.33 shows the Gradient Flow output and number of spectral clusters feature map for HyMap data at full hyperspectral resolution. In this case, Gradient Flow converged on 21 unique classes for the full image. With this number of clusters, the classification map does not look very accurate—there are many classes in the forest, grass, and foothill regions when ideally this regions would be described in just a few classes. Isodata converged on even more classes, 25, but perhaps produced a slightly more visual appealing map. However, using the Gradient Flow result to create the number of clusters feature map produces a result where the city, construction site, road, and area with a few houses all contain a larger number of classes than the forest, grass, or foothills.


Figure 3.34 shows the results of Gradient Flow on the HyMap data spectrally resampled to WV2 spectral resolution. The image was divided into four regions processed individually in Gradient Flow due to computing resources. The colormaps used to display the images are slightly different, however Gradient Flow found a similar number of classes for each of the four regions. The boundaries of the four regions are visible in the number of clusters feature map but did not negatively



(a) GRADIENT FLOW OUTPUT



(b) ISODATA OUTPUT




(c) HYPERSPECTRAL FEATURE MAP


Figure 3.33: (a) Gradient Flow classification image, (b) ISODATA classification image, and (c) number of spectral clusters feature map for HyMap data of Cooke City, MT.

effect the interest segmentation. In each of the four regions, a brighter tile, or larger number of classes, corresponds to increased spectral material variability and “interest” within a tile.

Gradient Flow found 63 spectral classes, shown in Figure 3.35 in the Hyperion data set of San Francisco. The feature map shows the tiles that are entirely water are dark and have very few classes. The large uniform regions of grass on both coast lines show up as the brightest. The regions are smaller than the 50 pixel square tile size, so multiple of these regions fall within a tile. The urban regions are not classified into many classes because spectral averaging within 30m pixels reduces the spectral separability of the classes. Still, the urban region all have a similar number of classes, different from both the water and the large natural material fields.




(a) GRADIENT FLOW OUTPUT




(b) MULTISPECTRAL FEATURE MAP

Figure 3.34: Gradient Flow output of HyMap data at WV2 spectral resolution (a) and number of spectral clusters feature map (c).



(a) GRADIENT FLOW OUTPUT



(b) HYPERSPECTRAL FEATURE MAP

Figure 3.35: TAD result on Hyperion data of San Francisco (a) and number of anomalous pixels (b).


3.5 Interest Segmentation

The following sections contain the results of the entire interest segmentation process on a variety of images. For each data set, all eight metrics were estimated on a per tile basis. K-means and Isodata were both used to produce a variety of interest maps by “clustering” the multi-band feature image into two or three interest segments. In some cases, using less than 8 metrics produced what could be more usable results for an analyst. The results are presented in a variety of ways. The two or three color segmentation maps are used to produce maps containing the original image data. In some cases, the brightness of the RGB image tiles are altered based on the segment the tile belongs to. In others, the color is altered such that the “interest” segment is left in true color display and the “no interest” segment is displayed in color-infrared.


3.5.1 HyMap data

Figure 3.36 shows the interest segmentation map of the HyMap data over Cooke City using five spectral features as input: PDD, PDDL, TEP, number of anomalies (NA), and number of spectral clusters (NC). In this case, K-Means was used to cluster the multi-band feature image into 2 interest classes. Green represents “interesting” and red represents “not interesting.” The green segment contains the city, the main and side roads, the construction site, the neighborhood of houses, and the region in the forest that may contain something of interest. It also contains a few forest tiles that are probably not interesting, a clearing in the grass on the bottom left, and a clearing in the forest on the top right. However, though these regions do not contain something as interesting as a construction site, they are unique when compared to the rest of the scene in terms of spectral diversity and distribution size/shape. In this case, approximately 78% of the image is put into the “not interesting” segment, significantly reducing the amount of area an analyst must inspect.


Figure 3.37 shows two interest segmentation maps, both divided into three interest clusters. Figure 3.37(a) shows the result of using Isodata to cluster the five-band feature image into three interest segments. Red is still “not interesting” while blue and green are both “interesting.” The large blue region corresponds to the city. The blue tiles in the neighborhood are tiles that do contain at least one house. The construction site falls completely in a blue tile. In this case, the blue “interest” segment can be labeled a “high priority” or “high clutter” interest region. The green segment contains those tiles that are more complex than the trees and grass, but perhaps less spectrally diverse than the blue segment. The blue region in the top



(a) RGB INTEREST ON



(b) TWO-INTEREST CLASS SEGMENTATION MAP




(c) RGB INTEREST OFF

Figure 3.36: Interest segmentation map from five spectral features (b) with the interest class turned on in (a) and off in (c).

center correspond to the region in the trees that had a high score in each of the five original features.

Figure 3.37(c) shows the interest segmentation map from using Isodata to cluster the eight-band feature image into three interest classes. In this case, the city is the only region that falls into “high priority interest” segment—the rest of the interesting regions, including the neighborhood, construction site, and most of the road, are in the green interest segment. However, none of the known interesting regions are in the red or “no interest” segment. The area in the forest that likely contains something of interest is in the “no interest” class. This is due to the inclusion of the three Gram Matrix spectral features. In both the peak estimated volume and area under the volume curve maps (See Figures 3.24 and 3.26), this region was not given


a high score because it is less spectrally diverse than other tiles in the image. These metrics require spectral diversity of a significant number of pixels in order to locate regions of interest. This area in the forest is unique among its neighboring pixels, but since it is surrounded by dark, tall trees, it does not have sufficient spectral diversity from the rest of the tile to provide a large score in these metrics, and thus it appears dark. The large area search problem requires that all possible regions of interest are detected, at the expense of accepting some false alarms. In this case, using fewer metrics can provide this better than using all eight metrics because the three additional metrics do not find the forest clearing.



(a) INTEREST SEGMENT MAP FROM FIVE FEATURES




(b) RGB IMAGE



(c) INTEREST SEGMENT MAP FROM EIGHT FEATURES

Figure 3.37: Interest segmentation map from five spectral features (a), RGB image (b) and 8 spectral features (c).



(a) INTEREST SEGMENT MAP FROM FIVE FEATURES



(b) RGB IMAGE


Figure 3.38: Interest segmentation map from 5 spectral features for HyMap data re-sampled to WV2 spectral resolution(a) and RGB image (b).

The results of interest segmentation on HyMap data spectrally resampled to match WV2 resolution using the eight band feature image are similar to the hyper-spectral result using five feature bands. All of the known interest areas are found, including the city, construction site, secluded buildings, and roads. The area with the forest clearing is also assigned into the “interesting” segment. In the multispectral case, 80% of the image was assigned into the “not interesting” class, leaving 20% for the analyst to focus on.

3.5.2 HYDICE

Urban Scene


Figures 3.39 (a-d) show the results of the feature clustering on HYDICE data into two and three interest segments. In both cases, approximately one third of the image area is put into a single interest cluster representing the “uninteresting” areas and two thirds of the image are possibly “interesting.” These regions do correspond to the known areas containing manmade materials or “interest.” In Figure 3.39(c-d), there is a third interest cluster that has regions labeled as both “interesting” and “not interesting” in Figure 3.39(a-b). This interest cluster, shown in blue in



(a) INTEREST SEGMENT MAP FROM FIVE FEATURES



(b) IMAGE MAP



(c) INTEREST SEGMENT MAP FROM FIVE FEATURES



(d) IMAGE MAP

Figure 3.39: Interest segmentation maps from five spectral features in two interest segments (a) and three interest segments (c) and visual maps (b) and (d) where true color (RGB) regions are “interesting” and false color (CIR) are “not interesting.”

Figure 3.39(c) and displayed in true color with altered brightness in Figure 3.39(d), primarily contains the area of the image containing the individual cars in the parking lot and road, or the small scale regions of interest; the larger interest cluster, colored green in Figure 3.39(c) and displayed in true color at full brightness in Figure 3.39(d), contains the large scale regions of interest including houses and other buildings.

Forest Scenes

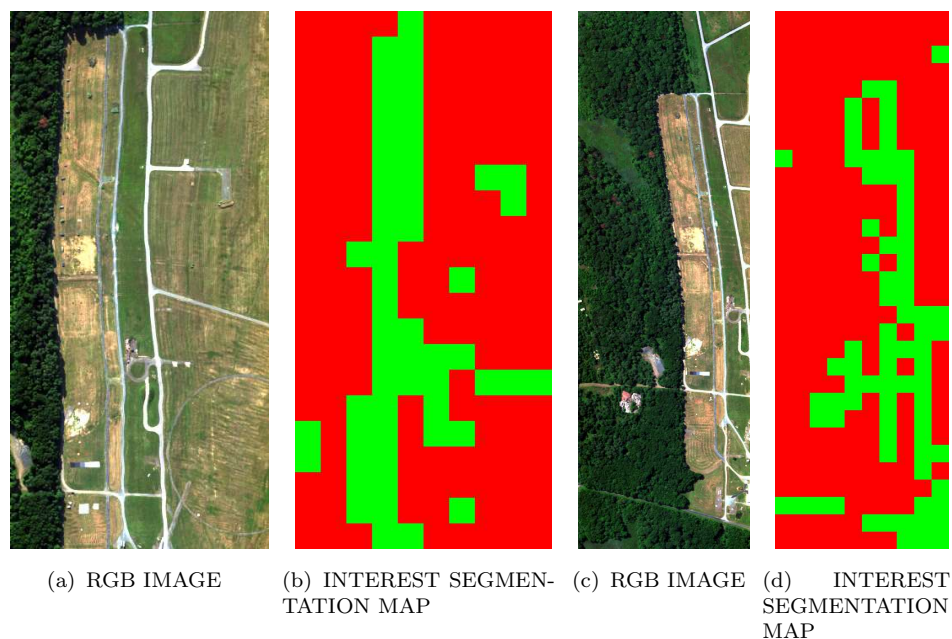


Figure 3.40: RGB image with exposed (a) and hidden (b) targets and corresponding interest segmentation maps of HYDICE forest imagery.

Figure 3.40 shows two HYDICE scenes of a predominantly natural scene collected a few days apart. In Figure 3.40(a), a group of camouflaged targets, tarps, and vehicles are exposed in the grass. In Figure 3.40(c), many of these targets have been moved into the treeline. In both images, the tiles containing the camouflaged nets, calibration targets, and some roads are put into the “interest” segment. The clearing with buildings on the right side of both images is also correctly identified as interesting. However, the targets that are small scale ROIs, whether exposed or hidden, are not found to be “interesting” by enough of the metrics to be included in the “interest” segment. For both images, 26% of the image was labeled “interesting” by this methodology.

3.5.3 AVIRIS

Figure 3.41 shows the true-color image of an area of the Rochester, NY coastline on Lake Ontario. Approximately one quarter of this image is “uninteresting” water—the rest is varying degrees of urban clutter. The large lake water region is correctly

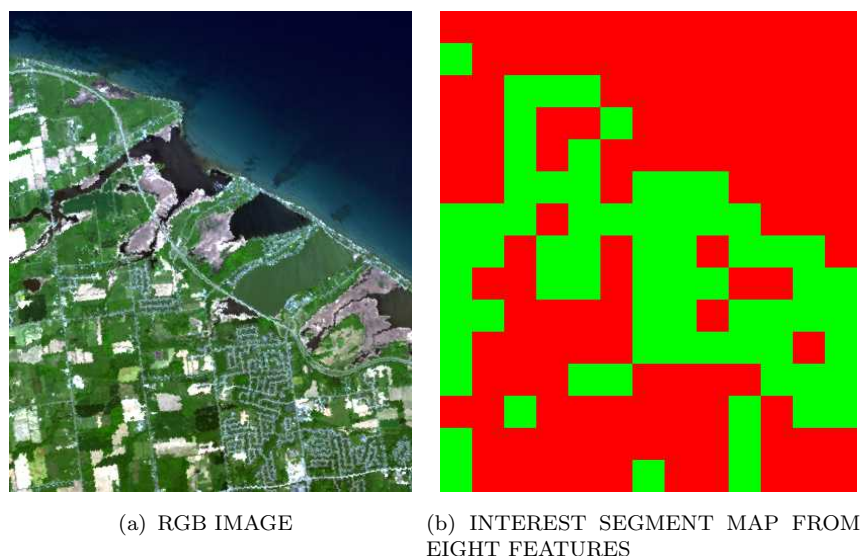


Figure 3.41: RGB image of AVIRIS data of Rochester, NY near Lake Ontario (a) and corresponding interest segmentation map (b).

put into a “no interest” segment. In many cases, tiles that fall entirely in the near-coast ponds are also correctly labeled “not interesting.” Tiles containing the coastline itself are labeled “interesting.” Some of the most cluttered land region is labeled “interesting,” approximately 35%, but much of the land is put into the same segment as the water. This is likely due to the effect of large GSD phenomenology on the presented metrics, in particular TAD and PDTL.




Figure 3.42: RGB image of AVIRIS data of Moffett Field, CA (a), corresponding interest segmentation map from one iteration (b), and two iterations (c).

Figure 3.42 shows an AVIRIS image of Moffett Field, CA. This image contains a small city in the lower right corner. The rest of the image is predominantly not interesting, but contains a significant amount of natural material variability and clutter, providing a challenging test for the interest segmentation process. Figure 3.42(b) shows the interest map from one iteration of K-means dividing the eight-band feature image into two interest segments. The urban area is correctly put into a separate “interesting” segment from most of the rest of the image. Ten additional tiles are also put into this “interest” segment. AVIRIS collects approximately 20m GSD data, so it is difficult to know whether these ten tiles are falsely or correctly labelled “interesting.” In this result, 20% of the image was labeled “interesting.” K-means was used a second time to segment the data into two interest segments, this time with two iterations, for Figure 3.42. In this result, seven of the ten individual tiles are removed from the “interest” segment, providing what is likely a more accurate interest segment map for an analyst with only 10% of the image labeled “interesting.”

