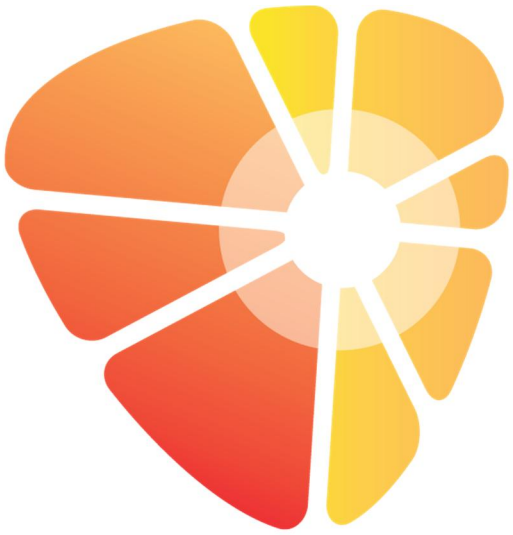

Scyven

Powered by Scyllarus

User Guide

For Scyven Version 1.3.0

Publication Date 24 June 2016

scyven.com

Table of Contents

1 Introduction

The Scyven Application has been created to provide both a convenient way of accessing the Scyllarus hyperspectral imaging engine, as well as a rich and fully featured Hyperspectral Image processing workflow.

Scyven additionally provides functionality for managing multiple Hyperspectral images in one session, for reading and writing Hyperspectral image files in ENVI, TIF (and GeoTIFF) and HSZ formats and for displaying the results of the Scyllarus pipeline processing.

2 The Main Scyven Screen

The default main screen of Scyven will look like this when first launched:

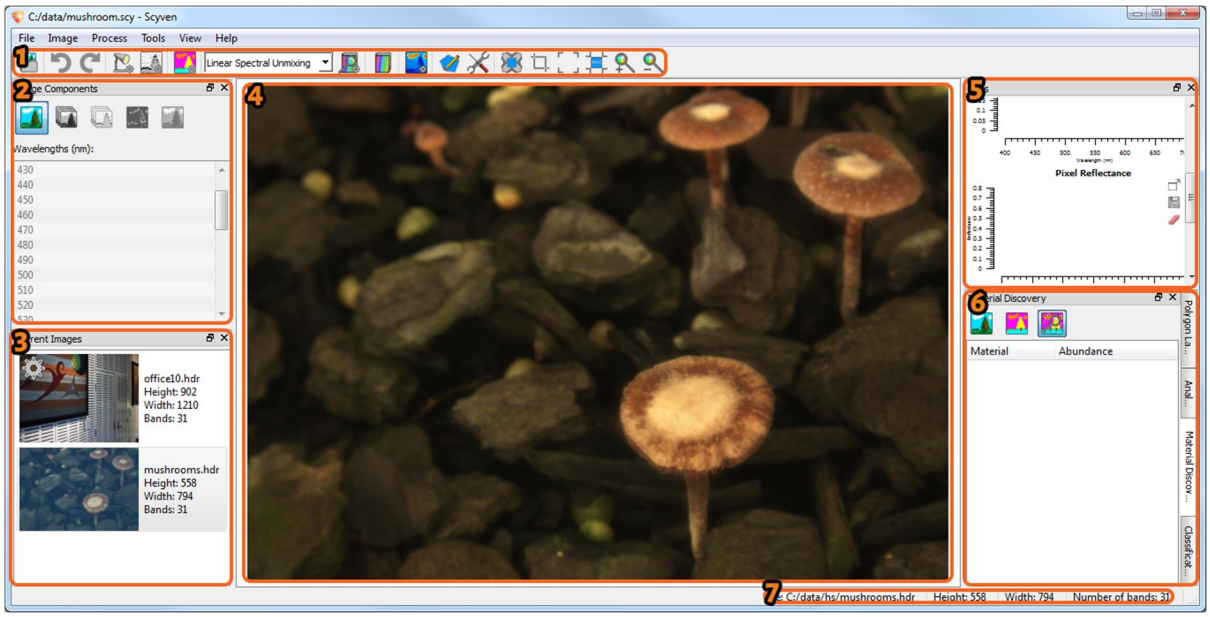

Figure 1: Scyven Main Screen

1. Main Toolbar
2. Image Components Browser
3. Current Images Browser
4. Image Display Window
5. Spectral Plots Panel
6. Analysis Tools Panel (Material Discovery, Classification, Analysis)
7. Status Bar.

All of the panels can be popped out, re-arranged, and closed. To re-open a closed panel, use the view menu.

Panels 1-6 are described in the following sub-sections. Panel 7 is the status bar - this displays information about the currently selected image (location on disk, and dimensions).

2.1 The Main Toolbar

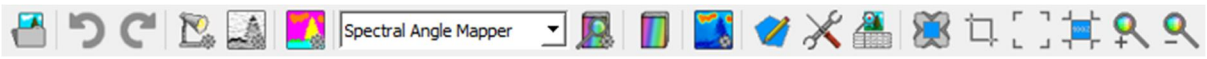

Figure 2: Main Toolbar

The items in the main toolbar are as follows (with links to the relevant sections in this document):

- Open Image: 4.2.6;
- Undo Image Operation: 5.2.1;
- Redo Image Operation: 5.2.1;
- Run Illuminant Recovery: 6.2.2;
- Run Reflectance Recovery: 6.2.2;
- Run Material Discovery: 6.2.2;
- Classification Selection pull-down and Run Classification: 7;
- Open Spectral Library Manager: 8.1;
- Run Analysis: 7;
- Draw a Polygon: 8.3.1;
- Open Scyllarus Settings: 6.2.1;
- Open Pixel Inspector: 8.2
- Resize Image: 5.2.4;
- Crop Image to Selection: 5.2.5;
- Fit Image to Image Display Window: 9;
- View Image at 100% (pixelwise): 9;
- Zoom in Image: 9;
- Zoom out Image: 9.

