Baker Bay, WA, Unsupervised Classification

An Analysis by Skyler Elmstrom for ESCI 442

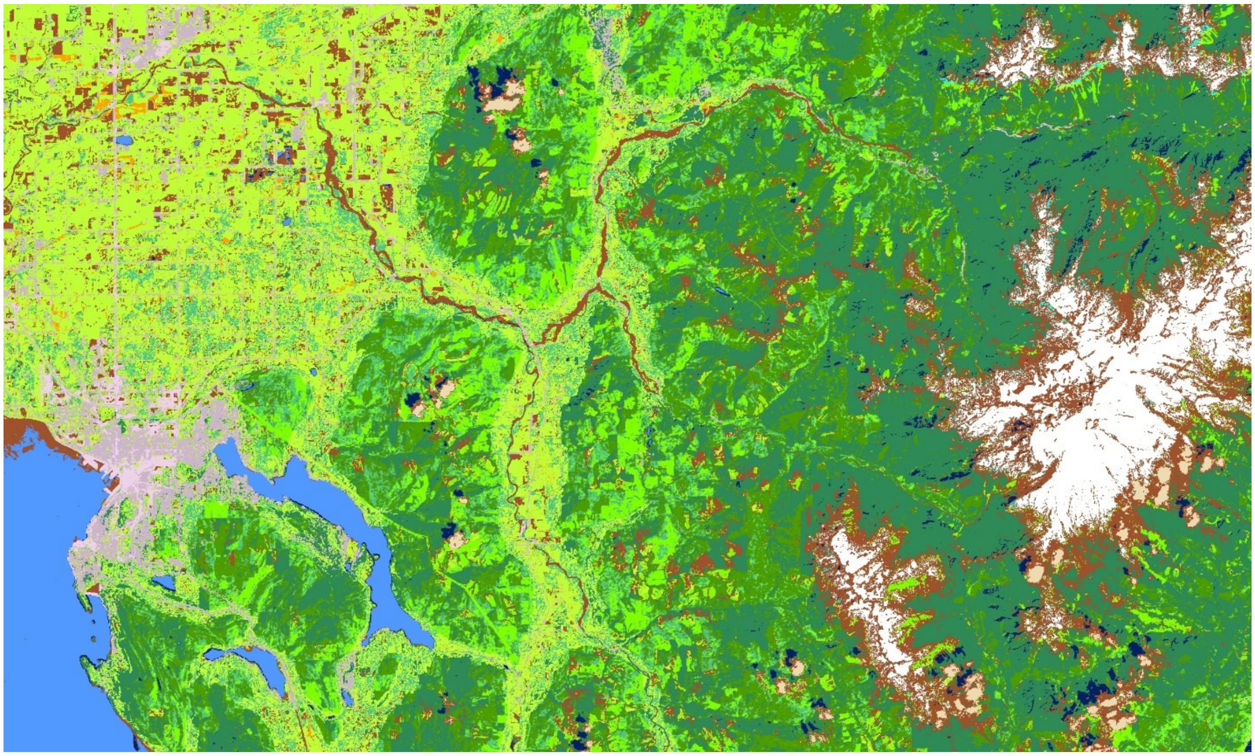

Figure 1. Baker Bay Area, WA, LANDSAT 5 scene from July 2011 following unsupervised classification using ENVI and road proximity, elevation, and slope conditional post-processing using ESRI's ArcMap.

Introduction

Extracting information from a satellite image has become a task any user with a modern desktop computer can accomplish in just a few automated steps. However, obtaining rigorously verified, categorically useful information tailored for a specific question or questions that is also aesthetically pleasing to look at requires far more patience and finesse. In this example, a LANDSAT 5 scene centered on the Bellingham Bay-to-Mount Baker area (referred to as Baker Bay in this document) underwent unsupervised classification with the assistance of ground-truth data followed by additional post-processing to correct misclassified features and noise due to spectral similarities of varying cover types (figure 1).

Methods

The image used in this analysis is a resampled LANDSAT 5 scene with a spatial resolution of 25m x 25m with minimal cloud cover from July 2011 centered at approximately N 48.800161, W 122.148154. This scene covers an area of 2378.25 km^2: the city of Bellingham, WA and Bellingham Bay mark the west side of the extent and Mount Baker marks the east.