3.5.4 Hyperion

The interest segmentation result for the Hyperion imagery of San Francisco, CA is shown in Figure 3.43. Approximately ten pure water tiles area assigned into the red “no interest” segment, as well as 3 other tiles. The two of these red tiles in the top half of the image are not water, but primarily consist of a single landcover type. In this result, most of the image, nearly 75%, is labeled interesting in either the green or blue segments. The four blue tiles are put into a separate segment, driven primarily by the high values in the anomaly map. These tiles correspond to the brightest points in the image, which were the most anomalous in the TAD result. The largest segment, colored green, contains every other tile. These tiles all consist of a mixture of materials within the tiles, giving increased scores for many of the features measured.

The results on a single flightline of Hyperion data over Oaxaca Valley, Mexico are shown split into three regions in Figure 3.44. The imagery was processed whole and only divided into three regions for display. K-means was used to divide the seven-band feature image into two interest segments. The number of anomalous pixels feature map was not used for this analysis due to multiple bad band lines of pixels in the image—nearly the whole image was anomalous. Approximately 50% of the full flightline is in the “interest” segment, shown in red. This is not as high a percentage as some of the other imagery used for this analysis, however this data has 30m pixels. This translates to nearly 27,000km labeled “not interesting”




Figure 3.43: Interest segmentation map (b) of Hyperion data of San Francisco, CA (a).

or lower priority on analysts task, which is a significant reduction. The interest segmentation process works best in the mountainous region, shown in Figures 3.44(a-b). This region is mostly tree covered, but does have some roads, small villages, buildings, and farms scattered around. When visually inspected, the red tiles in this region all corresponding to something unusual from the tree background. The arid region, shown in Figures 3.44(c-d), is the most spectrally complex region of the image. It contains a small urban area surrounded by farmland on the right side of the image. This region is correctly labelled “interesting.” The rest of the arid portion of the flightline is spectrally complex due to multiple types of exposed rock, varying altitudes, and some natural material. In this region, there are many tiles that may be falsely assigned into the red “interest” segment, but without ground truth or high resolution imagery, there is no way to know for sure. All of the small urban areas were correctly labeled “interesting.” The tropic region also has many scattered small towns, farms, and roads which are correctly assigned into the “interest” segment. Throughout the entire image, there are many roads winding up and down the mountains. From visual inspection, there are many areas along the road where buildings, farms, or small villages are set up. Examples of these regions are shown in Figure 3.45. Most of the red “interesting” tiles in this imagery corresponds to these roads.

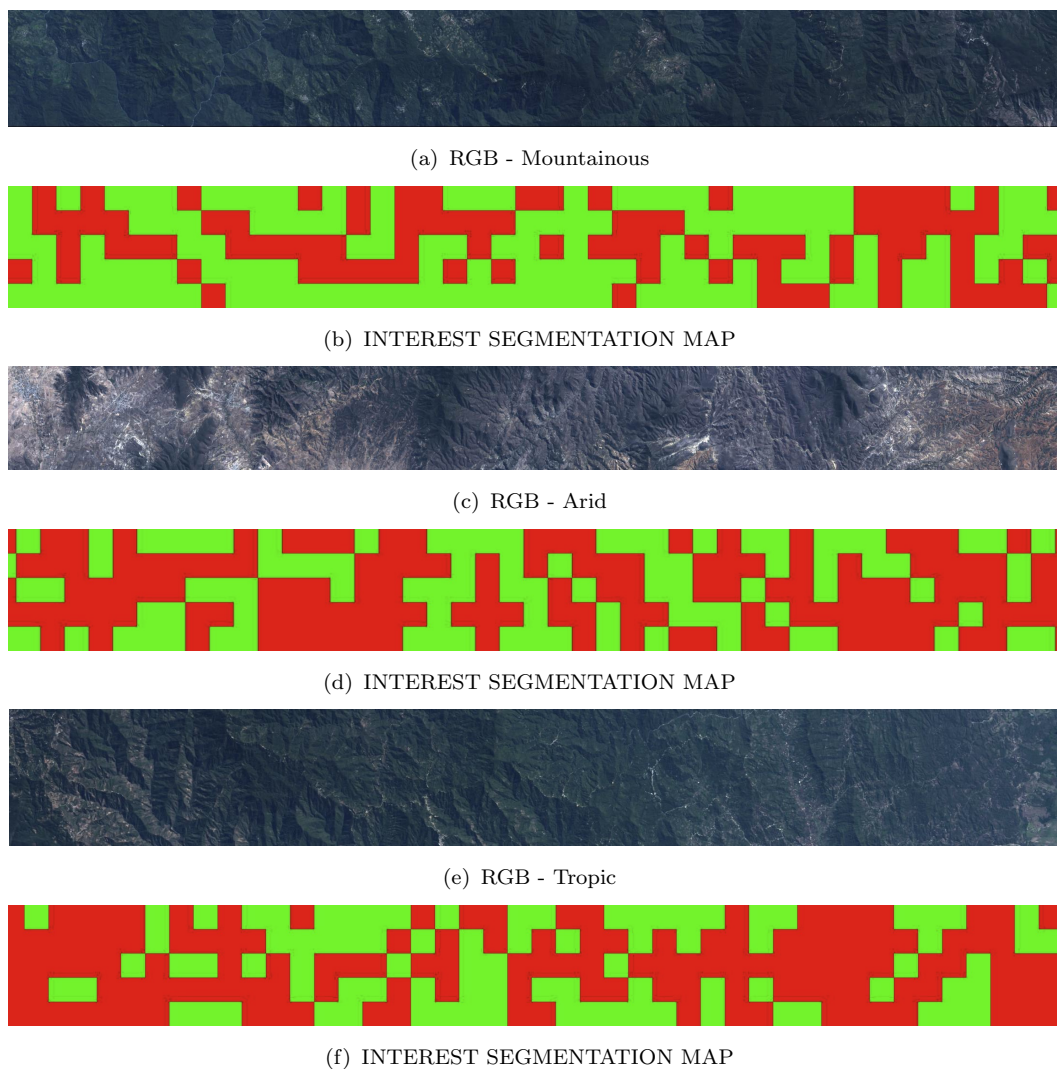


Figure 3.44: RGB images (a-c-e) and interest segmentation maps (b-d-f) for three regions of Hyperion imagery of Oaxaca Valley, Mexico.

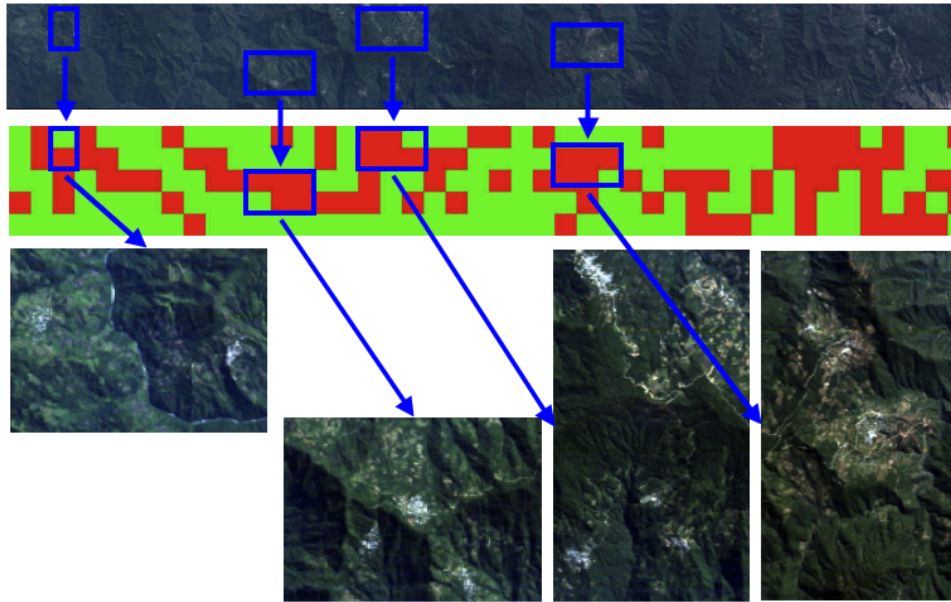


Figure 3.45: Interest segmentation map for Hyperion data of Oaxaca Valley, Mexico and correctly identified regions of “interesting” manmade activity.

3.5.5 True Multispectral: WV2 and Quickbird

Shortly after launch, DigitaGlobe released some sample imagery from the new WorldView2 satellite, both raw and pansharpened. A test region in a wooded area was used for interest segmentation. All of these maps were created by using K-means to combine the eight-band feature image. Figure 3.46(a-b) shows the interest segmentation results for the raw data at 30×30 pixel tiles, which corresponds to a $75\text{m} \times 75\text{m}$ tiles. Interest segmentation correctly puts most of the image, the trees, into a “no interest” segment, while the clearing with buildings and roads is in a separate “interest” segment. One tile containing the exposed river or road on the right is also put into the “interest” segment. The comparable result for the pansharpened imagery in Figure 3.46(c-d) uses 120 pixel tiles to cover the same area as the 30 pixel tiles in the raw imagery. The interest segmentation map is very similar—the clearing with buildings is found, but the tile with the road or river is not.

The benefit of using the pansharpened imagery is shown in Figure 3.46(e-f), where 40 pixel tiles are used to cover approximately $19\text{m} \times 19\text{m}$ square regions on the ground. In this case, the particular regions of “interest” are put into a region separate from the forest. In the large clearing, individual regions that have no

RAW WV2: 30 pixel/75m tiles



(a)

(b)


PANSHARPENED WV2: 120 pixel/75m tiles



(c)

(d)

PANSHARPENED WV2: 40 pixel/ 19m tiles



(e)


(f)

Figure 3.46: Interest segmentation maps (a,c,e) from five spectral features in two interest segments and visual maps (b,d,f) where true color (RGB) regions are “interesting” and false color (CIR) are “not interesting.”

manmade material are put into the “no interest” segment with the trees of the rest of the image, while all of the tiles containing road and buildings are kept in the “interest” segment. In this case, many of the tiles on the right side of the image that shown exposed road and/or river through the tree cover are also put in the “interest” segment, which are completely missed with larger tiles. This result shows that using pansharpened multispectral data can produce improved results over the raw data when an effective pansharpening algorithm is used.



(a) RGB IMAGE



(b) INTEREST SEGMENTATION MAP

Figure 3.47: RGB image (a) and Interest segmentation map for Quickbird imagery of an near the Esperanza forest fire containing a large agricultural and manmade region.

The large area scene shown in Figure 2.7 was collected by the Quickbird sensor over the site of the Esperanza forest fire near Cabazon, Ca in 2006. Three large sub-regions, each covering nearly 40km², with known small scale regions of interest were

used for interest segmentation. Figure 3.47 shows an RGB image of one subregion containing a large agricultural and manmade area. This data has 2.5m GSD and the interest segmentation map was made using all eight spectral features measured on 50 pixel tiles. At first glance, much of this image looks uninteresting. Approximately 88% of the image is assigned into the red, “no interest” segment. The small urban and agriculture areas are put into an “interest” segment separate from much of the rest of the image. These regions are obviously different from most of the background from visual inspection.





Figure 3.48: High resolution tiles corresponding to “interesting” tiles in the interest segmentation map in Figure 3.47.

Approximately 30 additional tiles are also put into the interest segment. Sub-meter GSD panchromatic data were taken coincidentally. Through inspection of the high resolution data, some of these “interesting” tiles are likely not interesting. Many do not appear interesting in the true-color image, but further inspection of the high resolution data shows otherwise. Figure 3.48 shows the corresponding tiles


in the high resolution image for six of the regions labelled “interesting” from the interest segmentation process. The lone green tile near the urban area and the agriculture region appear to be outside the boundaries of the heavily populated area. The high resolution imagery in Figure 3.48(a) shows the building that is not easily visible from the multispectral data. Figures 3.48(b-f) also show what appear to be small buildings not obvious in the multispectral data. Figure 3.48(d) is a fairly uninteresting tile with a few extremely bright pixels near the center. These are very anomalous and are likely driving the assignment of the corresponding tiles into the “interest class.” These pixels may be rated very anomalous because they are extremely bright or spectrally different, but with only high resolution panchromatic data there is no way to determine which is the cause.

Another region of the Esperanza collect is shown in Figure 3.49. The lower right corner of this subset contains a suburban neighborhood, but much of this image looks uninteresting and uninhabited beyond the network of roads through the mountains. Interest segmentation assigns approximately 83% of this image into the “no interest” segment. The suburban area is correctly segmented into the “interest” region along with many additional individual and clusters of tiles. The road network itself is not labeled as “interesting,” however many areas along the road contain cars or buildings that are correctly assigned into the “interest” segment. Example high resolution images of some of the “interesting” tiles are shown in Figure 3.50.

The region shown in Figure 3.51 covers the area where the wild fire originated near a small suburban and rural farming area. This scene also contains the known areas where the rescue teams were headquartered, including buildings and firetrucks, and suburban neighborhoods. Approximately 70% of this image was assigned into the “no interest” segment, significantly less area than either of the previously introduced subsets of this data set. This is for many reasons. Firstly, this image contains more large scale regions of interest including three large neighborhoods. These are correctly assigned into the green “interest” segment. The rest of the “interest” segment appears to be scattered around the entire image. Upon inspection of the high resolution imagery, nearly all of these tiles contain some amount of manmade activity, including buildings, firetrucks, cars, and power line towers. Examples of this are shown in Figure 3.52.