2.2 Panels

2.2.1 Image Components Browser

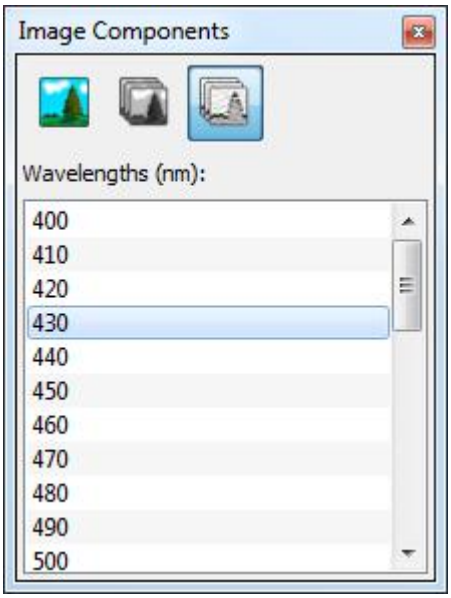

The Image components Panel (Area 2 in Figure 1: Scyven Main Screen) allows you to view different image components and bands.

The five buttons at the top of the window allow you to view the different image components. Clicking one will change the view in the main window. From left to right they are:

- Color Image;
- Irradiance (by band);
- Reflectance (by band);

The highlights/specularity map and the shading map image components can be found under the View menu. To view a specific band of Irradiance or Reflectance, select that band number from the list after clicking the relevant button. You can also scroll through bands using the up and down arrows on the keyboard, once one band has been selected. If an irradiance image is opened, the Reflectance, Highlights and Shading will not be available for un-processed images (as they are recovered by the processing). If a reflectance image is opened, the Highlights and Shading components will not be available. For more information on each Image Component, see section 5.1.

2.2.2 Image Display Window

The Image Display Window (Area 4 in Figure 1: Scyven Main Screen) displays the currently viewed image (according to the selection in the Image Components Browser) along with any Material overlays and polygons. Selections for cropping, drawing polygons and viewing plot data are made within this area using the mouse. You can zoom the image using the mouse scroll wheel.

2.2.3 Current Images Browser

Figure 3: Current Images Browser

The Current Images Browser (Area 3 in Figure 1: Scyven Main Screen) lists all the images that are currently open in your workspace. The currently viewed image is selected within the list (the top image in Figure 3: Current Images Browser. Basic image details are displayed next to each image: the filename, and dimensions.

If the image has already been processed, a small cog icon will appear on it.

Mousing over the preview of an image will display the processing parameters that were used in a tooltip (Figure 4: Processing Information Tooltip).

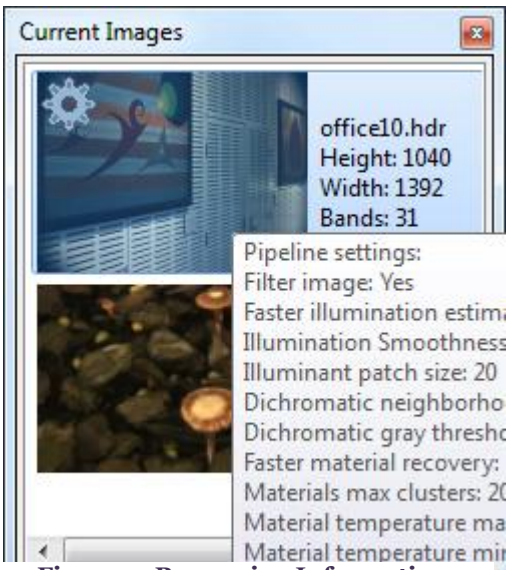

Figure 4: Processing Information Tooltip

2.2.4 Polygon Label Browser

The Polygon Label Browser (Area 6 in Figure 1: Scyven Main Screen) is used to view user created polygons. It operates in much the same way as the Material Discovery Browser (2.2.5). If you want to read about creating and using polygon labels, please see section 0.

2.2.5 Material Discovery Browser

The Material Discovery Browser (Area 6 in Figure 1: Scyven Main Screen) is used to view where material clusters lie upon the image. Once an image has been processed with the 'Run Material Discovery' button (see section 6), or a previously processed image has been loaded, the Material Discovery Browser will become populated with a list of materials.

The three buttons can be used to quickly turn ALL the material overlays on and off and toggle transparency:

- The left button removes all material overlays,
- The middle button displays all material overlays,
- The right button toggles the layer transparency – when on, materials will have a transparency that is proportional to their 'abundance' at that pixel.

