



Approaches to Incorporating RHSeg Image Segmentation into Cropland Extent Mapping

James C. Tilton, NASA GSFC, Greenbelt, MD

Jun Xiong, USGS, Flagstaff, AZ

Richard Massey, NAU, Flagstaff, AZ

August 16, 2017

Overview

The original goal was to work with each analysis team to incorporate RHSeg into their analysis approaches with the hope of achieving the following benefits:

1. Removing "salt and pepper" noise from the classification maps,
2. Eliminating partially mapped fields with parts of the fields unmapped, and
3. Overall improvement in classification accuracy.

Unfortunately, I ended up actively exploring approaches for this incorporation with only four teams: the teams led by Mutlu Ozdogan (with Aparna Phalke), Russ Congalton (with Kamini Yadav), Jun Xiong, and "Teki" Sankey (with Richard Massey).

Overview (cont'd)

The previously listed interactions led to incorporating RHSeg into the analysis approaches for only two of the teams:

1. For the analysis of cropland extent of Africa by Jun Xiong, and
2. For the analysis of cropland extent for Mongolia and for North America by “Teki” Sankey and Richard Massey.

This presentation describes the collaboration between myself, Jun Xiong and Richard Massey in developing two different approaches to incorporate RHSeg into cropland mapping.

The obstacles to working with the other groups included my poor health during much of the project tenure and the difficulties of interfacing RHSeg with the Google Earth Engine – which turned out to be the preferred analysis platform.

HSeg Background

HSeg produces a hierarchical set of image segmentations with the following characteristics:

A set of segmentations that

1. consist of segmentations at different levels of detail, in which
2. the coarser segmentations can be produced from merges of regions from the finer segmentations, and
3. the region boundaries are maintained at the full image spatial resolution

The HSeg algorithm is fully described in:

James C. Tilton, Yuliya Tarabalka, Paul M. Montesano and Emanuel Gofman, “Best Merge Region Growing Segmentation with Integrated Non-Adjacent Region Object Aggregation,” *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 50, No. 11, Nov. 2012, pp. 4454-4467.

Version 1.64 of RHSeg/HSeg:

The analysis flow of the RHSeg algorithm:




Fig. 2. L_r is determined as the number times the input image must be subdivided to achieve a small enough image size for efficient processing with HSeg.

The analysis flow of the recursive function rhseg(L, X):




Fig. 3. N_{min} is equal to $\frac{1}{4}$ the number of pixels in the subimage processed at the deepest level of recursion.

Issues to be Addressed when Incorporating RHSeg

RHSeg normally produces a hierarchically related set of image segmentations over a range of segmentation detail. Any analysis approach incorporating RHSeg needs to include a method for selecting the hierarchical level that produces the appropriate level of image segmentation detail for the application.

The simplest approach is to find a method for selecting a particular merge threshold at which to stop the RHSeg region growing process. More complicated approaches perform a “pruning” of the segmentation hierarchy tree based on individual region characteristics such as region standard deviation or other more sophisticated texture measures.

Jun ended up simply selecting a pair of merge threshold values based on visually evaluating RHSeg segmentation hierarchy results for a few representative data sets - and then evaluated the final classification results for these two cases.

Richard devised a more quantitative approach for selecting a single merge threshold for different portions of his analysis areas.

Issues to be Addressed when Incorporating RHSeg

The selected segmentation results may be further refined using the “hswo” program. This program performs region object merging starting from the selected region object segmentation. Regions of size up to a selected “minimum mapping unit size” are merged with the most similar neighboring region. This selectively removes small regions from the segmentation result.

The selected segmentation result must then be classified. This can be done either by computing the region mean image and performing an object-based classification, or by performing a pixel-based classification on the original data and combining it with the selected segmentation result by a “plurality vote” over the per-pixel classification for each region (or alternatively, a percent cropland map that can be thresholded later to form a cropland extent map).

Applying RHSeg to Improving Africa Cropland Extent Analysis

Jun Xiong and I first experimented with 10m resolution Sentinel data, but despite great looking results, Jun decided to use 30m resolution data for our RHSeg analysis in order to reduce processing time, data transfer time and data storage requirements.

Jun produced 2 season 30m resolution data sets (NIR, red, green, blue – total 8 bands) for 1919 1d x 1d grids covering most of Africa. These scenes were about 3720x3720 in size.