(a) RGB IMAGE



(b) INTEREST SEGMENTATION MAP

Figure 3.49: RGB image (a) and Interest segmentation map for Quickbird imagery of an area near the Esperanza forest fire containing a small suburban area and many large areas with no activity.





Figure 3.50: High resolution tiles corresponding to “interesting” tiles in the interest segmentation map in Figure 3.49.



(a) RGB IMAGE



(b) INTEREST SEGMENTATION MAP

Figure 3.51: RGB image (a) and Interest segmentation map for Quickbird imagery of an near the origin of the Esperanza forest fire containing multiple large and small scale regions of interest.




Figure 3.52: High resolution tiles corresponding to “interesting” tiles in the interest segmentation map in Figure 3.51.

Chapter 4

CHANGE DETECTION USING PDTL

4.1 Point Density Change Detection

The short study in Schlamm, *et al.* (2009b) demonstrated how the PDD, PDTL, and TEP change as more manmade pixels are included in the scene.[51] A large hyperspectral scene containing varying types of natural landcover and a small urban area were selected. Three areas of different natural material background were used as base tiles. From the same image, the manmade pixel spectra were extracted and collected into a library. The manmade pixels were randomly selected to replace random natural material pixels in the tiles, for an increasing number of pixels. The tail length as a function of the percentage of manmade materials is shown in Figure 2.42. The trend for each tile is identical. When just a few manmade pixels are added to the tiles (less than 2-3%), they are at a large distance from the center of the distribution, and therefore lengthen the tail of the point density plot. As more manmade pixels are added, the center of the distribution is pulled towards the manmade pixels and the tail length gradually shortens. The difference in PDTL between the raw hyperspectral tiles and the same tiles with injected pixels of manmade materials is the foundation for the change detection algorithm described below.

The change detection process from point density analysis is straightforward and is outlined in Figure 4.1. The two images must be registered: any user preferred method will suffice. Subpixel registration is not explicitly required and was not used in the results presented below. Because this method performs the change detection on tiles and not individual pixels, it is more robust to misregistration errors. A few

more or less pixels of the majority material in the plot will not affect the PDTL, but the corresponding tiles must approximately cover the same region. Overlapping regions from the two images of the same spatial dimensions should be selected out of the registered scenes for the subsequent processing.

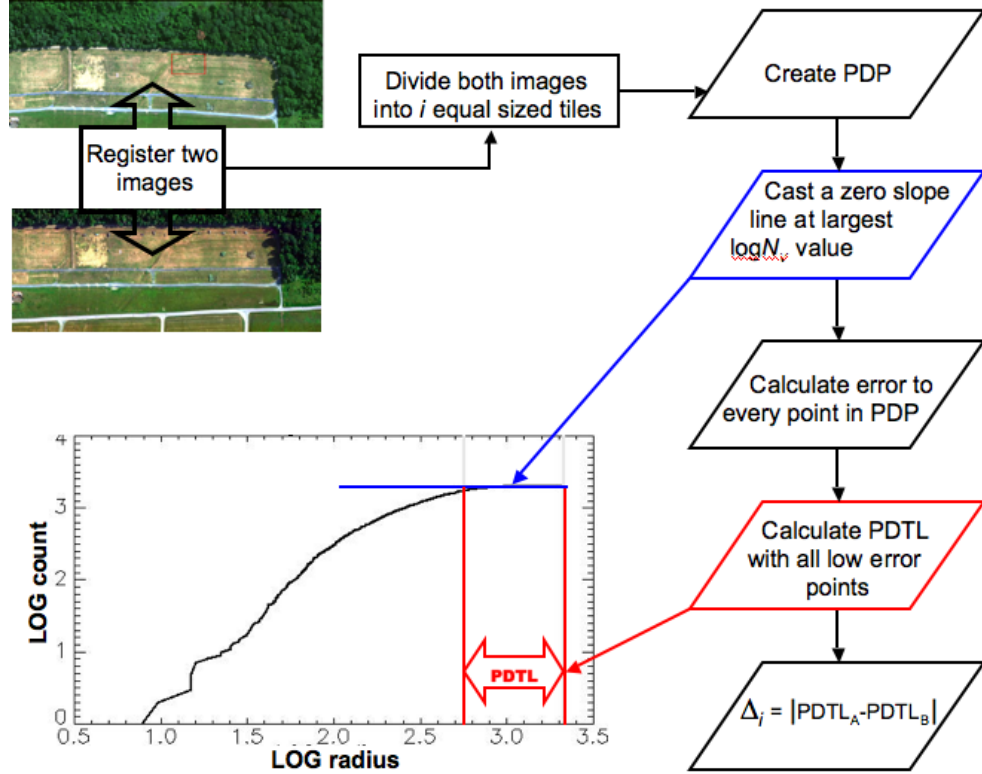


Figure 4.1: Change detection via PCTL estimation process.

The two images are divided into equal sized tiles, typically containing 1000-2000 pixels. The regions that do not fall into tiles are discarded. In this processing, square tiles were used, however rectangular tiles will also work. The data for the PCTL are generated for each tile as described in Schlamm, *et al.* (2009)[53]. We create a zero-slope line at the largest $\log N$ value and calculate the error between every data point and the zero slope line. The data points at small radii will have high errors while the data points at large radii will have low error. The error is thresholded and only those points in the PCTL with error less than the selected threshold are identified as the tail. By ignoring the incline portion of the PCTL, we are suppressing the common background between the two images. The PCTL is the approximate

distance, in radius units, from the first to last points identified as being in the tail. This is repeated for every tile in both images.

After the PDTLs are found for each tile in both images, the Δ metric is computed as

$$\Delta_i = |PDTL_A - PDTL_B| \quad (4.1)$$

where i is the index of corresponding tiles in image A and B . The difference in PDTL for an image tile of the same region before and after change is illustrated in Figure 4.2. When a target is present, the PDTL is significantly longer than in the target absent case.

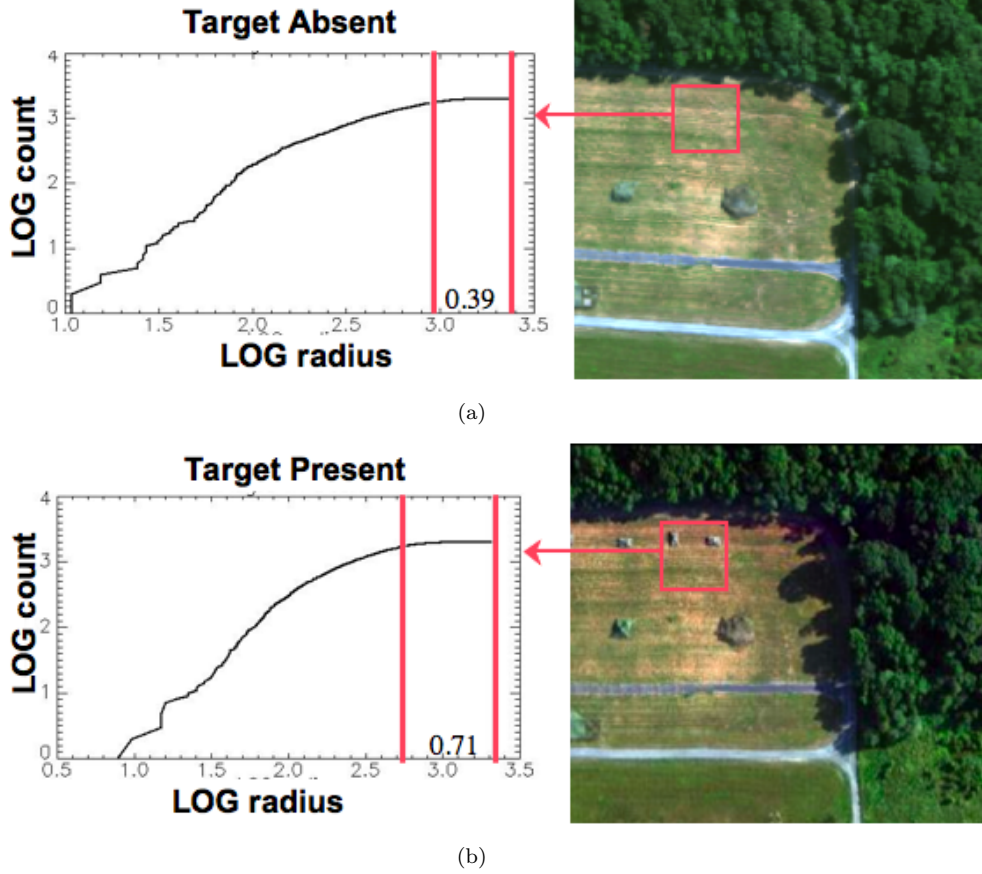


Figure 4.2: PDPs of two Hydice tiles illustrating the difference in PDTL between the target absent (a) and target present (b) cases.

For display purposes, the RGB bands from one of the tiled images are saved. The $\Delta PDDL$ metrics computer for the entire image are scaled between 0 and 1. For each tile, the scaled $\Delta PDDL$ is multiplied by the RGB values in the corresponding tile before the image is stitched together. The bright tiles in the resulting image correspond to tiles where a change has likely occurred (*i.e.*, $\Delta PDDL$ is large) between the two images while a change probably did not occur between the images if the tile is dark (*i.e.*, $\Delta PDDL$ is small). Before multiplication, the scaled $\Delta PDDL$ may be thresholded. If this is done, after multiplication all tiles below the threshold appear black in the detection map while all tiles above the threshold appear in full color and brightness.

The change detection method based on point density analysis was developed to assist an image analyst looking for change. The algorithm can be run with little user interaction and provide a detection map that is easily readable. This method is tile-based, not pixel-based and as a result, the detection map indicates where a change has likely occurred and where the analyst should focus his or her attention by altering the brightness of individual image tiles. At the same time, the method eliminates regions where a change did not occur by turning tiles off (or black), reducing the area to which an analyst must devote attention. The computationally expensive part of the process is calculating the distance from the centroid of the distribution to all other pixels. However, because the image is split into tiles before this is done, there are fewer computations. For purely change detection purposes, the dimension does not need to be estimated and the user can ignore or discard the region of the PDP that does not correspond to the tail. Because each tile is processed individually apart from final display, parallel processing may be used to further increase the efficiency and decrease the run time.

Typically, in order to observe small, single, and mixed pixel changes between images, pixel level differencing is necessary. However, as this method uses a new representation of the data in the point density space, these small changes are preserved and do appear as changes in the detection map. Also common in change detection is assigning one image to be the reference image and another the detection image as well as a prediction step, but this is not necessary in this method.

The first test of the algorithm's ability to distinguish areas of change from static regions between two images was done on imagery with synthetically generated changes. Figure 4.3 shows the resulting image where the tile brightness is scaled by the $\Delta PDDL$ value. In the areas of the image where no change was implanted the $\Delta PDDL$ is zero and the image appears black. The tiles where a change was implanted have larger $\Delta PDDL$ values. The more pixels that changed in a particular tile between the two images, the larger the $\Delta PDDL$ value is and the brighter the tile




Figure 4.3: Change detection result for HyMap image pair with implanted manmade material changes.

is displayed. All areas with known change did have $\Delta PDDL$ values above zero and were found by this change detection methodology. The darker tiles with non-zero $\Delta PDDL$'s all had only one or two pixel changes, while the brighter tiles had between 5 and 10 pixel changes.

Figure 4.4 shows the raw change detection result from the presented algorithm and the thresholded results at various thresholds on Hydice data. The raw result




Figure 4.4: Change detection results for high altitude HYDICE data at various threshold values.

shows the varying amount of change between the two images. Brighter tiles contain more significant change and have larger $\Delta PDDL$ values than darker tiles. The thresholded results show all of the tiles at full brightness with $\Delta PDDL$ values above the set threshold. As the threshold increases, the number of tiles displayed reduces as

fewer tiles have larger ΔPDTL 's. All of the tiles containing the known vehicle target changes do appear, along with some tiles that do not have known changes. Some of these are false alarms and there is no true change between the two images. Others are actually changing, mostly due to the shadows created by different illumination conditions; however, these changes are not interesting.




Figure 4.5: Metric value in a tile as a function of the subjective amount of change contained in the tile for Hydice forest data at full spectral resolution and a detection map where the tile brightness is scaled by the metric value.

Many of the results presented below are detection maps where the ΔPDTL is thresholded at a variety of values. In all cases, as the threshold increases, more tiles are eliminated from the detection map. Figure 4.5 shows the ΔPDTL of the approximately 1400 pixel tiles in the Hydice forest data as a function of the subjective amount of change that occurred between corresponding tiles in the two data sets. Also shown is a detection map where brightness of the tile corresponds to the ΔPDTL . The tiles labelled “strong change” contain multiple pixels of significant change including one or more targets appearing during one flight and not the other. “Some change” in a tile corresponds to tiles containing single pixel changes, natural change due to differences in illumination, or tiles with poor registration. This pattern is true for the hyperspectral case in Figure 4.5 as well as the multispectral cases.




Figure 4.6: Thresholded detection maps of HYDICE forest data at full spectral resolution (170 bands).

Figure 4.6 shows the results from forest radiance data at full spectral resolution. When a low threshold is set, approximately one third of the image is eliminated, while all of the tiles containing targets are still visible. As the threshold increases, more of the image is eliminated, including some known targets in the upper left hand side of the clearing. The forest tiles of this image are not eliminated until the threshold is extremely high. However, this is not necessarily a false alarm. These two images were taken at different times of day. One of the flightlines contains very strong shadows in and near the forest, while the other does not. The changes in illumination are true changes and appear as such in this methodology, though they are not interesting changes. The effects of data preprocessing methods for shadow suppression are under investigation.