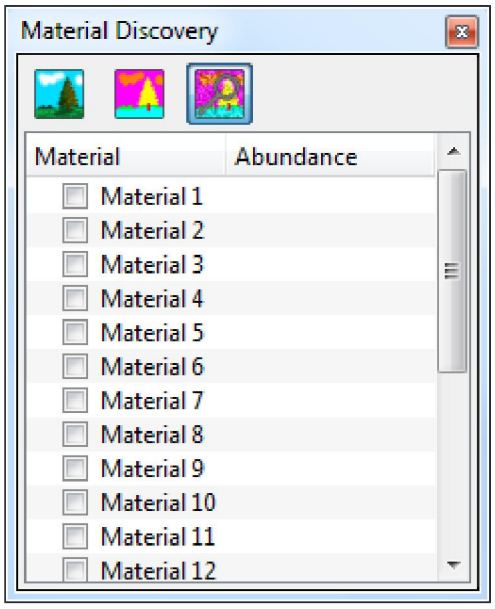

Figure 5: Image Materials Browser

You can select a few materials to see where they lie upon the image:

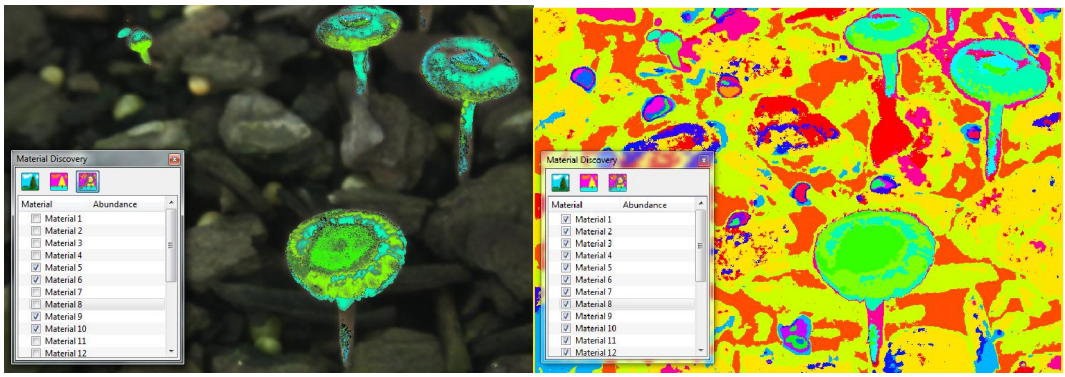

Figure 6: The same image with (a) 3 materials selected and (b) all materials selected.

Any combination of materials may be selected or viewed using the checkboxes, and the overlays can be easily removed by pressing the 'remove all materials' button.

2.2.6 Classification Browser

The Classification Browser works in the same way as the Material Discovery Browser. Once an image has been processed with the 'Run Classification Button' (see section 7) the Classification Browser will become populated with a list of Classes. The buttons in the Classification Browser operate in the same way as the Material Discovery ones (see section 2.2.5)

2.2.7 Analysis Browser

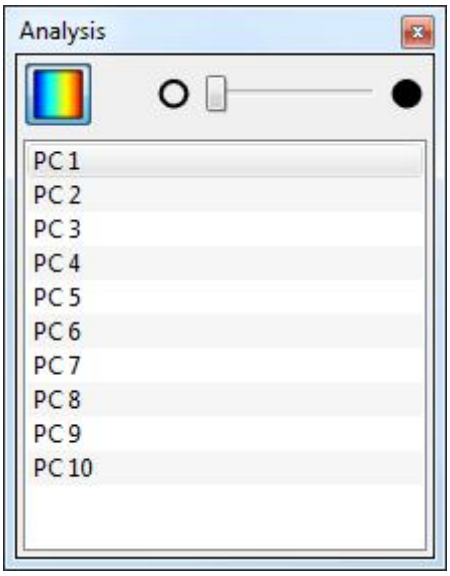

Figure 7: Analysis Browser

The Analysis Browser is used to view the results of analysis algorithms performed on images. It is currently used to view the results of Principal Component Analysis (PCA). Once an image has been processed with PCA using the 'Run Analysis' button (see section 7), the Analysis Browser will become populated with a list of Analysis results.

To view the PCA results, click the result labels (PC1, PC2 etc). The Colour Toggle button can be used to turn the colourisation of the results on and off, and the scale slider is used to adjust the scale of the result. More information about using and interpreting the results of PCA can be found in section 7.

2.2.8 Plots

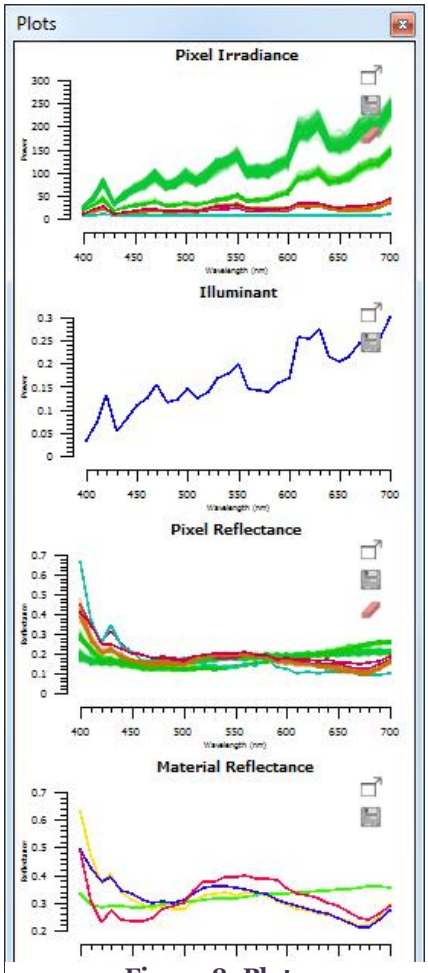

Figure 8: Plots

The plots panel can be used to view spectral profiles of images quickly within the Scyven Main Window. Clicking on the image in the Image Display Window adds data to the plots. You can pop out each plot by clicking the 'pop-out' button in the upper right corner. You can save each plot using the save button, and the Pixel plots can be cleared using the eraser button.