The methods of the analysis performed follow the guidelines set out by Wallin (2018) using Harris Geospatial's ENVI in classic mode. I first ran an ISODATA unsupervised classification to generate spectral classes and assigned them preliminary information classes based on comparisons of the classified raster values to the true color and NIR composite. A "best guess" was made for each of the classes I was interested in (table 1). Information classes such as snow, water, clouds, shadows, and urban areas were generally distinguishable without additional assistance but forest types and clear-cuts, soil and rock, crops and pastures/grass, and Alpine vegetation required additional ground-truth data to discern accurately. Several information classes that were initially created contained several spectral classes. This was particularly problematic for correctly identifying clear-cuts when compared to bare soil and other forest types; pasture to crops; urban to soil and rock; shadow to water; urban to clouds; urban to residential; and soil to alpine vegetation. Ground-truth reference data provided by Wallin consists of 19 training datasets collected from 2009 through 2017. I condensed these datasets by cross-tabulating and combining similar classes such as clear-cuts from 1979-1992 and 1992-2005 for a total of 13 training datasets. Following combination, I imported these training data to ENVI as “Regions of Interest” and generated a confusion matrix to assess the accuracy of my preliminary classification. I visually inspected problem classes and found obvious flaws that were difficult to correct with reclassification alone such as crops at high altitudes, urban areas at high altitudes and along unpopulated river banks, and shadow misclassified as water on sloped surfaces. These inconsistencies were mitigated using conditional models —developed by the Western Washington University Spatial Analysis Lab— in ESRI’s ArcMap. First, I ran my classified scene through a roads conditional model to reduce confusion between urban/residential and soil/rock: urban classified values were changed to rock if they were not found within 100 meters of road. Second, I used a slope raster to correct for shadows misclassified as water: water values found where the slope was 0 were changed to shadow. Finally, I used an elevation raster to improve the quality of crop, pasture, clear-cut, and alpine classification based on several assumptions: agriculture and croplands are found below 200m, forestry activity and clear-cuts occur at 200-1500m, and alpine areas are found above the tree line at 1500m. Lastly, I generated a new confusion matrix to assess the accuracy of my post-classification processing in ENVI and ArcMap. Clouds and Shadows were not assessed for accuracy because no training data exists for them.

Results

The results of my analysis and post-processing resulted in an overall 57.4% overall classification accuracy when comparing ground-truth data to the predicted values based on the information classes I created (table 2). This value seems unchanged compared to my unprocessed accuracy of 57.5% achieved before conditional manipulations in ArcMap; the post-processed result, however, is a visual and logical improvement. After combining information classes —now totaling 11 classes; not including clouds and shadow— that contained similar spectral classes such as conifer and deciduous forest and agricultural pasture, grass, and crops my overall classification accuracy was 65.8% (table 3). There are three information classes that are notably poor in user classification accuracy —the accuracy of ground-truth pixels being identified correctly compared to the total number of pixels classified in that category: soil and rock with 3.85%, alpine vegetation with 0%, and 2005-2011 clear-cuts with 0% (table 3). These classes were spectrally mixed and often fell into more than one information class. Additionally, these classes had less training data than most classes available in the ground-truth datasets: soil and rock had 3 total and 2 were correctly assigned; alpine vegetation had 4 total and 0 were correctly assigned; and 2005-2011 clear-cut had 20 total and 0 were correctly assigned. Overall, production accuracies —predicted values matching ground-truth values divided by the total number of pixels in each class— exceeded 50% except 1992-2005 clear-cut values, 2005-2011 clear-cut values, and alpine vegetation (table 3).

