

Feature Extraction with Example-Based Classification Tutorial

In this tutorial, you will use Feature Extraction to extract rooftops from a multispectral QuickBird scene of a residential area in Boulder, Colorado. Feature Extraction provides a quick, automated method for identifying rooftops, saving an urban planner or GIS technician from digitizing them by hand. If you need more information about a particular step, click the blue Help button to access ENVI Help.

Files Used in This Tutorial

Tutorial files are available from our [ENVI Tutorials](#) web page. Click the **Feature Extraction** link to download the .zip file to your machine, then unzip the files. You will use the file named `qb_colorado.dat` for this tutorial.

This is a pan-sharpened QuickBird image (0.6-meter spatial resolution) of Boulder, Colorado, acquired on 04 July 2005. QuickBird images are courtesy of DigitalGlobe and may not be reproduced or distributed without prior written permission of DigitalGlobe.

Background

Feature Extraction uses an object-based approach to classify imagery, where an object (also called *segment*) is a group of pixels with similar spectral, spatial, and/or texture attributes. Traditional classification methods are pixel-based, meaning that spectral information in each pixel is used to classify imagery. With high-resolution panchromatic or multispectral imagery, an object-based method offers more flexibility in the types of features to extract.

The workflow involves the following steps:

- Dividing an image into segments
- Computing various attributes for the segments
- Creating several new classes
- Interactively assigning segments (called *training samples*) to each class
- Classifying the entire image with a K Nearest Neighbor (KNN), Support Vector Machine (SVM), or Principal Components Analysis (PCA) supervised classification method, based on your training samples.
- Exporting the classes to a shapefile or classification image.


Start the Workflow

1. From the menu bar, select **File > Open**.
2. Navigate to `feature_extraction`, and select the file `qb_colorado.dat`. Click **Open**. The image is displayed at full resolution.
3. From the **Optimized Linear** drop-down list in the toolbar, select **Linear 2%**. This type of stretch brightens the image, making it easier to see individual features.
4. From the Toolbox, select **Feature Extraction > Example Based Feature Extraction Workflow**. The Data Selection panel appears.
5. The filename is already listed in the **Raster File** field. Click **Next**. The Object Creation panel appears.


Segment the Image

Segmentation is the process of dividing an image into segments that have similar spectral, spatial, and/or texture characteristics. The segments in the image ideally correspond to real-world features. Effective segmentation ensures that your classification results are more accurate.

1. Use the drop-down list in the main toolbar to zoom the image to **200% (2:1)**.
2. Enable the **Preview** option. A Preview Window appears, showing the initial segments from the image, colored in green. The following image shows an example of a Preview Window centered over a residential area:



The initial segmentation shown in the Preview Window delineates the rooftops, but there are too many segments, which could potentially increase processing time. Use a combination of segmentation and merging to reduce the number of segments while providing clear boundaries around rooftops.


3. The rooftops in this image appear much darker and provide better contrast against surrounding features in a near-infrared image. Applying segmentation to only the near-infrared band will produce cleaner boundaries around the rooftops. In the Object Creation panel, click the  button next to **Select Segment Bands**. Then select the **NIR** band and click **OK**.

The **Edge** segmentation method draws lines along the strongest intensity gradients, making it an effective edge detector. The **Full Lambda Schedule** merging method then combines adjacent segments, as you increase the merging value.

You can achieve reasonable results with the following settings:

- Segment Algorithm: **Edge**
- Scale Level: **35**
- Merge Algorithm: **Full Lambda Schedule**
- Merge Level: **80**

The following image shows the resulting segmentation using these settings:



If you were to leave the **Scale Value** at the default value of 50, some rooftop segments would be combined with segments representing adjacent backyards or trees because they have a similar intensity. So a value of 35 works well to separate the two features.

4. When you are satisfied with the segmentation, click **Next**. ENVI creates and displays a segmentation image (also called the Region Means image in the Layer Manager). Each segment is assigned the mean spectral values of all the pixels that belong to that segment.

Select Training Samples

In this step, you will select representative samples of different features and assign them to various classes before performing a supervised classification.