For more information on the Spectral Plots, see section 6.2.5

3 Quick Tutorial

Scyven has a built-in tutorial that can be accessed via the Help menu. The tutorial will walk you through the process of opening an image, pre-processing, running Scyllarus, and then analysing the results. The image used for the tutorial is the 'mushrooms.hdr' image (shown in screenshots in this manual) which is included with the Scyven Installer.

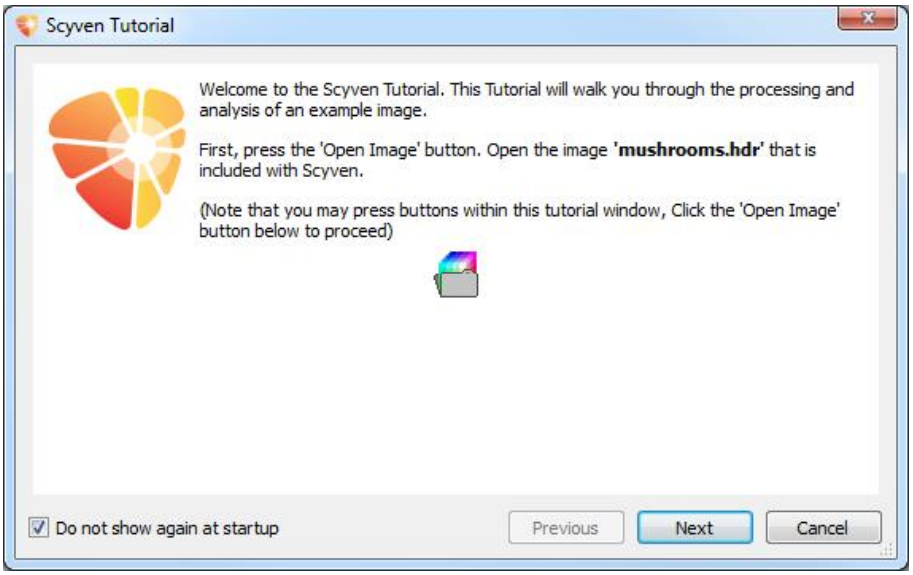

Figure 9: Scyven Tutorial

4 File Types and the File Menu

4.1 Hyperspectral Image Files

Scyven supports loading and saving of Hyperspectral image files. When an image is loaded, the color reconstruction is shown in the Image Display Window.

The following section contains a brief overview of the formats that Scyven is compatible with.

4.1.1 HDR/ENVI Image Files

See http://www.exelisvis.com/docs/ENVIHeaderFiles.html for information about this file format.

There are two related files to describe a HDR/ENVI Image File:

- The Header file, which has extension `.hdr`.
- The image data file, which can have any extension.

When you open an `.hdr` file, Scyven will search for another file with the same base name and treat the first one it finds as the related image file. If either file cannot be located or opened, the process will fail.

4.1.2 HSZ Image Files

See http://scyllarus.research.nicta.com.au/data/hsz-data-format/ for detailed information about this NICTA-developed file format.

The HSZ format stores the Hyperspectral Image in the decomposed form that is calculated when the image is processed. It also stores the results of the processing (material classification, reflectance, shading etc). The format is compressed, and takes up significantly less space per image than an HDR/ENVI format image, but there is some loss of information involved (see the link above for more information about this).

4.1.3 TIFF Image Files

See http://www.gdal.org/frmt_gtiff.html for detailed information about this file format.

TIFF Images can contain multiple bands of information as well as associated geo-spatial information. When you open a TIFF image in Scyven, the default wavelengths are set as 0, 1 2 .. etc as the TIF image files do not store this information. If you would like to manually add the wavelengths information, create a .hdr file following the ENVI format specified above. The file

must have the same basename as the TIFF file associated with it. When you open the TIFF file, Scyven will look for any matching HDR and use that to determine the wavelengths associated with the TIFF image. The easiest way to create the .hdr file is to copy and modify one from an existing HDR/ENVI image pair.

4.1.4 Workspaces

Workspaces allow you to save and load collections of images, results, and annotations quickly and easily, as well as across machines and platforms.

A Scyven Workspace file contains Images and session parameters describing how those image files were being used and displayed by Scyven at the time the Workspace was saved. Workspace files may be created in two different ways (use the option in Scyllarus -> Settings -> Output):

1. Non-Portable format – This stores the current configuration of the workspace along with settings, processed image results and polygons within a `'.scy'` file, and contains references to the locations of the original images on your filesystem.
2. Portable Format – This stores everything as above, plus the original images within the `'.scy'` file. This will result in a large workspace file and will take longer to create.

Workspace files are ordinary Zip files – you can open them with any Zip program (e.g. Windows zip utility) and inspect the contents. The following items may be contained within the file, depending on what was present in the Workspace when it was saved:

- A metadata file called 'scyven.dat'. This file is not human readable, and contains information about the Workspace state, Image list and processing settings for all the images;
- For each image that was processed, the components and processing results for that image (abundances, materials, shading, specularity map, Illuminant);
- If the Workspace is portable, copies of the original images.