Figure 4.7: Thresholded detection maps of Hydice forest data at WorldView2 spectral resolution (8 bands).

Similar effects are seen when using multispectral version of these images, but there are some key differences. The multispectral results, shown in Figures 4.7 (a) and (b) preserve most of the larger targets in the upper left hand side of the clearing at higher thresholds, but lose a few of smaller targets in the lower portion

of the image. The larger targets are lost at lower thresholds in the hyperspectral while the smaller targets are preserved. Shadows also have less of an effect in the multispectral cases, as more shadow forest tiles are eliminated at lower thresholds. The line illustrates the general downward trend of the Δ PDTL as a function of amount of change in the tile. Tiles with any amount of change in general have significantly higher Δ PDTL values compared to those with no noticeable change.

The forest radiance imagery contains three large targets near the roads that were not moved into the woods for the second image. There is also a white target in the center of the lower half of the image that was left out in the open. These targets frequently appear as anomalous in anomaly detection routines. However, as these particular targets were left unchanged between the two images, they do not appear as changes. This indicates that in a region where there may exist an established building or more clutter, something that is different from the background and constant will not cause a large Δ PDTL value in this methodology. This is one way in which the method is different from the simple subtraction of anomaly maps.




Figure 4.8: Thresholded detection maps of HyMap data at full spectral resolution (126 bands).

Figure 4.8 shows results of the change detection methodology applied to HyMap data with approximately 1600 pixels in each tile. This scene has more natural material clutter and contains regions of manmade clutter. Some areas change, including a field with targets and multiple moved cars, between the two images, but much of the content is static as the images were taken only a few hours apart. With a low threshold set, approximately one third of the image is preserved. Some of these areas are in the downtown Cooke City, MT area where many cars were moved around over the course of the day. The bright tiles in the wooded and grassy areas are likely false alarms, whether due to illumination or misregistration changes. When the threshold is raised, most of these false alarms are eliminated. What remains at a higher threshold are the areas where targets were removed before the second image was taken and areas in the city where there are known changes due to cars driving in and out of the scene.

Initial application of this change detection methodology to large area search is shown in Figure 4.9 on 4-band Quickbird images of the Indonesia coastline before and after the tsunami of 2004. The tsunami caused significant damage to the coastline and some inland areas. The raw change detection map displaying the PDTL for all tiles scaled between zero and one in Figure 4.9(c) shows that little has changed inland and the most significant changes are in the water and on the coast. The map in (d) shows the top 1% of changes with the largest PDTL scores in the image.

There are more than ten boats in the open water visible in the first image that are not present in the second image. These are found by this change detection methodology. An example is shown in Figure 4.10. The brightest tiles correspond to the two boats. In one case, the boat falls on the boundary between two tiles, so both tiles are given a high score.

Another example is shown in Figure 4.11 where the boat is present in the upper left hand corner of the first image and absent from the second image. This result also shows some of the significant damage to the coastline. The brightest tile in this region corresponds to the significant change caused by damage to what originally was an inlet from the ocean becoming more like a mouth of a bay with only a small peninsula of land left. Additional coastline damage is shown in Figure 4.12. Also apparent in the water region of the PDTL map is a long line of strong changes in the lower left hand corner of the image. This area, shown in Figure 4.13, contains a line of buoys and is possibly a large fishing net. After the tsunami, these features are gone.

Though most of the tsunami damage was done to the coastline, some regions inland have PDTL detections from this methodology. One area in particular is shown in Figure 4.14. Before the tsunami, this region had a small river or stream

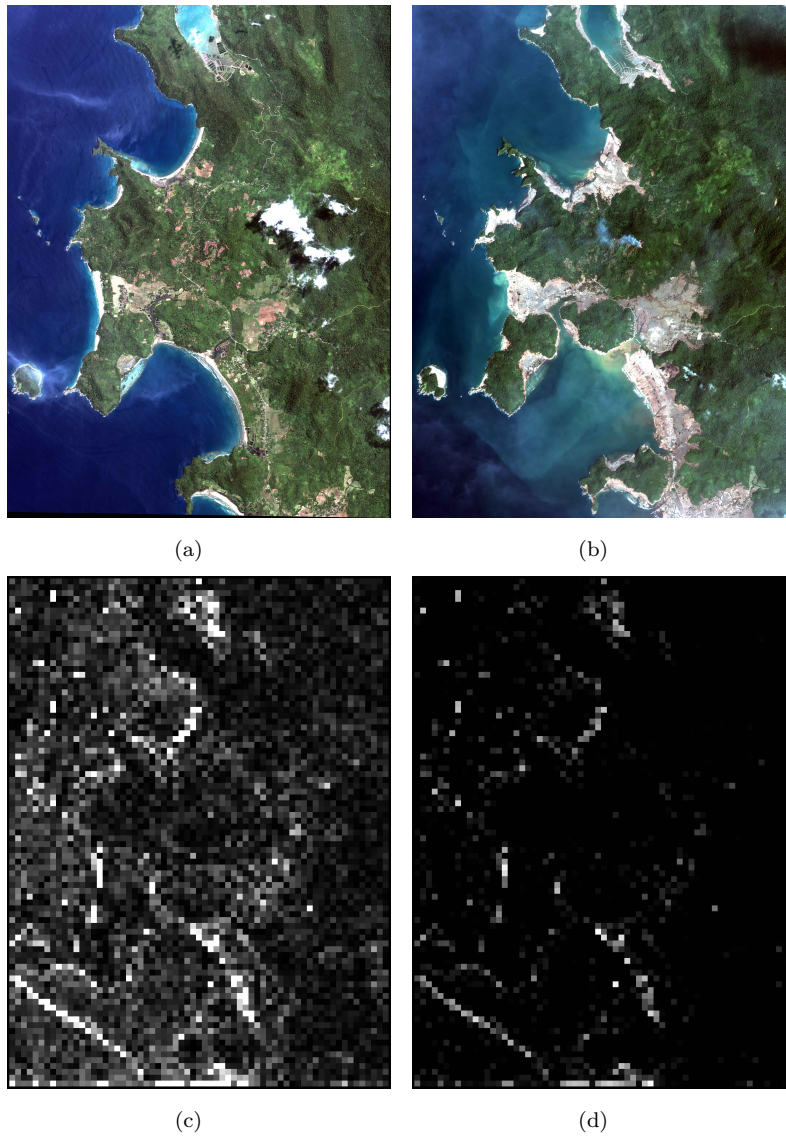


Figure 4.9: Indonesia coastline before (a) and after (b) the tsunami of 2004. Raw PDDL map (c) and scaled map showing only the top 1% of detected changes (d).

on the left connected to an ocean inlet. The tsunami flooded this stream with water and completely wiped out regions of houses and crops. The brightest tile in this region corresponds to a bright tower in one image completely missing from another image. Another inland region is shown in Figure 4.15(a). This area does not show significant flood damage from the tsunami, but there is a single tile with significant

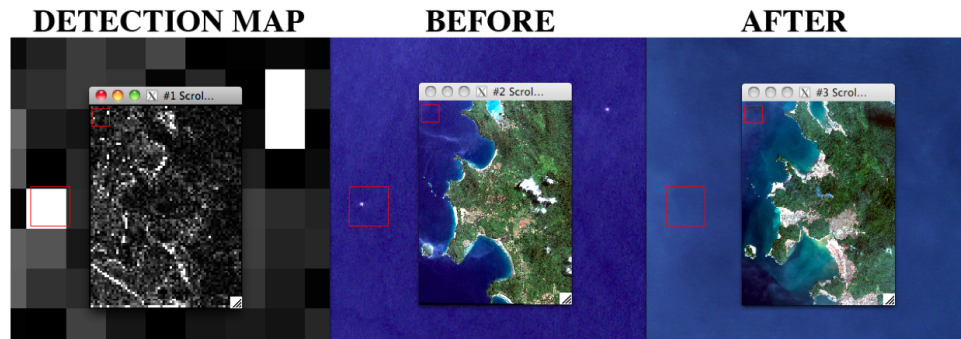


Figure 4.10: Change detection results for a region off the coast on Indonesia with two boats.




Figure 4.11: Change detection results for a coastal region of Indonesia showing a boat and coast damage.

change. Upon close inspection in (b), on the edge of a clearing in this region in the second image there are clusters of bright red and blue pixels where there was previously nothing. This is likely a camp that was set up far inland either before the tsunami hit Indonesia or set up afterwards as a place for displaced people to go.

Ideally, a universal threshold could be set on the PDTL for all data sets. Below this threshold, the tile will not contain changes and the tiles can be removed from further analysis. More realistically, the threshold will be set for a particular image or an adjustable scale could be implemented for interactive analysis. In the case of the HYDICE forest data, in order to detect all of the known targets, a very low threshold (0.05 or under) should be set. This eliminates approximately one third of the image. The optimal threshold of this image varies slightly between the hyper and multispectral cases. The COMPASS data however loses few targets when a



Figure 4.12: Change detection results for an area of significant coastal damage in Indonesia.




Figure 4.13: Change detection results for a region in the ocean off the coast of Indonesia containing a line of buoys suspended in the water in one image missing from the later image.

high threshold is set and have nearly identical detection maps across varying spectral resolution. A threshold around 0.3 can be set to eliminate eighty percent of the image from further analysis. Alternately, the detection map can be displayed similarly to target or anomaly maps, for example scaling the display so only the highest 1% of the image is visible. This change detection methodology was implemented into an ENVI GUI as part of RIT Toolkit for ENVI currently in development.



Figure 4.14: Change detection results on an inland region of Indonesia showing tsunami damage to a crop region.



(a)



(b)

Figure 4.15: Change detection results on an inland region of Indonesia showing a rescue camp (a) and a close up of the blue tents (b).

Chapter 5

SUMMARY

This dissertation consists of improving the greater understanding of the phenomenology of spectral data in the native hyperspace and developing a variety of features that can be applied to the full distribution of spectral data in order to locate “interesting” regions of manmade activity and regions that are not “interesting.” Spectral image analysis typically involves one or more of the following: target/anomaly detection, classification, and change detection. These typically involve either statistical or geometric restraints being place on the data. The algorithms developed for this research are predominantly data driven and avoid placing assumptions on the data, including the TAD and Gradient Flow clustering. The point density approach used here does involve some assumptions, but they are valid for local areas of the distribution. In addition, when the assumption is violated, more information regarding the distribution is manifested in the plot and can be used for additional analysis.

5.1 Future Work

5.1.1 Dimension Estimation and Subject Phenomenology

Section 2.3 introduced using an estimate of the inherent dimension of the spectral distribution of data as a way to indicate the presence of manmade activity within a scene. With more work, inherent dimension estimation may be useful as a way to quantify other information about the data. The phenomenology of spectral data in the native hyperspace is still not well understood. Figure 5.1 shows the estimated dimension from the PDP, box-counting algorithm, and PCA as a function of time of year for Hyperion data over the same mountainous region in Oaxaca Valley, Mexico. The graph plot shows that between November and December, the estimation




Figure 5.1: Estimated dimension for different months during the year.

dimension for all three algorithms increases significantly. This corresponds to when the rainy season starts in the Oaxaca Valley. After the season change occurs, the estimated inherent dimension begins to decrease. In this case, the increased dimension estimate may correspond to the season change, the health of the trees, or the saturation of the ground. The inherent dimension may be useful for determining how conditions of the ground relate to the distribution of the data.

5.1.2 Gram Matrix Volume for Improved Dimension Estimation

Section 2.4.2 introduces the Gram matrix, G , and its unique property,

$$|G| = V^2, \quad (5.1)$$

where V is the volume of the parallelepiped that encloses the data. This provides a simple and direct way for calculating the volume of a distribution. The main assumption behind the point density approach to spectral image analysis is

$$V \propto N_V \quad (5.2)$$

which is approximately true in a dense region of the spectral distribution, but is not true of the distribution as a whole. We use this assumption to simplify our method by counting points, not calculating volumes. Using the Gram matrix approach to estimate the volume of the distribution at increasing radii from the center may be a way to avoid the assumption in 5.2 and improve the results of dimension estimation from the point density approach.

5.1.3 Additional Analysis Using the PDP

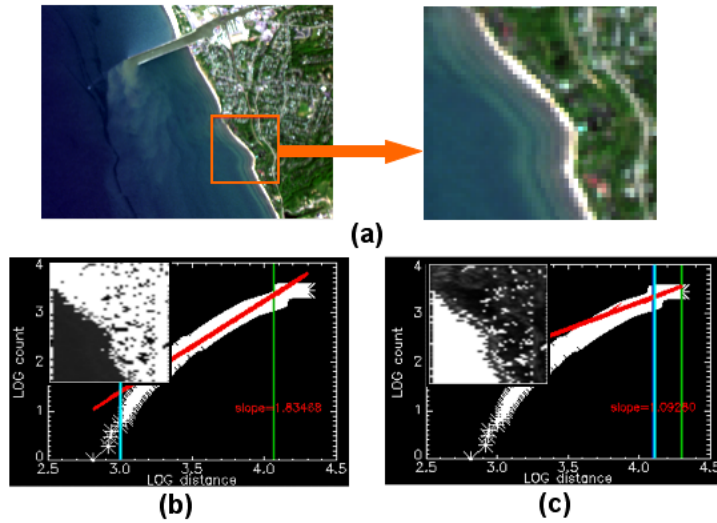


Figure 5.2: (a) RGB image of an AVIRIS hyperspectral flightline over Rochester, NY and the particular region of interest used in processing. (b) Dimension estimation plot of the scene and a greyscale map showing that the land pixels are those contained in the inclining portion of the graph. (c) Dimension estimation plot and greyscale map showing that the water pixels are contained in the plateau of the graph. Only those points between the vertical lines in the graph are used to estimate the dimension in (b) and (c).