When segmentation is complete, the Example-Based Classification panel appears with one undefined class (New Class 1). As you move the mouse around the segmentation image, the objects underlying the cursor are highlighted in cyan. You may need to click once on the image to activate this function.

1. Disable the **Preview** option.
2. In the Class Properties table, change the **Class Name** to **Rooftop** and press the **Enter** key.


3. Click on at least 20 different segments that represent rooftops. Try to pick a variety of sizes, shapes, colors, and intensities. The more training samples you select, the better the results from the classification. The following image shows an example:



The following are some tips for selecting training samples:


- Click again on a segment to remove it from selection.
- If individual segments are hard to discern from each other in certain areas, enable the **Show Boundaries** option to draw boundaries around each segment.
- If the segmentation image does not provide enough detail to determine if segments represent rooftops (versus a driveway or backyard, for example), uncheck the Region Means image in the Layer Manager. The original QuickBird image will display instead. The above image shows this view.
- If you pan or zoom around the image, remember to click the **Select** icon  in the main toolbar before selecting training regions.

Next, you should define several more classes that are not rooftop.

1. In the Example-Based Classification panel, click the **Add Class** button .
2. In the Class Properties table, change the new class name to **Grass** and press **Enter**.


3. Select **Grass** in the left side of the panel, then select at least 20 training samples from the image that represent grassy areas such as backyards, fields, and parks.
4. Repeat steps 1-3 for the following classes: **Concrete** (curbs and driveways) and **Road** (asphalt only). Select at least 20 training samples each, and change the class colors as desired. The following shows an example of the classification panel; your results may be different:



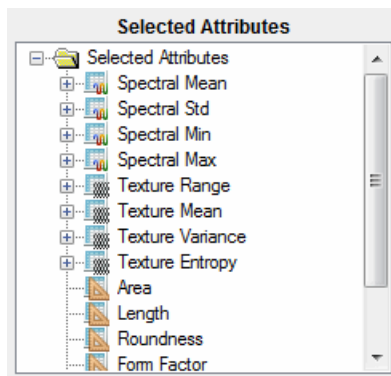
5. Click the **Save an Example file** button , and select an output folder and filename for the training regions you have defined. You can restore this file later if you want to continue where you left off.

Select Attributes for Classification

For each segment, ENVI computes various spatial, spectral, and texture attributes. In this step, you can choose which attributes to use in the supervised classification. By default, all attributes will be used. See "List of Attributes" at the end of this tutorial for definitions of all available attributes.

1. Select the Attributes Selection tab in the Example-Based Classification panel.
2. For this tutorial, you can let ENVI determine the most appropriate attributes to classify with by clicking the **Auto Select Attributes** button . After a brief moment, the


Selected Attributes column updates to show which attributes will be used. The following image shows an example; your results may be different.



Select a Classification Method

Feature Extraction offers three methods for supervised classification: K Nearest Neighbor (KNN), Support Vector Machine (SVM), or Principal Components Analysis (PCA). For this tutorial, you will use KNN, which classifies segments based on their proximity to neighboring training regions. It is more rigorous than the other methods, and it more accurately distinguishes between similar classes.

1. Use one of the Zoom tools in the toolbar to zoom to **100%**.
2. Enable the **Preview** option in the Example-Based Classification panel. A Preview Window appears with the current classification results. As you make changes to the training data, attributes, and classification settings, the classification results will automatically update. Move the Preview Window around the image or resize it to view results for different areas. The following image shows an example.



3. Disable the **Region Means** option in the Layer Manager, to hide the segmentation image. The original QuickBird image will display underneath the Preview Window.
4. Right-click inside the Preview Window and select **Blend**. The classification image gradually fades in and out to show the QuickBird image. Use the ◀ and ▶ buttons to slow down or speed up the blend, respectively. Click the Pause button || to stop the blend.

Black segments are those that the classifier could not determine a suitable class for, so they remain unclassified. The **Allow Unclassified** option under the Algorithms tab controls whether or not to force all segments into one of the classes you defined.