4.2 The File Menu

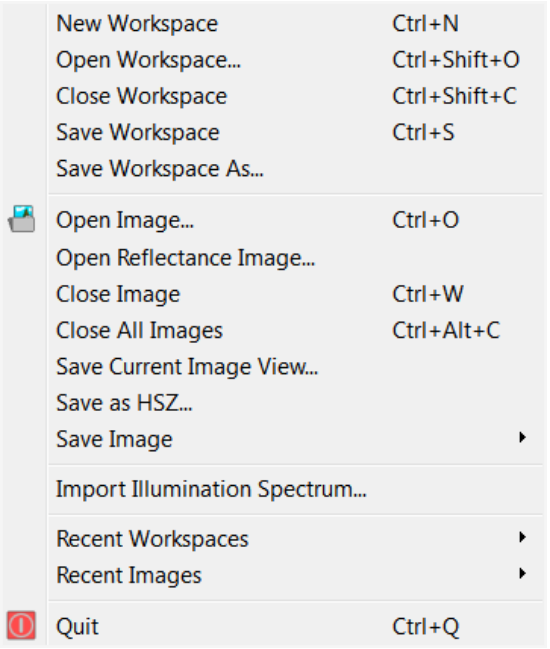

Figure 10: The File Menu

4.2.1 New Workspace

The New Workspace item will create a new workspace within Scyven. If there is already a workspace open, you will be prompted to save it before it is closed.

4.2.2 Open Workspace

The Open Workspace item will launch an open file dialog. Select the Workspace file you would like to open within this dialog and click Open to open it. Depending on the size of the Workspace this may take some time, and a progress bar may appear. Opening a workspace will prompt you to save your current workspace, if there is anything present in it.

4.2.3 Close Workspace

The Close Workspace item will close all images and results from the current Workspace, after prompting you if you would like to save first. After the workspace is closed, Scyven is reverted to its initial state (as if just opened).

4.2.4 Save Workspace

The Save Workspace item will, if the Workspace has been previously saved, update the Workspace save file on disk to reflect what is currently open in Scyven. If the Workspace has not yet been saved, it will behave as the 'Save Workspace As…' item.

4.2.5 Save Workspace As

The Save Workspace As item will launch a save file dialog. After choosing the folder you would like to save your Workspace to, enter a filename and click 'Save'. A Workspace file will be created (Workspace files are described in section 4.1.4).

4.2.6 Open Image

The Open Image item will launch an open file dialog. Select the Image file you would like to open (HDR/ENVI, TIFF or HSZ) within this dialog and click 'Open' to open it. Depending on the size of the Image this may take some time. Opening an image will add it to the current

Workspace. Images that have the 'REFLECTANCE' tag in the description field of the header will be treated as Reflectance images (see next section)

If you open an image with more than 100 million elements in it, the 'Opening Large Image' dialog will appear. Large images can consume a lot of memory. If when loading an image your computer runs out of memory, in most cases it will crash. You can use the 'Opening Large Image' dialog to select a subset of the image you are opening in order to avoid this (Or if you have enough RAM, just click 'OK' to continue opening the image).

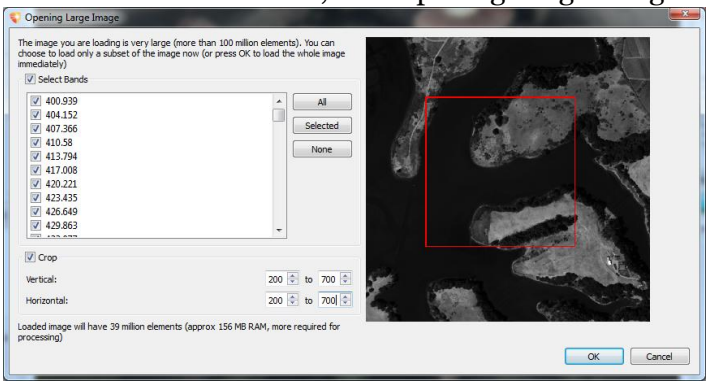

Figure 11: Opening a large image

4.2.7 Open Reflectance Image

The Open Reflectance Image item will launch an open file dialog. Select the Image file you would like to open (HDR/ENVI or TIFF) within this dialog and click 'Open' to open it. Depending on the size of the Image this may take some time. Opening an image will add it to the current Workspace.

Using this option will force Scyven to treat the image source as Reflectance rather than Irradiance. Note that all Reflectance images saved using Scyven will have a 'REFLECTANCE' tag within the header description field so opening the image with either image open option will have the same effect.

When an image is opened as reflectance, Scyven will still generate a 'colour' view of the image. Options that rely on the irradiance image will be disabled, though all processing that makes use of the reflectance image may be performed.

4.2.8 Close Image

The Close Image item closes the currently selected image, and removes it from the current Workspace.

4.2.9 Close All Images

The Close All Images item closes all open images, and removes them from the current Workspace.

4.2.10 Save Current Image View

The Save Current Image View item allows you to save the view of the Image Display Window, as formats JPG, PNG and BMP. The image that is saved is essentially a screenshot of the Image Display Window, i.e. all annotations (materials, polygons) will be flattened and saved on the image. This can be useful for saving visualizations, and individual results for viewing later, as well as saving colour reconstructions in formats that can be viewed by external programs.

4.2.11 Save As HSZ

The Save As HSZ item will launch a save file dialog. After choosing where you would like to save your file, and the filename, the HSZ file will be written to disk. See section 4.1.2 for more information about HSZ files. Note that you cannot save a HSZ file without processing the image first, due to the way the image is stored as a combination of components within the file. Furthermore, you cannot save a HSZ file if the image was opened as reflectance.

4.2.12 Save Image

The Save Image item allows you to save Irradiance or Reflectance in HDR/ENVI format. Choosing either option will open a save file dialog prompting for a save location and a filename. Clicking save will write the file to disk. The Irradiance option saves the original image data, depending on any resizes or crops that have been performed. This can be useful if you want to crop or resize images and then save them back in HDR format. The Reflectance option saves the 'raw' reflectance cube, and can be useful for further analysis in external programs. Note that the reflectance cube that is loaded from a HSZ file is compressed, and hence not the 'original' reflectance that is calculated when processing an image.