The pixels mapped to the transition and tail regions of the PDP are spectrally different from those in the incline used to estimate the dimension. To demonstrate this, the PDP of an AVIRIS image near Rochester, NY was generated. Figure 5.2(b) shows the PDP for a region of land and water with a thumbnail image of the region where the points on the incline and transition region are turned white. These pixels correspond to nearly all the land pixels in the region. Figure 5.2(c) shows the same plot with the thumbnail highlighting all of the tail pixels, which are predominantly

the water pixels. This indicates that a PDP generated for an area with multiple landcover types may be used to cluster the data based on relative densities and locations in the spectral distribution.

Figure 2.30 shows a screen shot of the Point Density Analysis GUI created for this research. In this example, a tile containing mostly grass and exposed soil with a single vehicle parked is used to create the PDP. The pixels most spectrally different from the center of the distribution are turned red in a greyscale image in the GUI. These correspond to the anomalous pixels in the image. The PDP may be used for anomaly detection in addition to clustering.

5.1.4 Additional Features

Additional research is being conducted into locating anomalous objects from a thresholded anomaly map and using an additional processing step to locate the connected anomalous components. Currently a method is being developed to locate spectrally close and spatially connected components through matrix methods. Morphological processing or segmentation could also be used to locate anomalous objects from an anomaly map. This information can then be used to create another feature map corresponding to the number of anomalous objects. In this case, two tiles with an anomalous object would be given the same score regardless of the number of pixels that make up the object.

If a tile has a few classes, more information, like the size of the class, may be helpful in determining the overall interest in the tile. For example, the dimension estimate of a particular class may be used to indicate the amount of interest in the class. If the contained classes are spectrally near each other in the hyperspace, they are likely less interesting than two spectrally different classes. Vegetation is well modeled by a Gaussian distribution while manmade materials are not. A measure of how well the class fits to a multivariate normal distribution can be used to determine the contents of the class. The number of outliers in the class can be used for this purpose. It has been shown that elliptical T-distributions can better model hyperspectral data than Gaussian distributions and may also be used to generate measures such as the weight of class tails, to indicate the interest of the data.

5.1.5 Geometry of Hyperspectral Distributions

Methods using the topology or graph of the data have been developed into dimension estimation techniques, anomaly detectors, and land cover classifiers. However, the way that the data lie in the hyperspace can be interpreted in more than just

dimension. For example, several methods assume that the data lie in a convex hull. However, often there are large holes with no data within this hull. One idea is to try to estimate how much of the convex hull is empty. This may be estimated by measuring the local density around points in the hyperspace to determine the proportion of how many of the points are in a low density area and how many are in a high density area.

5.1.6 Spatial Features: Object Identification and Spatial Frequencies

Typically spatial information is ignored in hyperspectral image analysis. The spatial resolution of hyperspectral imagery is usually low enough that there are no strong edges or other useful spatial characteristics. As technology improves, the spatial resolution of hyperspectral information is increasing and spatial information may be useful to discriminate between interesting and uninteresting areas. For example, a tile filled with grass or trees likely does not have any straight edges, which indicates it is not interesting. Manmade regions will likely have straight roads, buildings, or other structures and are more interesting.

When edge detection is used in hyperspectral analysis, it is simply either applied to one representative band of the image, an average of bands, or each band separately. In the latter case, the resultant edge maps are combined into a single map. [7] There have been methods presented that take the entire hyperspectral cube as input and create a resulting edge map directly. Bakker and Schmidt[7] present a method based on 2D Laplacian convolution and extend it to the full spectral extent of the scene using either the spectral angle, intensity difference or Euclidean distance as the distance measure. Gong and Bi[19] describe a method of edge detection using the proportion images from unmixing methods, but this requires the assumption that each pixel is a mixture of pure pixels. Current research is attempting to expand a gradient-based method for RGB image segmentation for multi- and hyperspectral imagery.[45, 40]

The spatial texture of manmade regions also differs from the texture of natural materials. Haralick co-occurrence matrices are the mostly widely used method for measuring texture, however they are designed to be used on single band images. [55] Extending the co-occurrence matrix technique to hyperspectral imagery is computationally expensive. Other methods for determining texture features include Gabor filters, fractal dimensions, and wavelets.

Recently, methods involving 3D Fourier transforms of the entire hyperspectral cube have been investigated for including texture information in classification. Bau,

etal. (2008), [12] use Gabor filters in the 3D frequency space of a hyperspectral scene for classification. A Gabor filter is a sinusoid modulated by a Gaussian function in the radiance or spatial space. In this method, a Gabor filterbank (made of the Fourier transforms of the Gabor filters) is applied to the 3D spectral/spatial Fourier transform of a region; each filter will measure a single number corresponding to how much power is located at a particular frequency region. The filter responses and the mean spectrum of the region are put into vector form and the Mahalanobis distance is used as a classifier. With this process, a region is classified based on the spatial and spectral features in a small region.

The fractal nature of image texture has been investigated for use in discriminating natural and manmade materials in two dimensional data. [41, 70] One characteristic of fractals is that they are self-similar at different scales. A smooth or repeated patterned surface has essentially the same fractal dimension at different scales. Peli (1990) uses the fractal dimension as a function of scale, or fractal signature, to discriminate between natural and manmade surfaces. The fractal dimension of natural surfaces is constant or slow varying as a function of scale, while a region containing manmade materials change rapidly. Zwiggelaar (1995) comes to a similar conclusion, but he computes the fractal dimension in the spatial frequency domain, while Peli (1990) computes it based on the digital count values. Both methods are based on application to single-band images.

Texture features are commonly derived from imagery and used to improve classification. Manian, *etal.* (2007), derive statistical and wavelet features from a hyperspectral image to improve land cover classification. [34] The typical statistical measures for this purpose are the average, standard deviation, average deviation of the gradient, and residual energy. These are computed locally and may also be calculated to isolate spatial features in the horizontal or vertical directions. They also derived texture features from the subimages of a wavelet transform to measure the regularity and homogeneity of the texture. These features are derived for every pixel in every band of the image. The optimal spatial and spectral features are selected from the large set and used to classify the image.

Methods based on the 2D Fourier transform of each spectral channel or 3D Fourier transform of a spatial tile from a hyperspectral cube may be useful for incorporating spatial information as a feature for interest segmentation. Transforming the data into frequency space provides a cube that will describe the spatial frequencies at all wavelengths. A smooth region of no variability in the original scene (at any wavelength) will correspond to the DC component of the Fourier transform. A region with a significant amount of texture, like a tree canopy, will correspond to variability over a very broad range of spatial frequencies. Manmade objects are usu-

ally linear spatially; this will manifest in the 3D Fourier transform as straight lines in the spatially orthogonal direction. These features are shown in Figure 5.3, where the spread of the city Fourier transform is predominately in the vertical direction while the spread of the forest Fourier transform is in every direction. Using the 2D or 3D Fourier transform will provide a new space in which characterization of what is occurring spatially in the region.




Figure 5.3: Fourier transforms of city (left) and forest (right) hyperspectral data.

Appendix A

Additional PDP Feature Maps

The feature maps from the point density approach to spectral image analysis for the HSI and MSI data used for interest segmentation are shown below. The development of the point density approach to inherent dimension estimation can be found in Sections 2.3 and 3.1.1. The expansion of this method to measure additional features measured PDP is described in Section 2.4.1. The brightness of a tile corresponds to the estimated feature value. The feature is large for a bright tile and small for a dark. Sample feature maps and analysis is provided in Section 3.1.2. The feature maps below were used to create the interest segmentation maps in Chapter 3.

A.1 Dimension

A.1.1 Hyperspectral




Figure A.1: RGB image (a,c) and PDD feature maps (b,d) of Hydice Forest imagery with exposed and hidden targets.

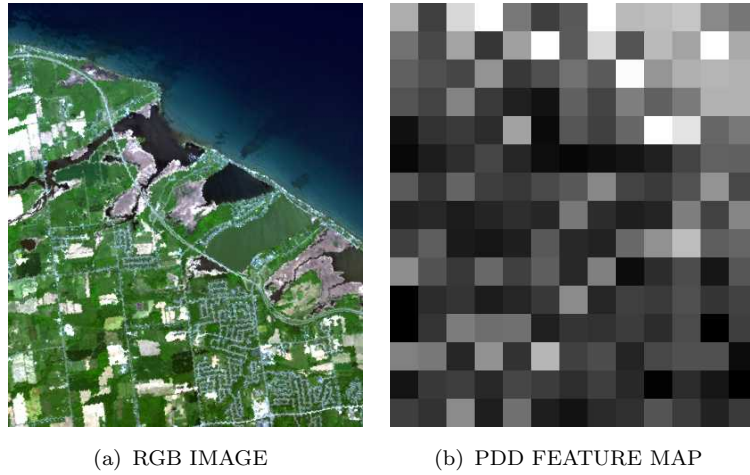


Figure A.2: RGB image of AVIRIS data of Rochester, NY near Lake Ontario (a) and corresponding PDD feature map (b).

A.1.2 Multispectral




Figure A.3: RGB image of AVIRIS data of Moffett Field, CA (a) and corresponding PDD feature map (b).






Figure A.4: RGB image of sample WV2 data of a forested region (a) and corresponding PDD feature map (b).



(a) RGB IMAGE




(b) PDD FEATURE MAP

Figure A.5: RGB image (a) and PDD map (b) for Quickbird imagery near the Esperanza forest fire containing a large agricultural and manmade region.



(a) RGB IMAGE




(b) PDD FEATURE MAP

Figure A.6: RGB image (a) and PDD map (b) for Quickbird imagery near the Esperanza forest fire containing a suburban region and multiple small scale regions of interest.



(a) RGB IMAGE



(b) PDD FEATURE MAP

Figure A.7: RGB image (a) and PDD map (b) for Quickbird imagery near the origin Esperanza forest fire containing many small and large scale regions of interest.

A.2 PDTL

A.2.1 Hyperspectral

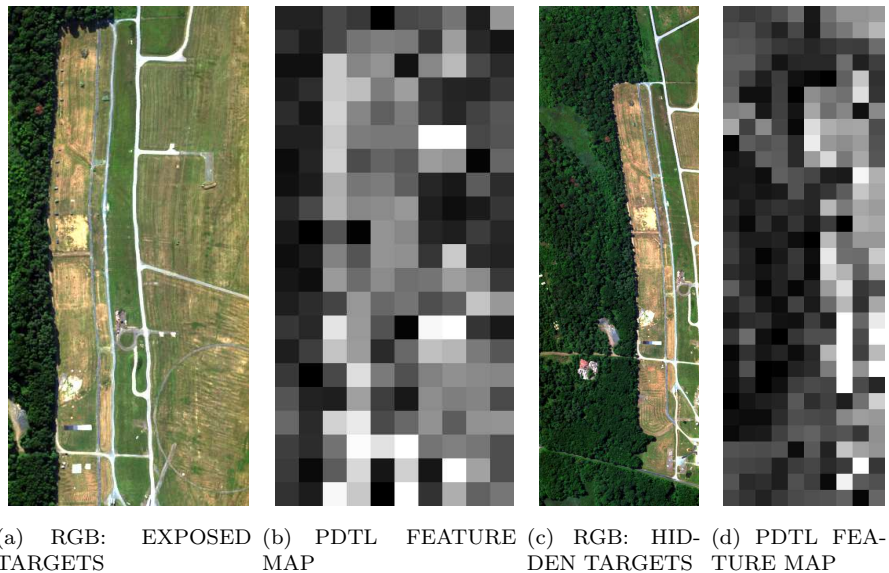


Figure A.8: RGB image (a,c) and PDTL feature maps (b,d) of Hydice Forest imagery with exposed and hidden targets.




Figure A.9: RGB image of AVIRIS data of Rochester, NY near Lake Ontario (a) and corresponding PDTL feature map (b).




Figure A.10: RGB image of AVIRIS data of Moffett Field, CA (a) and corresponding PDTL feature map (b).

A.2.2 Multispectral






Figure A.11: RGB image of sample WV2 data of a forested region (a) and corresponding PDTL feature map (b).




(a) RGB IMAGE




(b) PDTL FEATURE MAP

Figure A.12: RGB image (a) and PDTL map (b) for Quickbird imagery near the Esperanza forest fire containing a large agricultural and manmade region.



(a) RGB IMAGE




(b) PDTL FEATURE MAP

Figure A.13: RGB image (a) and PDTL map (b) for Quickbird imagery near the Esperanza forest fire containing a suburban region and multiple small scale regions of interest.



(a) RGB IMAGE



(b) PDTL FEATURE MAP

Figure A.14: RGB image (a) and PDTL map (b) for Quickbird imagery near the origin Esperanza forest fire containing many small and large scale regions of interest.

A.3 TEP

A.3.1 Hyperspectral

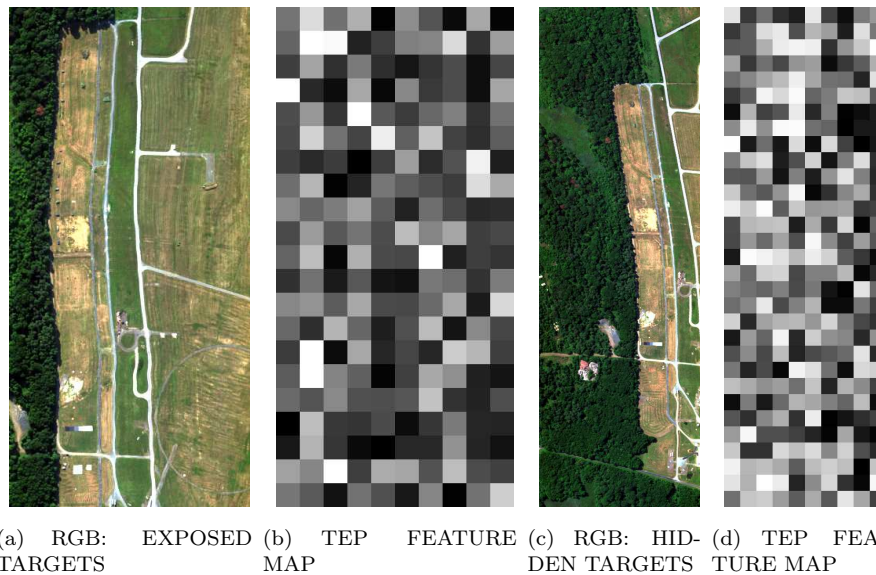


Figure A.15: RGB image (a,c) and TEP feature maps (b,d) of Hydice Forest imagery with exposed and hidden targets.