You should focus on the accuracy of the Rooftop classification, since that is the feature you are extracting. The initial classification shown in the Preview Window shows many segments incorrectly identified as Rooftop. You have a couple of options to reduce these errors.


Adjust KNN Settings

1. Select the Algorithms tab.
2. Try increasing the **Threshold** value. The default value is 5 percent, which means segments that have less than 5 percent confidence in each class are set to "unclassified." As you increase the **Threshold** slider, the classifier will allow more unclassified segments. As you decrease the **Threshold** slider, the classifier forces more segments into classes, thus creating more opportunity for misclassification.

3. Increase the **Neighbors** value to **3**. A higher value takes into account more neighbors when choosing a target class and should reduce noisy or irrelevant features.


Collect More Training Samples

If you have segments that are persistently being misclassified, you can assign them to known classes by following these steps:

1. Select the Examples Selection tab.
2. On the left side of the Example-Based Classification panel, select the class name that you want to assign the segment to.
3. Ensure that the **Select** icon  is active in the toolbar.
4. In the image display, click on the segment to assign it to the selected class.

Define a New Class

You will probably notice that some trees and grassy areas are incorrectly classified as Rooftop. Try creating a new class called **Trees** and select training samples for that class:

1. In the Example-Based Classification panel, click the **Add Class** button .
2. In the Class Properties table, change the new class name to **Trees** and press **Enter**.
3. Select **Trees** in the left side of the panel, then select at least 20 training samples of trees, including their shadows.

Continue to experiment with these options until you are satisfied with the classification of rooftops. Then click **Next**. In the next step, you will export a Rooftop class to a raster file.

Set Output Options

The Output panel lets you export various components of the supervised classification to raster images and/or shapefiles. For this tutorial, you will export a classification image to ENVI raster format, then overlay the classes on the original image. By default, the classification image will include all of the classes that you defined.

1. In the Export Vector tab, disable the **Export Classification Vectors** option.
2. Click the Export Raster tab.
3. Enable the **Export Classification Image** option, and select an output directory to save the file. Change the output filename to **qb_colorado_classimage**.

4. Click **Finish**. ENVI adds a new layer called `qb_colorado_classimage.dat` to the Layer Manager. To view this layer over the original image, use the Transparency slider in the main toolbar to increase the transparency of the classification image.

The example-based classification did a fairly good job of extracting rooftops, but like any automated method, it did not extract every segment that represents rooftops. In some areas, the texture and brightness of rooftops was similar to adjacent yards, so the two features may have been combined into one class. Or, other features may have been misclassified as Rooftop. If you are not satisfied with your results, try experimenting with different segmentation and merging values, try a different classification method, or collect more training samples.

List of Attributes

Spectral Attributes

Spectral attributes are computed on each band of the input image. The attribute value for a particular segment is computed from input data band where the segmentation label image has the same value (i.e., all pixels in the same segment contribute to the attribute calculation).

Attribute	Description
Spectral_Mean	Mean value of the pixels comprising the region in band x
Spectral_Max	Maximum value of the pixels comprising the region in band x
Spectral_Min	Minimum value of the pixels comprising the region in band x
Spectral_STD	Standard deviation value of the pixels comprising the region in band x

Texture Attributes

Texture attributes are computed on each band of the input image. Texture attribute computation is a two-step process where the first pass applies a square kernel of pre-defined size to the input image band. The attributes are calculated for all pixels in the kernel window and the result is referenced to the center kernel pixel. Next, the attribute results are averaged across each pixel in the segment to create the attribute value for that band's segmentation label.

Attribute	Description
Texture_Range	Average data range of the pixels comprising the region inside the kernel (whose size you specify with the Texture Kernel Size parameter in segmentation)
Texture_Mean	Average value of the pixels comprising the region inside the kernel
Texture_Variance	Average variance of the pixels comprising the region inside the kernel
Texture_Entropy	Average entropy value of the pixels comprising the region inside the kernel

Spatial Attributes

Spatial attributes are computed from the polygon defining the boundary of the segment, so band information is not required.