4.2.13 Import Illumination Spectrum

The illumination spectrum for the currently viewed image can be imported in the form of a CSV file or a text file where numbers are separated by whitepace. Note that the number of elements in the CSV or text file must match the number of bands in the image.

4.2.14 Recent Workspace / Images

Scyven will keep a history of recently opened Workspaces and Images for convenience.

4.2.15 Drag and Drop

Scyven supports drag and drop operations for opening irradiance Images and Workspaces. To perform this action, open Scyven and then drag an Image or Workspace from a file explorer window to anywhere within the Scyven window.

5 Images and the Image Menu

5.1 Hyperspectral Images

The data captured by a hyperspectral camera is typically stored as an 'image cube'. The image cube is named as such because the structure of the data is generally thought of as three dimensional. As in a traditional image, two of the dimensions are spatial, that is, they make up the pixels in the image. The third dimension is spectral – for each pixel there will be a number of wavelength samples. In a traditional colour image the depth of this dimension is three, (R, G, B). In a Hyperspectral image this depth is equivalent to the number of channels (e.g. 7 -100).

Scyven supports images of any size and any number of bands (well, at least 2, and no more than 2,147,483,647). HDR images saved in various formats are supported (all common formats are supported).

Hyperspectral images can be decomposed into the following components. Display toggles for each can be found in the Image Components Browser or the View menu:

- Irradiance – This is the original image as captured by a camera. The irradiance image can be viewed band by band
- Reflectance – The reflectance cube can be viewed band by band, with the brightness of each pixel representing the reflectance of the material at that location (brighter is higher reflectivity)
- Highlights / Specularity Map – This image is a single channel, and shows the areas of the image that contain specular highlights. These are the parts of the image that are highly saturated on all bands.
- Shading Map – The shading map shows the general shape of the scene, with lighting removed. You would expect a surface that was partially in shadow to have a uniform intensity in the shading map.

5.2 The Image Menu

The Image menu contains operations that are used to manipulate images before processing.

Figure 12: The Image Menu

5.2.1 Undo and Redo

The Undo menu option can be used to reverse changes made to images using Image operations. You can undo Crop, Resize and Remove band operations, one by one in the reverse order that you performed them. You can also redo operations (re-perform the last undone operation) using the Redo menu option or button. Note that if you perform the Undo operation, any processing that was performed previously will be cleared, as the image will be changed.

5.2.2 Reload Original Image

The Reload Original Image item will discard all processing results, and image manipulation operations (crop, resize) that have been performed on an image, and reload the original file from disk. This is essentially the same as closing the file and re-opening it, however as the image data file will be already open in memory, using this operation will be faster.

5.2.3 Remove Band

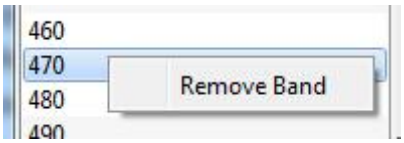

Figure 13: Remove a Band

The Remove Band menu option will remove the currently selected bands in the Image Components Browser. This option is grayed out unless the Irradiance or Reflectance buttons are selected in the Image Components Browser (otherwise no band can be selected for removal). You can also remove bands by right clicking on the selected bands in the Image Components Browser and selecting 'Remove Bands'.

5.2.4 Resize

The Resize item will launch the Resize Image dialog. To resize an image, enter the new size you would like the image to be in the width and height boxes. (To enter different scales for each dimension, click the small link in the middle of the dialog. Note that this will distort the image). For example, resizing an image to 50% will half its length and width, and therefore will contain 25% of the pixels of the original before resizing. An Image can only be resized to be smaller than the original.

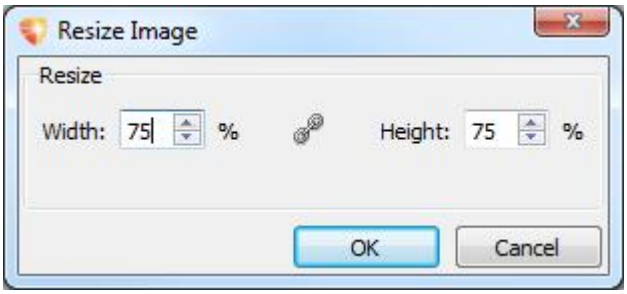

Figure 14: Resize Dialog

5.2.5 Crop to Selection

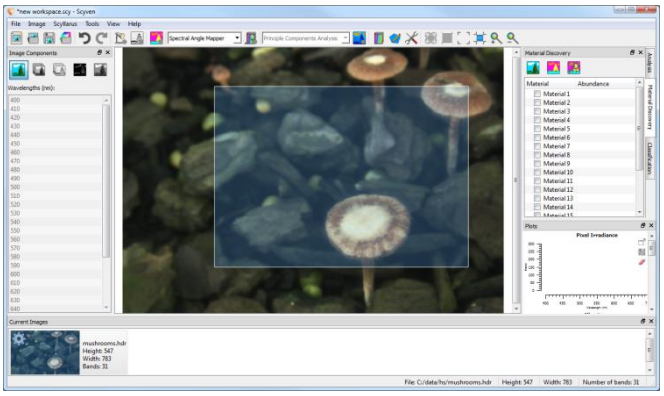

Figure 15: Area to be cropped to selected

The Crop to Selection item will crop the image to the area that is selected in the Image Display Window when it is clicked. To select an area on the image, drag (click and hold the left mouse button) from any point on the image (any of the four corners of the area you would like to crop to) to the diagonally opposite corner. Then release the mouse button, and select the 'Crop to Selection' item or button to crop the image. The selection box will disappear when you crop the image. When cropping, you can use the floating co-ordinate box near the mouse cursor to ensure desired dimensions.