Discussion

No classification can be 100% accurate; the whole purpose of classifying an image is to generalize thousands or millions of features into "information classes" and extract information specific to the task at hand. Spectral classes —sets of pixels with spectrally similar brightness values— often fall within one or several information classes. In this instance, several of my information classes had little to no spectral variability while others had significant and complicated spectral variability. For example, table 2 shows that at only 1 pixel out of 10 identified as snow by ground-truth was classified as not snow; snow and ice reflect nearly all light in the visible and near-infrared spectrum and hence have a very distinct spectral signature. Water also has a similarly unique spectral signature; this is evident in the overwhelming classification accuracy of 88.54%. Residential and urban information classes have much more spectral variability and confusion. Table 2 shows that nearly 50% of residential pixels were classified as urban and about 20% of urban pixels were classified as residential. This suggests the two classes are spectrally similar and some error will be introduced when trying to separate them. Confusion was greatest when attempting to distinguish clear-cuts from conifer or deciduous forest. As these clear-cuts recover over time, the vegetation in a clear-cut begin to resemble vegetation more characteristic of the forests they occupy or even other types of vegetation throughout the scene. Conversely, recently clear-cut land appears very similar to bare soil or rock. Some inherent error can be expected from creating information classes, but other types of error also reduce the classification accuracy of any scene. The ground-truth data was collected over a wide temporal span —2009-2017— while the scene being analyzed was captured in July of 2011. This may introduce anomalous training data that does not match the seasonality of a feature or does not exist in the same place as it once had, drastically altering the spectral comparison between what is seen in the image and what was or would be seen. It may be worthwhile to rerun this analysis using only ground-truth data collected during summer or restrict the data used for ground-truth to 2010-2012 records. Some reduction in accuracy may due to the small number of training data for some information classes such as agricultural pasture and grass, alpine vegetation, and soil and rock. These features are difficult to ground-truth because they have an additional dimension of variability —land use and land cover. Pasture could be repurposed for crops, alpine vegetation could be overwhelmed by snow and ice, and soil or rock could be covered by vegetation or developed into built up lands.

With any classification, there is always room for improvement. I was particularly disappointed with how after running conditional post-processing, the number of pasture/grass fields evaporated. Table 2 shows that production accuracy was just over 6% for pasture after post-processing yet there was 100% user accuracy —all the pasture training data fell within the group of pixels classified as pasture. Prior to conditional post-processing, my initial production accuracy was over 50%, suggesting that there really was not as much pasture/grass as it seemed or there was a problem with the thresholds used in my ArcMap model to identify pasture and crop lands. Both of those possibilities can and should be tested. In this case, having a higher resolution image of pasture and croplands from another sensor that was taken during the same timeframe as my LANDSAT image could be used to create more training data for both pasture and crops. This would allow me to check the model's output as well as my own classification accuracy. Ultimately, I ended up merging these two information classes because they are spectrally similar. However, with more information these two classes could be separated again. Another instance of misclassification exists between conifer and deciduous forest; both have similar spectral qualities. Although the ground-truth training datasets provided many conifer and deciduous forest data points, my classification still did not perform well in distinguishing these two information classes. I would attempt this classification again with a different approach to address this outcome: a high-resolution image or images of forest areas would help me create a narrower information class for conifer and deciduous and help me confirm ground-truth data; I would add a new information class called "mixed forest" to help me filter out pixels of forest that contain less dominant conifer or deciduous canopy or sparse canopy; developing more conditional post-processing models may aid in identifying different canopy types as well.

Another metric I have available for evaluating my classification's accuracy and efficacy is the total area each information class covers in my LANDSAT scene (table 4). Additional trials of this classification could be averaged with my area calculations to provide a more objective estimate but an overall accuracy of nearly 70% suggests these numbers can be used to get a generalized sense of each information class's extent.

Conclusion

My classification of the Baker Bay scene from 2011 produced reasonable results: the overall classification accuracy after post-processing the scene in ArcMap using conditional calculations for proximity to roads, elevation, and slope and combining spectrally similar classes was 65.8%. Achieving an accuracy greater than 70% would likely require a loss of information through combination of spectrally similar information classes or additional post-processing steps in ArcMap to reduce confusion between information classes. Areas with the greatest potential for improvement in this classification include crops and pasture and forest types. Moving forward, it may be worthwhile to narrow the focus of the classification and mask out areas that are out of the bounds of the questions being asked and visually compare images with higher spatial resolution to features identified in the LANDSAT scene to improve classification accuracy.

Figures and Tables

Table 1. Classification Scheme For Baker Bay, WA Scene

Class Number	Class description
1	Residential
2	Urban or Built up lands
3	Ag. Pasture/Grass
4	Crops
5	1973-1992 Clearcuts from Boyce
6	1992-2005 Clearcuts from Boyce, Grace, and Wallin
7	2005-2011 from Wallin
8	Deciduous forest
9	Conifer forest
10	Water
11	Soil/rock
12	Alpine veg., non-forest
13	Snow/ice
14	Clouds
15	Shadow