Attribute	Description
Area	<p>Total area of the polygon, minus the area of the holes.</p> <p>If the input image is pixel-based, the area is the number of pixels in the segmented object. For a segmented object with 20 x 20 pixels the area is 400 pixels.</p> <p>If the input image is georeferenced, the area is in the map units of the input image. For a segmented object with 20 x 20 pixels, where the input image pixel resolution is 2 meters, the total area is 1600 square meters (400 pixels x 2 meters x 2 meters).</p>
Length	<p>The combined length of all boundaries of the polygon, including the boundaries of the holes. This is different than the Major_Length attribute.</p> <p>If the input image is pixel-based, the length is the number of pixels. For a segmented object with 20 x 20 pixels, the length is 80 pixels.</p> <p>If the input image is georeferenced, the length is in the map units of the input image. For a segmented object with 20 x 20 pixels, where the input image pixel resolution is 2 meters, the length is 160 meters (80 pixels x 2 meters).</p>
Compactness	<p>A shape measure that indicates the compactness of the polygon. A circle is the most compact shape with a value of $1 / \pi$. The compactness value of a square is $1 / 2(\sqrt{\pi})$.</p> <p>$Compactness = \sqrt{4 * Area / \pi} / \text{outer contour length}$</p>
Convexity	<p>Polygons are either convex or concave. This attribute measures the convexity of the polygon. The convexity value for a convex polygon with no holes is 1.0, while the value for a concave polygon is less than 1.0.</p> <p>$Convexity = \text{length of convex hull} / \text{Length}$</p>
Solidity	<p>A shape measure that compares the area of the polygon to the area of a convex hull surrounding the polygon. The solidity value for a convex polygon with no holes is 1.0, and the value for a concave polygon is less than 1.0.</p> <p>$Solidity = Area / \text{area of convex hull}$</p>

Attribute	Description
Roundness	<p>A shape measure that compares the area of the polygon to the square of the maximum diameter of the polygon. The "maximum diameter" is the length of the major axis of an oriented bounding box enclosing the polygon. The roundness value for a circle is 1, and the value for a square is 4 / pi.</p> $\text{Roundness} = 4 * (\text{Area}) / (\text{pi} * \text{Major_Length}^2)$
Form_Factor	<p>A shape measure that compares the area of the polygon to the square of the total perimeter. The form factor value of a circle is 1, and the value of a square is pi / 4.</p> $\text{Form_Factor} = 4 * \text{pi} * (\text{Area}) / (\text{total perimeter})^2$
Elongation	<p>A shape measure that indicates the ratio of the major axis of the polygon to the minor axis of the polygon. The major and minor axes are derived from an oriented bounding box containing the polygon. The elongation value for a square is 1.0, and the value for a rectangle is greater than 1.0.</p> $\text{Elongation} = \text{Major_Length} / \text{Minor_Length}$
Rectangular_Fit	<p>A shape measure that indicates how well the shape is described by a rectangle. This attribute compares the area of the polygon to the area of the oriented bounding box enclosing the polygon. The rectangular fit value for a rectangle is 1.0, and the value for a non-rectangular shape is less than 1.0.</p> $\text{Rectangular_Fit} = \text{Area} / (\text{Major_Length} * \text{Minor_Length})$
Main_Direction	<p>The angle subtended by the major axis of the polygon and the x-axis in degrees. The main direction value ranges from 0 to 180 degrees. 90 degrees is North/South, and 0 to 180 degrees is East/West.</p>
Major_Length	<p>The length of the major axis of an oriented bounding box enclosing the polygon. Values are map units of the pixel size. If the image is not georeferenced, then pixel units are reported.</p>
Minor_Length	<p>The length of the minor axis of an oriented bounding box enclosing the polygon. Values are map units of the pixel size. If the image is not georeferenced, then pixel units are reported.</p>
Number_of_Holes	<p>The number of holes in the polygon. Integer value.</p>

Attribute	Description
Hole_Area/Solid_Area	<p>The ratio of the total area of the polygon to the area of the outer contour of the polygon. The hole solid ratio value for a polygon with no holes is 1.0.</p> <p>$\text{Hole_Area/Solid_Area} = \text{Area} / \text{outer contour area}$</p>