Figure A.16: RGB image of AVIRIS data of Rochester, NY near Lake Ontario (a) and corresponding TEP feature map (b).




Figure A.17: RGB image of AVIRIS data of Moffett Field, CA (a) and corresponding TEP feature map (b).

A.3.2 Multispectral






Figure A.18: RGB image of sample WV2 data of a forested region (a) and corresponding TEP feature map (b).



(a) RGB IMAGE




(b) TEP FEATURE MAP

Figure A.19: RGB image (a) and TEP map (b) for Quickbird imagery near the Esperanza forest fire containing a large agricultural and manmade region.



(a) RGB IMAGE




(b) TEP FEATURE MAP

Figure A.20: RGB image (a) and TEP map (b) for Quickbird imagery near the Esperanza forest fire containing a suburban region and multiple small scale regions of interest.



(a) RGB IMAGE



(b) TEP FEATURE MAP

Figure A.21: RGB image (a) and TEP map (b) for Quickbird imagery near the origin Esperanza forest fire containing many small and large scale regions of interest.

Appendix B

Additional Gram Matrix Feature Maps

The feature maps from the Gram matrix complexity approach to spectral image analysis for the HSI and MSI data used for interest segmentation are shown below. The development of the Gram matrix approach to estimating spectral image complexity is described in Section 2.4.2. The brightness of a tile corresponds to the estimated feature value. Example feature maps and analysis is provided in Section 3.2. The feature, in this case number of endmembers, volume, and area, is large for a bright tile and small for a dark. These feature maps were used to create the interest segmentation maps in Chapter 3.

B.1 Number of Endmembers

The NEM feature maps for the AVIRIS Moffett Field, Ca data and one Quickbird image of Esperanza Forest Fire are not displayed because the same number of endmembers were found for every tile the entire image.

B.1.1 Hyperspectral

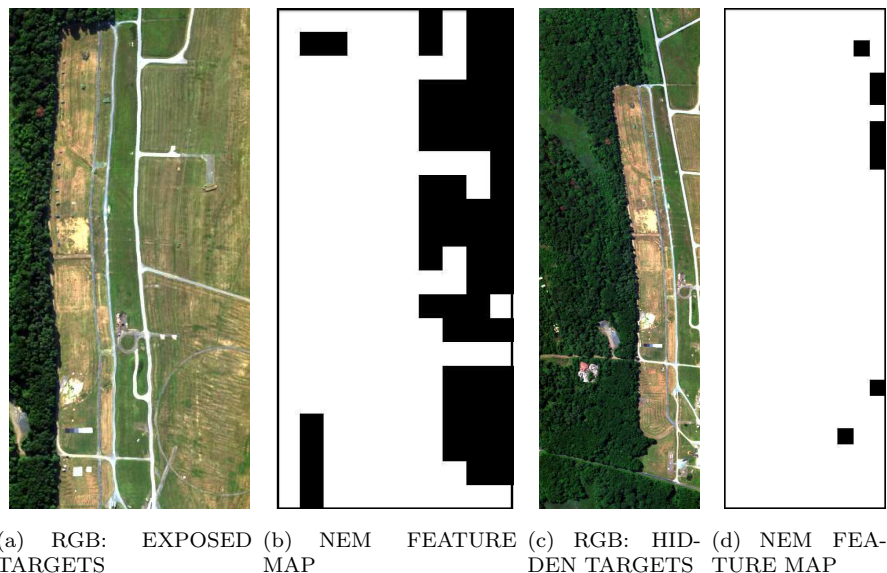


Figure B.1: RGB image (a,c) and NEM feature maps (b,d) of Hydice Forest imagery with exposed and hidden targets.

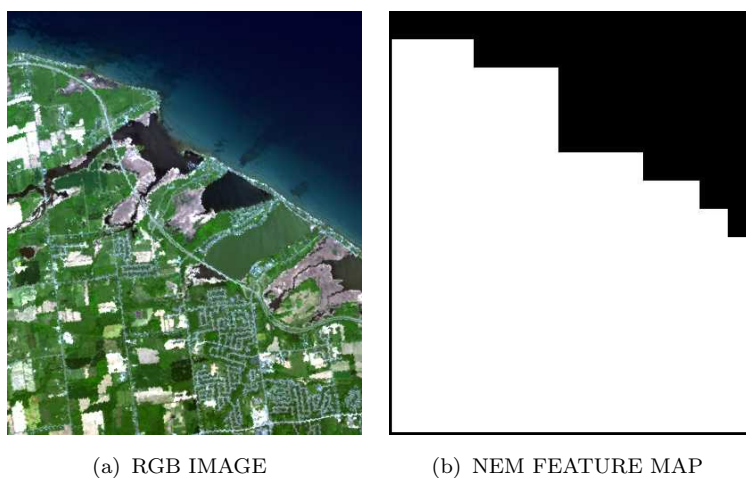


Figure B.2: RGB image of AVIRIS data of Rochester, NY near Lake Ontario (a) and corresponding NEM feature map (b).

B.1.2 Multispectral






Figure B.3: RGB image of sample WV2 data of a forested region (a) and corresponding NEM feature map (b).



(a) RGB IMAGE




(b) NEM FEATURE MAP

Figure B.4: RGB image (a) and NEM map (b) for Quickbird imagery near the Esperanza forest fire containing a suburban region and multiple small scale regions of interest.



(a) RGB IMAGE




(b) NEM FEATURE MAP

Figure B.5: RGB image (a) and NEM map (b) for Quickbird imagery near the origin Esperanza forest fire containing many small and large scale regions of interest.

B.2 Peak Volume

B.2.1 Hyperspectral



(a) RGB: EXPOSED TARGETS (b) PV FEATURE MAP (c) RGB: HIDDEN TARGETS (d) PV FEATURE MAP

Figure B.6: RGB image (a,c) and PV feature maps (b,d) of Hydice Forest imagery with exposed and hidden targets.




Figure B.7: RGB image of AVIRIS data of Rochester, NY near Lake Ontario (a) and corresponding PV feature map (b).




Figure B.8: RGB image of AVIRIS data of Moffett Field, CA (a) and corresponding PV feature map (b).

B.2.2 Multispectral






Figure B.9: RGB image of sample WV2 data of a forested region (a) and corresponding PV feature map (b).



(a) RGB IMAGE




(b) PV FEATURE MAP

Figure B.10: RGB image (a) and PV map (b) for Quickbird imagery near the Esperanza forest fire containing a large agricultural and manmade region.



(a) RGB IMAGE




(b) PV FEATURE MAP

Figure B.11: RGB image (a) and PV map (b) for Quickbird imagery near the Esperanza forest fire containing a suburban region and multiple small scale regions of interest.



(a) RGB IMAGE



(b) PV FEATURE MAP

Figure B.12: RGB image (a) and PV map (b) for Quickbird imagery near the origin Esperanza forest fire containing many small and large scale regions of interest.

B.3 Area Under Volume Curve

B.3.1 Hyperspectral

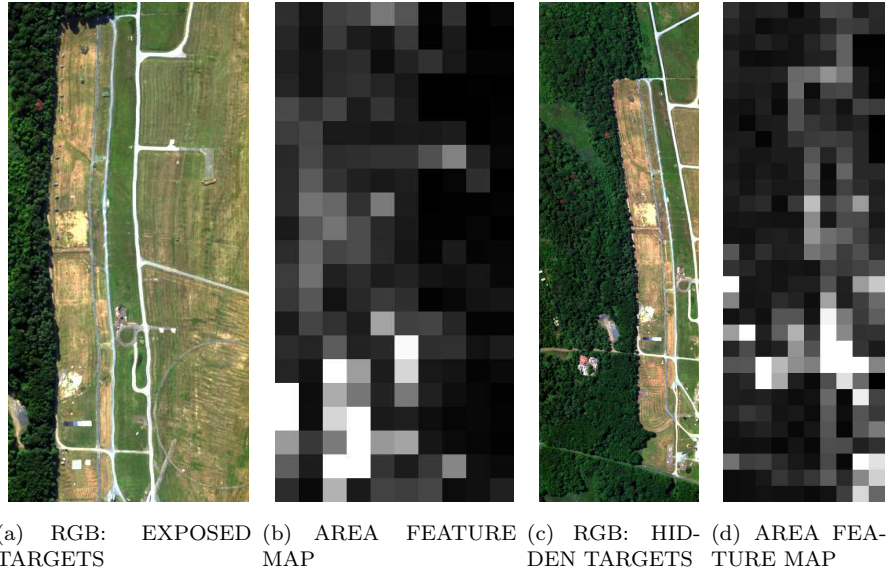


Figure B.13: RGB image (a,c) and AREA feature maps (b,d) of Hydice Forest imagery with exposed and hidden targets.

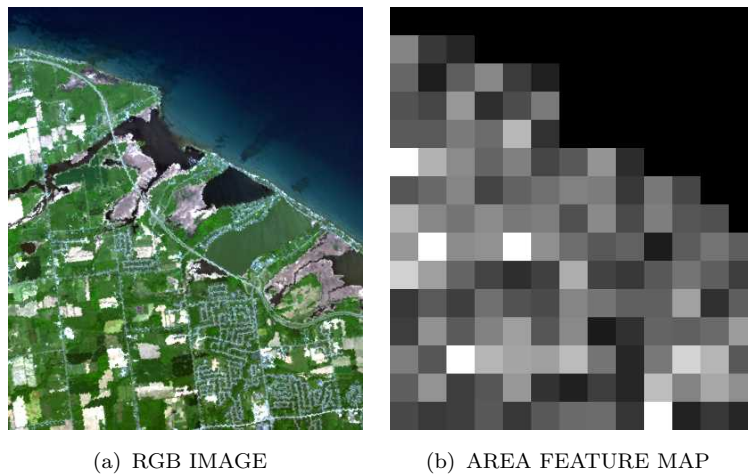


Figure B.14: RGB image of AVIRIS data of Rochester, NY near Lake Ontario (a) and corresponding AREA feature map (b).




Figure B.15: RGB image of AVIRIS data of Moffett Field, CA (a) and corresponding AREA feature map (b).

B.3.2 Multispectral






Figure B.16: RGB image of sample WV2 data of a forested region (a) and corresponding AREA feature map (b).




(a) RGB IMAGE




(b) AREA FEATURE MAP

Figure B.17: RGB image (a) and AREA map (b) for Quickbird imagery near the Esperanza forest fire containing a large agricultural and manmade region.



(a) RGB IMAGE




(b) AREA FEATURE MAP

Figure B.18: RGB image (a) and AREA map (b) for Quickbird imagery near the Esperanza forest fire containing a suburban region and multiple small scale regions of interest.



(a) RGB IMAGE



(b) AREA FEATURE MAP

Figure B.19: RGB image (a) and AREA map (b) for Quickbird imagery near the origin Esperanza forest fire containing many small and large scale regions of interest.

Appendix C

Additional Anomaly Feature Maps

Additional TAD results and the corresponding feature maps are shown below. The TAD algorithm is described in Section 2.5. The brightness of a tile corresponds to the number of anomalous pixels found by TAD within a tile. Example feature maps and analysis is provided in Section 3.3. These feature maps were used to create the interest segmentation maps in Chapter 3.

C.1 Hyperspectral

In order to compare the performance of TAD against a standard anomaly detector, TAD and the version of the RX algorithm built into ENVI were applied to the urban portion of the Cooke City, MT scene, the forest radiance scene with targets in the trees, and the forest radiance scene with targets out in the open. Figure C.1 shows the anomaly maps for the urban portion of the Cooke City, MT scene with 2% of the scene being anomalies. The buildings in the city appear anomalous as the scene contains predominately natural materials like grass and trees. Everything that RX rates as an anomaly is also found as an anomaly with the TAD algorithm. With TAD, there are all much stronger or more anomalous anomalies. TAD even finds spots on the left side of the image as anomalous that RX cannot even find.

The results of TAD and RX on the forest radiance scenes are shown in Figures C.2 and C.3. The striping in the RX results is a consequence of noise in the image. TAD does not show any of the noise as anomalous. TAD picks up every target set up in the scene that RX finds, as well as many that RX does not show as anomalous. RX also shows the roads through the scene and the shadow areas in the trees to be anomalous, while TAD suppresses these structures.




Figure C.1: Anomalies detected using TAD (left), RX (center), on the Cooke City, MT scene (right).




Figure C.2: Anomalies detected on the forest radiance scene with targets out in the open (top) using TAD (center) and RX (bottom).




Figure C.3: Anomalies detected on the forest radiance scene with targets in shadow (left) using RX (center) and TAD (right).




Figure C.4: RGB image of HYDICE data with exposed targets (a), TAD result (b), and corresponding NANOM feature map (c).




Figure C.5: RGB image of HYDICE data with hidden targets (a), TAD result(b), and corresponding NANOM feature map (c).