Based on a visual inspection of about a dozen scenes processed by RHSeg to produce a wide range of hierarchical levels, we decided that RHSeg segmentation hierarchy results at merge thresholds 7.5 and 15.0 were likely to produce good results.

Applying RHSeg to Improving Africa Cropland Extent Analysis


We post-processed the RHSeg outputs at thresholds 7.5 and 15.0 with hswc using a minimum mapping unit of 6.

RHSeg usually took less than 15 minutes to process each of Jun's Africa images with 64 CPUs on Discover or Pleiades. Water masks were needed for images with large bodies of water to avoid long processing times for those scenes.

I wrote a C++ program (called combineff) that output the region crop percentage over each region. Jun labeled the entire region as cropland if the region was $\geq 85\%$ cropland and non-cropland if the region was $\leq 15\%$ cropland, and did not modify the cropland labelings when the region was between 15 and 85% cropland.

(The combineff program runs quickly on a single processor Linux workstation.)


1d x 1d gridding of Africa



- Total **1919** grids by a 1° by 1° for non-desert area.
- Each grid is about **3720 x 3720** pixels
- **2 season** mosaics (season 1: March-June, season 2: July-Oct) with NIR, red, green and blue bands.
- Utilizing both the Pleiades and Discover clusters, all segmentation processing is finished within **10** days.



Merging pixel-based and object-based



Google Imagery	Pixel-base classification
Rhseg Segmentation	Merging Results



Results of Applying RHSeg to Africa Cropland Extent Mapping

(Classification Accuracies with respect to the crowd source reference data set produced by “CroplandReference.”)

Without RHSeg/hsw0:

	Non-Cropland	Cropland	Total	User Accuracy %
Non-Crop	6704	45	6749	99.3%
Cropland	280	264	544	48.5%
Total	6984	309	7293	
Producer's Accuracy %	95.9%	85/4%		
Overall Accuracy %				95.5%

RHSeg/hsw0 with merge threshold = 7.5:

	Non-Cropland	Cropland	Total	User Accuracy %
Non-Crop	6692	57	6749	99.1%
Cropland	244	300	544	55.1%
Total	6936	357	7293	
Producer's Accuracy %	96.4%	84.0%		
Overall Accuracy %				95.8%

RHSeg/hsw0 with merge threshold = 15.0:

	Non-Cropland	Cropland	Total	User Accuracy %
Non-Crop	6636	113	6749	98.3%
Cropland	139	405	544	74.4%
Total	6775	518	7293	
Producer's Accuracy %	97.9%	78.1%		
Overall Accuracy %				96.5%

Applying RHSeg to Improving Cropland Extent Analysis for North America and Mongolia

Richard devised a quantitative approach to selecting the best threshold value for selecting the output from RHSeg.

The analyzed area was subset into study area zones. For each zone:



1. 35 to 50 field samples were hand digitized and labeled, distributed across each study area zone (average field sample size was about 250 pixels).
2. Landsat 5 TM maximum (85%) NDVI value composite data from two (or three) seasons in the nominal year 2010 was constructed as a seamless mosaic.

The bands included were bands 1, 2, 3, 4, 5, and 7 for two seasons (summer and spring) making a total of 12 bands. (For Mongolia, three seasons were used, making a total of 18 bands).


See the next two slides for more detail...

Study Area Zones

- 1) **USA** – 9 zones (22 sub-zones):
USDA farm resource regions 2000
- 2) **Canada** – 3 zones (8 sub-zones):
Canada census of agriculture 2011
and vegetation regions 1998
- 3) **Mexico** – 6 zones (9 sub-zones):
Mexico farm sizes map - INEGI
2007 (Instituto Nacional de
Estadística y Geografía)
- 4) **Central America & the Caribbean**
– 5 zones: FAO global agro-
ecological zones 2000
- 5) **Other-** Alaska, and Hawaii
- 6) **Mongolia**



RHSeg: Digitized field samples



- Individual fields are extracted within the cropland extent
- 35-50 digitized field samples per sub-zone

Applying RHSeg to Improving Cropland Extent Analysis for North America and Mongolia

The Landsat study zone mosaics were then subset into multiple 1 deg. by 1 deg. tiles for processing by RHSeg (with a 0.1 degree buffer on all sides).