Table 2. Post-Processed Confusion Matrix

Predicted Class	Ground Truth (Reference Data)													**Row Total**		User's Acc(%
	Residential	Urban	Pasture	Crops	Deciduous	Conifer	Water	Soil/Rock	Alpine Veg.	Snow/Ice	1973-92 Cle	1992-05 Cle	2005-11 Cle			
Residential	109	34	5	3	0	5	1	1	0	0	0	1	0	159		68.55
Urban	49	179	0	0	0	0	6	0	0	0	0	0	0	234		76.5
Pasture	0	0	7	0	0	0	0	0	0	0	0	0	0	7		100
Crops	19	20	89	67	14	54	0	0	0	0	6	4	0	273		24.54
Deciduous	5	0	4	4	8	5	0	0	0	0	15	17	0	58		13.79
Conifer	4	3	1	0	0	95	3	0	0	0	12	4	1	123		77.24
Water	0	0	0	0	0	0	85	0	0	0	0	0	0	85		100
Soil/Rock	2	13	3	12	0	3	1	2	4	1	1	3	7	52		3.85
Alpine Veg.;	0	0	0	0	0	0	0	0	0	0	0	0	0	0		0
Snow/Ice	0	0	0	0	0	0	0	0	0	9	0	0	0	9		100
1973-92 Cle	0	0	0	0	0	3	0	0	0	0	48	13	0	64		75
1992-05 Cle	0	0	0	0	1	0	0	0	0	0	10	35	12	58		60.34
2005-11 Cle	0	0	0	0	0	0	0	0	0	0	0	0	0	0		0
Column Total	188	249	109	86	23	165	96	3	4	10	92	77	20	1122	**Total Samples**	
														644	**Total # Correct**	
Prod Acc (%)	57.98	71.89	6.42	77.91	34.78	57.58	88.54	66.67	0	90	52.17	45.45	0	**57.4**	**Overall Accuracy**	

Table 3. Post-Processed and Combined Class Confusion Matrix														
Predicted Class	**Ground Truth (Reference Data)**											**Row Total**		User's Acc (%)
	Residential	Urban	Pasture/Crop	Forest	Water	Soil/Rock	Alpine Veg.	Snow/Ice	1973-92 Cle	1992-05 Cle	2005-11 Cle			
Residential	109	34	8	5	1	1	0	0	0	1	0	159		68.55
Urban	49	179	0	0	6	0	0	0	0	0	0	234		76.5
Pasture/Crop	19	20	163	68	0	0	0	0	6	4	0	280		58.21
Forest	9	3	9	108	3	0	0	0	27	21	1	181		59.67
Water	0	0	0	0	85	0	0	0	0	0	0	85		100
Soil/Rock	2	13	15	3	1	2	4	1	1	3	7	52		3.85
Alpine Veg.	0	0	0	0	0	0	0	0	0	0	0	0		0
Snow/Ice	0	0	0	0	0	0	0	9	0	0	0	9		100
1973-92 Cle	0	0	0	3	0	0	0	0	48	13	0	64		75
1992-05 Cle	0	0	0	1	0	0	0	0	10	35	12	58		60.34
2005-11 Cle	0	0	0	0	0	0	0	0	0	0	0	0		0
Column Total	188	249	195	188	96	3	4	10	92	77	20	1122	**Total Samples**	
												738	**Total # Correct**	
Prod Acc (%)	57.98	71.89	83.59	57.45	88.54	66.67	0	90	52.17	45.45	0	**65.8**	**Overall Accuracy**	

Table 4. Information Class Total Area			
Class	# Pixels	Total Area (km^2)	Percentage of Area
Residential	129541	80.96	3.4
Urban	30538	19.09	0.8
Pasture/Crops	705494	440.93	18.5
1973-1992 Clear-cut	543094	339.43	14.3
1992-2005 Clear-cut	256412	160.26	6.7
2005-2011 Clear-cut	3099	1.94	0.1
Conifer/Deciduous Forest	1277520	798.45	33.6
Water	153604	96.00	4.0
Soil/Rock	452810	283.01	11.9
Alpine Vegetation	8532	5.33	0.2
Snow/Ice	181015	113.13	4.8
Cloud	25547	15.97	0.7
Shadow	37994	23.75	1.0
Total	**3805200**	**2378.25**	100

Sources

Wallin, D. 2015. *Lab III: Unsupervised Classification with ENVI*. Retrieved from http://myweb.facstaff.wwu.edu/wallin/envr442/ENVI/442_unsup_class_ENVI.html#Step4. Accessed on 2/12/2018.