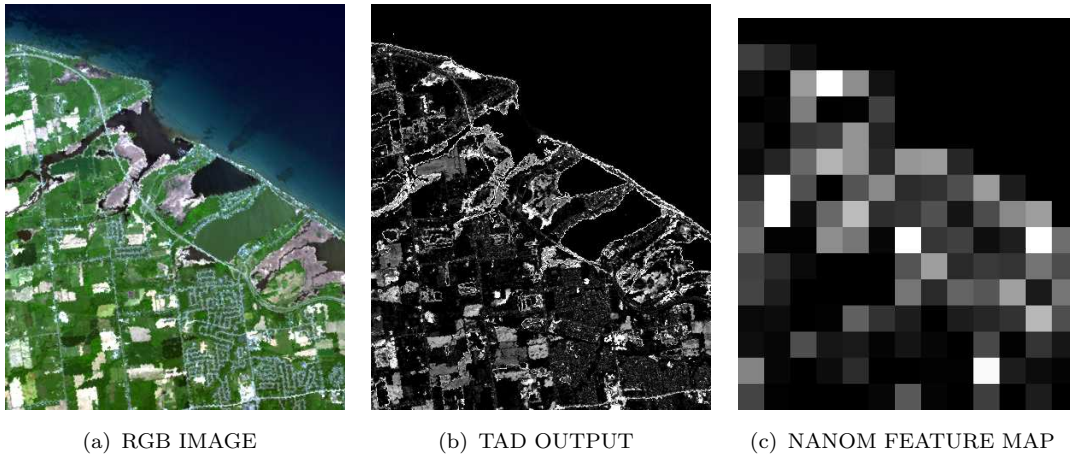


Figure C.6: RGB image of AVIRIS data of Rochester, NY near Lake Ontario (a), TAD result (b), and corresponding NANOM feature map (c).




Figure C.7: RGB image of AVIRIS data of Moffett Field, CA (a), TAD result (b), and corresponding NANOM feature map (c).

C.2 Multispectral






Figure C.8: RGB image of sample WV2 data of a forested region (a), TAD result (b) and corresponding PDTL feature map (c).



(a) RGB IMAGE



(b) TAD OUTPUT




(c) NANOM FEATURE MAP


Figure C.9: RGB image (a), TAD result (b), and NANOM map (c) for Quickbird imagery near the Esperanza forest fire containing a large agricultural and manmade region.



(a) RGB IMAGE



(b) TAD OUTPUT




(c) NANOM FEATURE MAP


Figure C.10: RGB image (a), TAD result (b), and NANOM map (b) for Quickbird imagery near the Esperanza forest fire containing a suburban region and multiple small scale regions of interest.



(a) RGB IMAGE



(b) TAD OUTPUT



(c) NANOM FEATURE MAP

Figure C.11: RGB image (a), TAD result (b), and NANOM map (b) for Quickbird imagery near the origin Esperanza forest fire containing many small and large scale regions of interest.

Appendix D

Additional Clustering Feature Maps

Additional Gradient Flow and ISODATA spectral clustering results and the corresponding feature maps are shown below. The Gradient Flow algorithm is described in Section 2.6. The brightness of a tile corresponds to the number of unique spectral clusters found by Gradient Flow within a tile. Example feature maps and analysis is provided in Section 3.4. In the case of the Quickbird imagery near the Esperanza forest fire, ISODATA was used as the unsupervised classification method to create these feature maps due to the size of the image and computing resources necessary for Gradient Flow to complete. These feature maps were used to create the interest segmentation maps in Chapter 3.

D.1 Hyperspectral

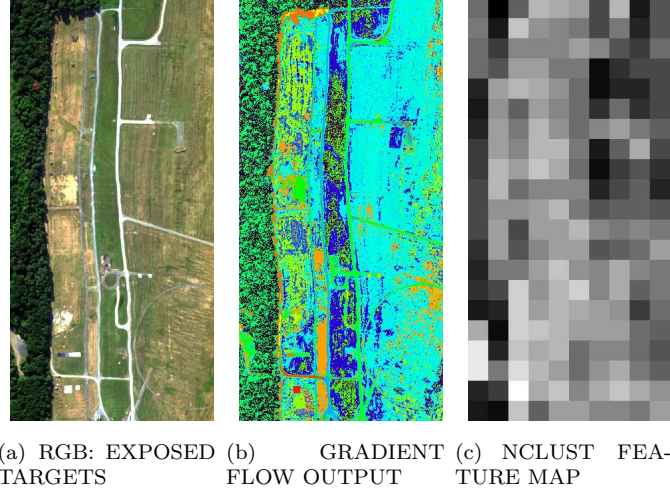


Figure D.1: RGB image of HYDICE data with exposed targets (a), Gradient Flow output (b), and corresponding NCLUST feature map (c).

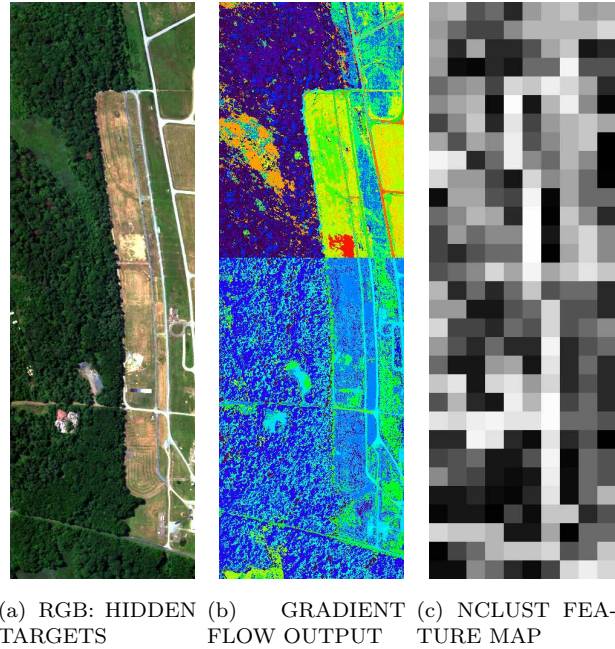


Figure D.2: RGB image of HYDICE data with hidden targets (a), Gradient Flow output (b), and corresponding NCLUST feature map (c).

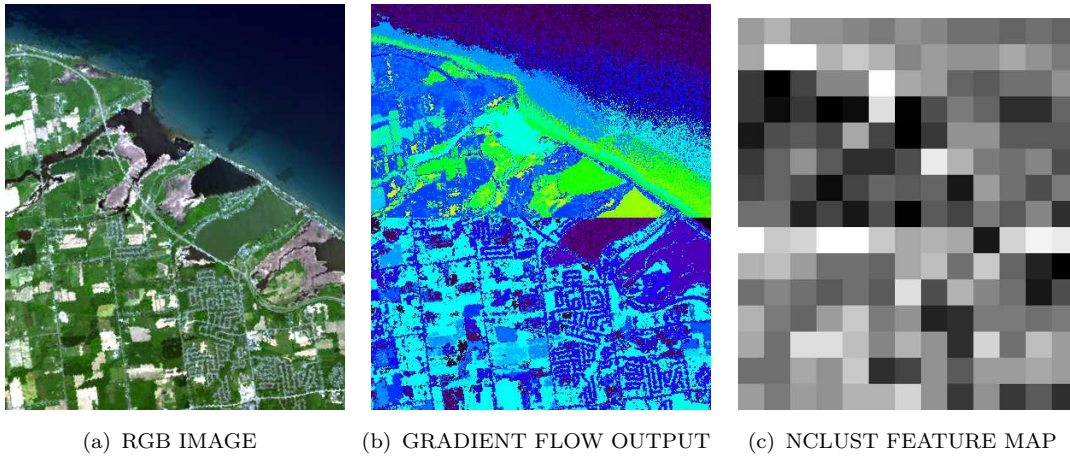


Figure D.3: RGB image of AVIRIS data of Rochester, NY near Lake Ontario (a), Gradient Flow output (b), and corresponding NCLUST feature map (c).

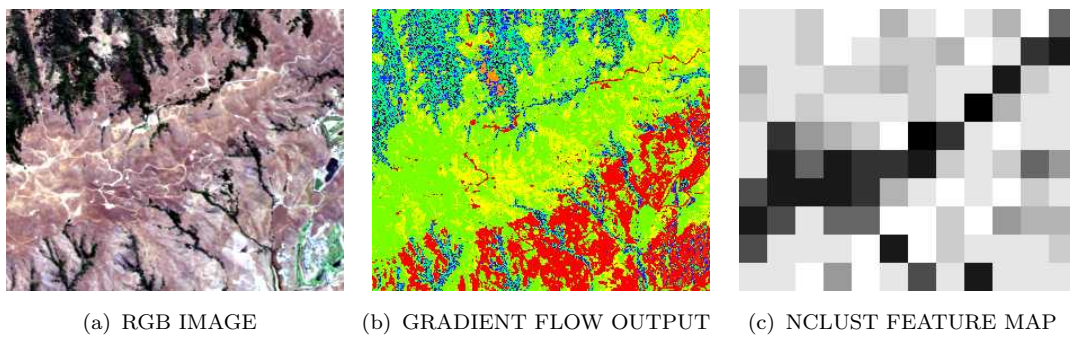


Figure D.4: RGB image of AVIRIS data of Moffett Field, CA (a), Gradient Flow output (b), and corresponding NCLUST feature map (c).

D.2 Multispectral






Figure D.5: RGB image of sample WV2 data of a forested region (a), Gradient Flow result (b) and corresponding NCLUST feature map (c).



(a) RGB IMAGE



(b) GRADIENT FLOW OUTPUT




(c) NCLUST FEATURE MAP


Figure D.6: RGB image (a), Gradient Flow output (b), and NCLUST map (c) for Quickbird imagery near the Esperanza forest fire containing a large agricultural and manmade region.



(a) RGB IMAGE



(b) ISODATA OUTPUT




(c) NCLUST FEATURE MAP


Figure D.7: RGB image (a), ISODATA result (b), and NCLUST map (b) for Quickbird imagery near the Esperanza forest fire containing a suburban region and multiple small scale regions of interest.



(a) RGB IMAGE



(b) ISODATA OUTPUT



(c) NCLUST FEATURE MAP

Figure D.8: RGB image (a), ISODATA result (b), and NCLUST map (b) for Quickbird imagery near the origin Esperanza forest fire containing many small and large scale regions of interest.

Acknowledgements

The path from being an amateur photographer in high school to earning a PhD was long . I wouldn't have made it all this way without help from a lot of people. Thanks to Dr. A for first putting a camera in my hands as a high school sophomore. I went to college to pursue Fine Art Photography, but discovered a longing for science and math after a presentation by Andrew Davidhazy. Terry Kessler made physics interesting—something I'd never thought possible—and his 5-hour night labs were a lot of fun. Nanette Salvaggio told me that it wasn't too late to apply to grad school and that I didn't need Modern Physics, and she was right. Everyone in DIRS helped me with something during the past few years, but Scott, Rolo, Niek, and Carl suffered through C++, UNIX, and IDL questions more than anyone. Likin Simone-Romero took me through two years of math classes and opened himself up to many discussions about math I am only beginning understand.

My whole committee played an important part in my graduate school career, above and beyond merely evaluating my worthiness to receive a degree. Emmett Ientilucci always had an objective eye and constructive questions—the first to say that maybe I proposed to do too much in my proposal. David Ross challenged me to explain what I was doing in a practical way. I owe a huge debt of gratitude to Carl Salvaggio, who stepped in at the last minute to serve on my committee and still had interesting questions and perspectives after just two weeks on board. Dave Messinger and Bill Basener kept me going on a project that often seemed like it would never work. Their excitement constantly re-energized me, getting me back to the drawing board and back on track. Luckily for all of us, things all seemed to work out pretty well in the end. Dave, thanks for always keeping your door open. It was always a surprise to hear what kind of music was playing in your office on any given day. Thanks also for convincing me that there isn't only one job out there for me.

I started off grad school by adopting a dog from ABC Basset Hound Rescue. Bogie gets me up every morning and makes sure that I get into bed at night; I couldn't imagine a better way to start and end each day. But I adopted more than a dog from ABC. I made friends by volunteering through the rescue and was proud to be the Volunteer Coordinator and later on the Board of Directors. There's always something more we can do for the hounds; focusing on them was a great way to distract myself from the everyday duties of schoolwork. Thanks to everyone at ABC, especially Gary and Shelly.

My friends and family entertained and supported me when I wasn't sitting in front of a computer. My phototech girls—Natalie, Danielle, Katelyn, and Julie—are

great friends. Natalie always has the ability to cheer me up. Danielle, Mike Z., and Jacqueline kept our office lively and fun. Kate and Natalie came over and baked with me on the night before my defense. My L-girls back home met up for many a long night at Dunkin Donuts. Laura edited my entire dissertation in exchange for crepes and coffee; you can't ask for a better friend than that. My best friend Dave has talked me through nearly every day in some way or another for 9 years. Neither of us have excuses anymore, so we're going to Coachella. Thank you Ben for walking Bogie in the mornings so that I could sleep in. We may not have celebrated our anniversary, but at least you got lots of scones!

Most of all I'd like to thank my mom. I've never gotten anywhere in life without her help and her trust. I certainly wouldn't have made it to RIT without her. Back in high school I was lazy, and she told me so. I thought I knew exactly what I wanted, and that certainly did NOT include any more math or science. I was determined to be an artist; I didn't want to apply to college outside of New York City, but she made me visit and apply to RIT. Clearly she knew something I didn't, because it turned out to be the best possible place for me. I can't imagine having spent the past 8 years anywhere else, and I can't imagine having done a single day of it without knowing that I could rely on her for guidance and support. Thank you for everything.