5.2.6 Normalise Band

The Normalise Band item is a toggle option to choose how the image is to be normalised for display in the Image Display Window. Choosing to normalize by band will result in each band of the Irradiance and Reflectance to be normalized (scaled) by the maximum value in the band being viewed. Choosing normalize by cube will use the maximum value in the whole image instead. Choosing the latter is useful for comparing different bands in the same image, while choosing to normalize by band can be useful for viewing dark bands.

5.2.7 Image Rendering

The way the image is rendered on Scyven can be controlled by the Image Rendering item. Three tools are available:

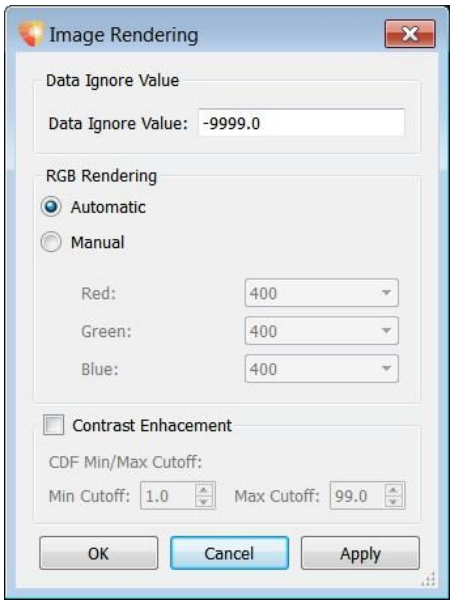

Figure 16: Image Rendering dialog

- Data Ignore Value: Pixel spectrums with all values set to the data ignore value will be rendered as 0 in the RGB image. Also, when undertaking contrast enhancement (see below), the Data Ignore Value is ignored during the contrast calculation.
- RGB Rendering: An image can be rendered in RGB automatically using a camera's color sensitivity function (default), or band values can be attributed to R, G and B manually.
- Contrast Enhancement: The contrast is enhanced by using histogram equalization. The minimum cutoff value is the percentage of the darkest pixels that are rendered as 0, and the maximum cutoff is the percentage of the brightest pixels that are rendered as 255.

6 Scyllarus and the Process Menu

Scyllarus is the technology behind Scyven that allows Scyven to process and analyze hyperspectral images. The Scyllarus image processing pipeline consists of a number of stages that act in turn to recover components and parameters of an image. The main stages of the processing are Illuminant Estimation, Reflectance Recovery, Material Discovery, classification and analysis

6.1 The Scyllarus Processing Pipeline

For more details on the Scyllarus processing pipeline, please see further information available at http://scyllarus.research.nicta.com.au .

6.1.1 Pre-processing

The pre-processing step can optionally apply spatial or spectral filters to the image to reduce noise. There are several filters available to choose from. (see section 6.2.1.2)

6.1.2 Illuminant Estimation

The NICTA illuminant recovery method is used to estimate the illumination power spectrum within the image. The illuminant function can be viewed in the spectral plots panel (6.2.5).

6.1.3 Reflectance and Component Recovery

The Component Recovery stage requires the illuminant as an input. This stage of the processing will calculate the reflectance cube, as well as the shading and specularity maps for the image.

6.1.4 Material Discovery

Material Clustering is performed on the reflectance cube. A deterministic annealing method is used by which several materials are found and clustered according to their reflectance spectra.

6.1.5 Material Classification

Material Classification is performed on the reflectance cube. Three methods for classification are currently available, Spectral Angle Mapping, Support Vector Machines and Linear Spectral Unmixing. Each of these methods takes an input of a Spectral Library and classifies each pixel in the image to one of the Spectral Library's classes.

6.1.6 Principal Component Analysis

Image analysis can be performed on the reflectance cube using Principal Component Analysis (PCA).

6.2 Process Menu

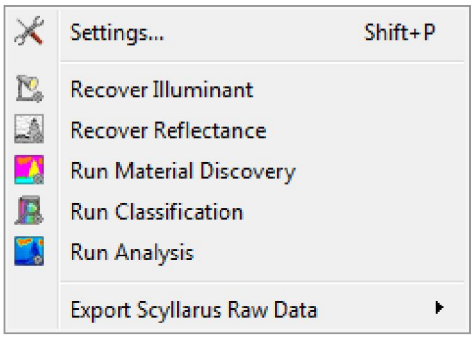

Figure 17: The Scyllarus Menu

6.2.1 Settings

There are several settings associated with the Scyllarus processing that can be accessed via the Settings Dialog. The Process settings are used to configure stages of the pipeline when you process an image. Clicking the Settings menu item will open up the Settings dialog.