Corresponding training field sample tiles were also generated from the training field sample vector data for the study zone.

For tiles that contained training fields, RHSeg was run so as to produce multiple hierarchical segmentation levels with the number of regions ranging from 220,000 regions down to 10,000 regions.

Applying RHSeg to Improving Cropland Extent Analysis for North America and Mongolia

The new “hsegrefcomp” program was then run on the RHSeg outputs and the training field sample tiles, producing an ASCII output table, for example:

h_level	classes	objects	merge thresh	sample label	sample pixels	object label	object pixels	overlap pixels	error_rate
0	220000	232069	7.06377	41	472	154558	272	223	0.400538
0	220000	232069	7.06377	11	333	205361	132	132	0.432258
1	136481	148297	9.66024	11	333	205361	132	132	0.432258
1	136481	148297	9.66024	41	472	154558	421	223	0.50056
2	87306	99275	13.0247	11	333	205361	132	132	0.432258
2	87306	99275	13.0247	41	472	133482	443	224	0.510383
3	71051	83272	14.8438	11	333	175755	252	252	0.138462
3	71051	83272	14.8438	41	472	133482	443	224	0.510383
4	61716	74217	16.2227	11	333	175755	252	252	0.138462
4	61716	74217	16.2227	41	472	133482	443	224	0.510383
5	56450	69114	17.1291	41	472	133482	443	224	0.510383
5	56450	69114	17.1291	11	333	175755	252	252	0.138462

the error rate is defined as (False positive pixels + False negative pixels)/(Total number of pixels in the sample). For the first line of the table this is 0.400538 (= ((472 - 223) + (272 - 223))/(472 + 272)).


Applying RHSeg to Improving Cropland Extent Analysis for North America and Mongolia

The new “find_best_thresh” program was then run on the outputs from hsegrefcomp for all tiles processed in the zone. This program produces a table of average and median error rate values over mg_thresh values. For example:

This table is imported into MS Excel and plotted (next slide).

mg_thresh_center	ave_error_rate	median_error_rate
2.5356	0.401465	0.371429
4.7857	0.363679	0.299838
7.0358	0.336315	0.346939
9.2859	0.310911	0.279748
11.536	0.325855	0.323529
13.7861	0.298126	0.299544
16.0362	0.311379	0.323529
18.2863	0.362195	0.323529
20.5364	0.315011	0.338235
22.7865	0.303601	0.338235
25.0366	0.402678	0.398295
27.2867	0.409878	0.397571
29.5368	0.276404	0.239146
31.7869	0.413968	0.397571
34.037	0.440116	0.581893
36.2871	0.406849	0.494
40.7873	0.305532	0.494

Applying RHSeg to Improving Cropland Extent Analysis for North America and Mongolia



Based on a minimum error rate in a stable portion of the curve, the threshold value of 14.0 was selected.

Applying RHSeg to Improving Cropland Extent Analysis for North America and Mongolia

RHSeg was then run on all 1 deg. by 1 deg. tiles so as to produce single segmentation result at the selected merge threshold value.

The RHSeg outputs were then post-processed with hsw0 using a minimum mapping unit of 9.

The combineff was then run on the outputs from RHSeg/hsw0 and the pixel-wise Random Forest classification of cropland extent, producing region crop percentage over each RHSeg/hsw0 region object.



Richard ended up using a threshold of 25% for determining whether a region object was cropland.

Richard also eliminated region objects of 9 pixels or less and 2000 pixels or more (the number of pixels in each region object is a standard output from RHSeg). Note that hsw0 previously eliminated regions of less than 9 pixels by merging them into the most similar neighboring region.


The overall object-based classification approach is summarized on the following slide...

Object-based classification: RHseg on NAU Cluster Monsoon

- Landsat 5 max NDVI composite
 - Two seasons: 90-180, 180-270
 - Bands per season: B2, B3, B4, B5, B7
- RHseg tiles:
 - 1deg x 1deg tiles
 - Each study area zone divided into tiles
- Digitized field boundaries to select segmentation hierarchy
- RHSeg objects:
 - Objects with >9 and <2000 pixels were extracted
 - Objects with $<25\%$ cropland pixels were removed



RHSeg: Improvements to the pixel-based classification



(Location: Zone 7 Texas, USA)