Bibliography

- [1] <http://edcsns17.cr.usgs.gov/eo1/hyperion.php>.
- [2] <http://www.intspec.com/products/hymap/overview/>.
- [3] <http://www.digitalglobe.com/>, 2009.
- [4] W. S. Aldrich, M. E. Kappus, Ronald G. Resmini, and P. A. Mitchell. Hydice postflight data processing. volume 2758 of *Proc. SPIE*, pages 354–363, 1996.
- [5] Charles M. Bachmann, Thomas L. Ainsworth, and Robert A. Fusina. Exploiting manifold geometry in hyperspectral imagery. *IEEE Trans. Geosci. Rem. Sens.*, 43(3):441–454, March 2005.
- [6] G. L. Baker and J. P. Gollub. *Chaotic Dynamics: An Introduction*, chapter 5: The characterization of chaotic attractors. Cambridge University Press, New York, first edition, 1990.
- [7] W.H. Bakker and K.S. Schmidt. Hyperspectral edge filtering for measuring homogeneity of surface cover types. *ISPRS Journal of Photogrammetry & Remote Sensing*, 56:246–256, 2002.
- [8] Amit Banerjee, Philippe Burlina, and Chris Diehl. A support vector method for anomaly detection in hyperspectral imagery. In *Transactions on Geoscience and Remote Sensing*, volume 44, pages 2282–2291. IEEE, August 2006.
- [9] R. W. Basedow, D. C. Carmer, and M. E. Anderson. Hydice system: implementation and performance. volume 2480 of *Proc. SPIE*, pages 258–267, 1995.
- [10] Bill Basener, Alexey Castrodad, David Messinger, Jennifer Mahle, and Paul Prue. A dynamical systems algorithm for clustering in hyperspectral imagery. In Sylvia S. Shen and Paul E. Lewis, editors, *Algorithms and Technologies*

- for Multispectral, Hyperspectral, and Ultraspectral Imagery XII*, volume 6966 of *Proc. SPIE*, 2008.
- [11] Bill Basener, David Messinger, and Emmett Ientilucci. Anomaly detection using topology. In Sylvia S. Shen and Paul E. Lewis, editors, *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XII*, volume 6565 of *Proc. SPIE*, 2007.
 - [12] Tien C. Bau, Subhadip Sarkar, and Glenn Healey. Using three-dimensional spectral/spatial gabor filters for hyperspectral region classification. In Sylvia S. Shen and Paul E. Lewis, editors, *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XII*, volume 6966 of *Proc. SPIE*, 2008.
 - [13] Charlene E. Caefer, Marcus S. Stefanou, Eric. D. Nielsen, Anthony P. Rizzuto, and Stanley R. Rotman. Analysis of false alarm distributions in the development and evaluation of hyperspectral point target detection algorithms. *Optical Engineering*, 46(7), July 2007.
 - [14] Francesco Camastra and Alessandro Vinciarelli. Estimating the intrinsic dimension of data with a fractal-based method. *IEEE Trans. Pattern. Anal. Mach. Intell.*, 24(10):1404–1407, October 2002.
 - [15] Mark J. Carlotto. A synergistic exploitation concept for wide area search. In Ivan Kadar, editor, *Signal Processing, Sensor Fusion, and Target Recognition XV*, volume 6235 of *Proc. SPIE*, 2006.
 - [16] M.J. Carlotto. Detection and analysis of change in remotely sensed imagery with application to wide area surveillance. *Image Processing, IEEE Transactions on*, 6(1):189–202, Jan 1997.
 - [17] Chein-I Chang. Exploration of virtual dimensionality in hyperspectral image analysis. In Sylvia S. Shen and Paul E. Lewis, editors, *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XII*, volume 6233 of *Proc. SPIE*, May 2006.
 - [18] Jose A. Costa and Alfred O. Hero. Geodesic entropic graphs for dimension and entropy estimation in manifold learning. *IEEE Trans. Signal. Process.*, 52(8):2210–2221, August 2004.
 - [19] Yanjun Gong and Dongmei Bi. The edge detection of hyperspectral image based on its proportion image. In Xun Hou, Wei Zhao, and Baoli Yao, editors,

- 27th International Congress on High-Speed Photography and Photonics*, volume 62792S. SPIE, 2007.
- [20] Peter Grassberger and Itamar Procaccia. Characterization of strange attractors. *Phys. Rev. Lett.*, 50(5):346–349, January 1983.
 - [21] Peter Grassberger and Itamar Procaccia. Measuring the strangeness of strange attractors. *Phys. Nonlinear Phenom.*, 9:189–208, 10 1983.
 - [22] David C. Grimm. Comparison of hyperspectral imagery target detection algorithm chains. Master’s thesis, Rochester Institute of Technology, 2005.
 - [23] J. C. Harsanyi and C. Chang. Hyperspectral image classification and dimension reduction: an orthogonal subspace project approach. In *IEEE Trans. Geosci. Rem. Sens.*, volume 32, pages 779–785, 1994.
 - [24] <http://aviris.jpl.nasa.gov>.
 - [25] <http://www.itervis.com/ProductServices/ENVI/Tutorials.aspx>.
 - [26] D. R. Hundley and M. J. Kirby. Estimation of topological dimension. In *Siam International Conference on Data Mining*. Society for Industrial and Applied Mathematics, 2003.
 - [27] J. Irvine. National image interpretability rating scales (niirs): Overview and methodology. volume 3128 of *Proc. SPIE*, pages 93–103, 1997.
 - [28] Michael Kirby. *Geometric Data Analysis: An Empirical Approach to Dimensionality Reduction and the Study of Patterns*, pages 8–11. John Wiley & Sons, INC, 2001.
 - [29] A. Kostrikin and Y. Manin. *Linear algebra and geometry*. Gordon and Breach Science Publishers, 1997.
 - [30] Heesung Kwon and Nasser M. Nasrabadi. Kernel rx-algorithm: A nonlinear anomaly detector for hyperspectral imagery. In *Transactions on Geoscience and Remote Sensing*, volume 43, pages 338–397. IEEE, 2005.
 - [31] Justin D. Kwong, David Messinger, and W. D. Middleton. Hyperspectral clustering and unmixing for studying the ecology of state formation and complex societies. volume 7457 of *Proc. SPIE*, 2009.

- [32] David Landgrebe. Hyperspectral image analysis. *IEEE Signal Process. Mag.*, 19(1):17–28, January 2002.
- [33] Ryan H. Lewis and Anthony A. Harkin. A network-theoretic approach to hyperspectral image classification. *JARS*, in preparation, 2010.
- [34] Vidya Manian and Luis O. Jimenez. Land cover and benthic habitat classification using texture features from hyperspectral and multispectral images. *Journal of Electronic Imaging*, 16(2), 2007.
- [35] David B. Marden and Dimitris Manolakis. Modeling hyperspectral imaging data using elliptical contoured distributions. In Sylvia S. Shen and Paul E. Lewis, editors, *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XII*, volume 5093 of *Proc. SPIE*, pages 253–262, 2003.
- [36] Joseph Meola and Michael T. Eismann. Methods for multi-temporal change detection. volume 7334, page 73340L. SPIE, 2009.
- [37] David Messinger. Dynamic analysis of spectral imagery for improved exploitation. NGA University Research Initiative, 2007.
- [38] David Messinger, Amanda Ziemann, Bill Basener, and Ariel Schlamm. A metric of spectral image complexity with application to large area search. *submitted to JARS*, 2010.
- [39] P. A. Mitchell. Hyperspectral digital imagery collection experiment: Hydice. volume 2587 of *Proc. SPIE*, pages 70–95, 1995.
- [40] A. E. Ononye, A. Vodacek, and E. Saber. Automated extraction of fire line parameters from multispectral imagery. *Journal of of Remote Sensing of Environment*, 108:179–188, 2007.
- [41] T. Peli. Multiscale fractal theory and object characterization. *J. Opt. Soc. Am.*, 7(6):1101–1112, June 1990.
- [42] Irving S. Reed and Xiaoli Yo. Comparative adaptive multiple-band cfar detection of a optical pattern with unknown spectral distribution. In *Transactions on Acoustics, Speech, and Signal Processing*, volume 38, pages 1760–1770. IEEE, October 1990.

- [43] Ronald G. Resmini. A tool for the nonparametric characterization of the geometry of spectra in hyperspace. volume 7334 of *Proc. SPIE*, 2009.
- [44] John A. Richards and Xiuping Jia. *Remote Sensing Digital Image Analysis*, chapter 6: Multispectral transformations of image data. Springer, 3rd edition, 1999.
- [45] E. Saber, A.M. Tekalp, and G. Bozdagi. Fusion of color and edge information for improved segmentation and edge linking. In *Image and Vision Computing*, volume 15, pages 769–780, 1997.
- [46] L. L. Scharf and L. R. McWhorter. Adaptive matched subspace detectors and adaptive coherence. In *20th Proceedings Asilomar Conference on Signals and Systems*, pages 114–117, 1996.
- [47] A. Schaum and A. Stocker. Advanced algorithms for autonomous hyperspectral change detection. In *Applied Imagery Pattern Recognition Workshop, 2004. Proceedings. 33rd*, pages 33–38. IEEE, Oct 2004.
- [48] Alan P. Schaum. Spectral subspace matched filtering. volume 4381 of *Proc. SPIE*, pages 1–17, 2001.
- [49] Alan P. Schaum and Alan Stocker. Hyperspectral change detection and supervised matched filtering based on covariance equalization. volume 5425, pages 77–90. SPIE, 2004.
- [50] Ariel Schlamm, David Messinger, and Bill Basener. Geometric estimation of the inherent dimensionality of a single material cluster in multi- and hyperspectral imagery. In Sylvia S. Shen and Paul E. Lewis, editors, *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XII*, volume 6966 of *Proc. SPIE*, 2008.
- [51] Ariel Schlamm, David Messinger, and Bill Basener. Effect of manmade pixels on the inherent dimension of natural material distributions. In Sylvia S. Shen and Paul E. Lewis, editors, *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery*, Proc. SPIE, April 2009.
- [52] Ariel Schlamm, David Messinger, and William Basener. Change detection in spectral imagery using point density estimation. *submitted to JARS*, August 2009.

- [53] Ariel Schlamm, David Messinger, and William Basener. Geometric estimation of the inherent dimensionality of a single and multi-material clusters hyperspectral imagery. *JARS*, 3:033527, Feb 2009.
- [54] Ariel Schlamm, Ronald G. Resmini, David Messinger, and William Basener. A comparison study of dimension estimation algorithms. volume 7695 of *Proc. SPIE*, 2010.
- [55] John R. Schott. *Remote Sensing: The Imaging Chain Approach*, chapter 10.2: Issues of Dimensionality and Noise. Oxford University Press, 2nd edition, 2007.
- [56] Kari Sentz and Scott Ferson. Combination of evidence in dempster-shafer theory. *Sandia National Laboratories*, SAND 2002-0835, April 2002.
- [57] Sylvia S. Shen. Spectral quality metrics for vnir and swir hyperspectral imagery. volume 5093 of *Proc. SPIE*, 2003.
- [58] C. G. Simi, E. M. Winter, M. M. Williams, and D. C. Driscoll. Compact airborne spectral sensor (compass). volume 4381 of *Proc. SPIE*, pages 129–136, 2001.
- [59] R. Simmons, T. Elder, D. Stewart, E. Cincotta, C. Kennedy, and R. V. Nostrand. General spectral utility metric for spectral imagery. volume 5806 of *Proc. SPIE*, 2005.
- [60] M. O. Smith, S. L. Ustin, J. B. Adams, and A. R. Gillespie. Vegetation in deserts: A regional measure of abundance from multispectral images. *Journal of Remote Sensing of Environment*, 31:1–26, 1990.
- [61] David Snyder, John Kerekes, Ian Fairweather, Robert Crabtree, Jeremy Shive, and Stacy Hager. Development of a web-based application to evaluate target finding algorithms. In *IEEE Trans. Geosci. Rem. Sens.*, 2008.
- [62] Marcus S. Stefanou and John Kerekes. A method for assessing spectral image utility. In *IEEE Trans. Geosci. Rem. Sens.*, 2008.
- [63] Christopher M. Stellman, Frederick M. Olchowski, and Joseph V. Michalowicz. War horse (wide area reconnaissance - hyperspectral overhead real-time surveillance experiement). In *Automatic Target Recognition*, volume 4379 of *Proc. SPIE*, pages 339–346, 2001.

- [64] J. Sweet, J. Granahan, and M. Sharp. An objective standard for hyperspectral image quality. In *AVIRIS Conference*, 2000.
- [65] James Theiler. Estimating fractal dimension. *J. Opt. Soc. Am.*, 7(6):1055–1073, June 1990.
- [66] James Theiler. Quantitative comparison of quadratic covariance-based anomalous change detectors. *Appl. Opt.*, 47(28):F12–F26, 2008.
- [67] Jason E. West. Matched filter stochastic background characterization for hyperspectral target detection. Master’s thesis, Rochester Institute of Technology, 2005.
- [68] Jason E. West, David Messinger, and John Schott. Comparative evaluation of background characterization techniques for hyperspectral unstructured matched filter target detection. *Journal of Applied Remote Sensing*, 1:013520, 2007.
- [69] Amanda Ziemann, David Messinger, and William Basener. Iterative convex hull volume estimation in hyperspectral imagery for change detection. In Sylvia S. Shen, editor, *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XVII*, volume 7695 of *Proc. SPIE*, 2010.
- [70] R. Zwiggelaar and C. R. Bull. On the use of local and scalable fourier transforms, fractal dimension information, and texture segmentation. In *Image Processing and its Applications*, volume 410, pages 100–104. IEEE, July 1995.