6.2.1.1 Pipeline Stages and Filter Options

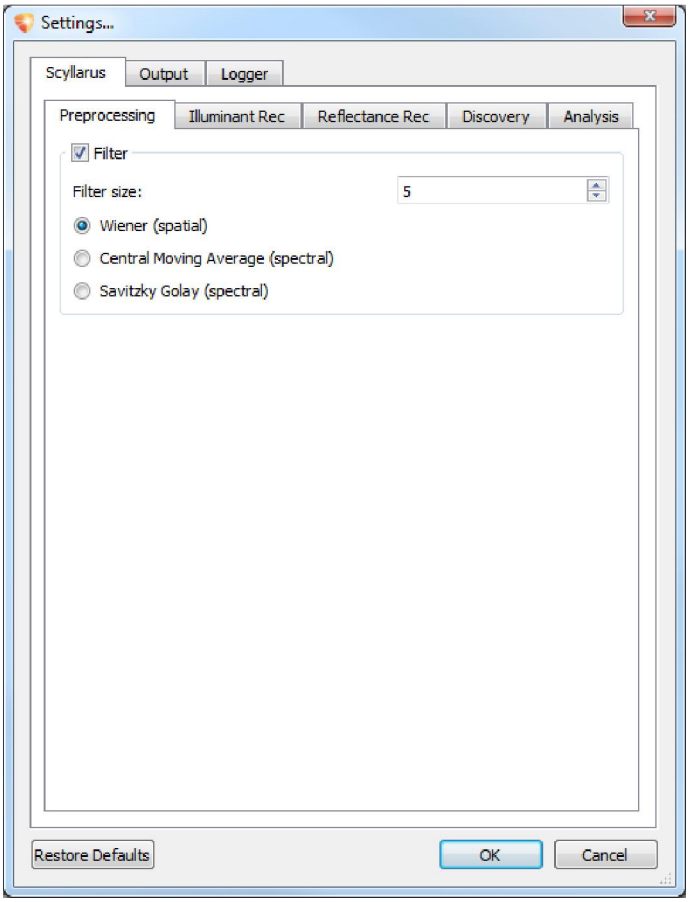

Figure 18: Process settings window

The Scyllarus Panel is used to control how the Scyllarus processing pipeline is executed upon images.

6.2.1.2 Advanced Pipeline Settings

Clicking the tabs at the top of the Scyllarus Settings will show the different Advanced Pipeline Settings (Figure 19: Pipeline Settings for Illuminant Estimation, Component Recovery and Material Clustering). These settings control how the Scyllarus processing will execute. You can adjust them to tune your results, but the default values are appropriate for most images.

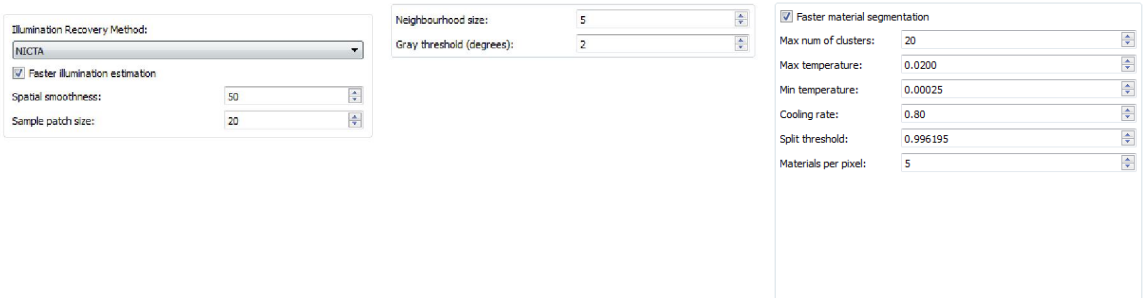

Figure 19: Pipeline Settings for Illuminant Estimation, Component Recovery and Material Clustering

The options are as follows:

Pre-processing panel (filters):

- 'Wiener' Filter (spatial) – this option can be used to filter out the noise in an image before processing. This can affect the quality and granularity of the results, so it is good to try processing your images with the option enabled or disabled to see which gives the best results for your case. The Wiener filter has a configurable window size, which must be an odd number 3 or greater.
- 'Savitzky-Golay' Filter (spectral) – this option can be used to spectrally smooth an image before processing. As above, this can affect the results so it is best to test and see if using this feature results in an improvement. The Savitzky-Golay filter has a configurable window size of 5, 7 or 9.
- 'Central Moving Average' Filter (spectral) – this option can be used to spectrally smooth an image before processing. The Central Moving Average filter has a configurable window size which must be an off number 3 or greater, and is better suited to smoothing images with large numbers of bands.

Illumination Estimation Panel:

- Illumination Recovery Method – Several different options are available for this. The default option is suitable for most images.
- Faster Illumination Estimation – For images with more than 50 bands, if this option is on (checked) an optimisation that significantly improves execution speed will be employed. There is a small (less than 1%) difference in the result when using this method.

- Spatial Smoothness – Specifies the maximum number of locations to sample when determining the Illuminant.
- Sample Patch Size – Specifies the size of the areas sampled when determining the Illuminant.

Component Reconstruction Panel

- Neighbourhood size – The size of the sliding window used to calculate the components.
- Gray Threshold – Threshold value used when determining specularities.

Materials Clustering Panel

- Faster Material Segmentation – Use a quick method to get an estimate of the material segmentation.
- Max Number of Clusters – The maximum number of materials that will be classified by the clustering algorithm. For example, when set to 20, there will be 20 or less clustered regions in the result.
- Max / Min temperatures and Cooling Rate – The maximum and minimum temperatures define the starting and ending points of the annealing algorithm, with the cooling rate being the step size of each iteration. For example, with the default settings the Temperature (starting at the max = 0.02) will be cooled by the cooling factor (0.8) every iteration until the temperature is below the minimum (0.00025).
- Split Threshold – The split threshold is the threshold used to determine cluster similarity. If set to 1, any two compared clusters, even if identical, will be deemed different, while a lower threshold will be more generous (0.8, for example would result in very few splits being made).
- Materials Per Pixel – When performing the final stage of the material clustering algorithm, each pixel in the image is broken down into a linear combination of several of the identified materials in the image. When using the default (5) this will result in each pixel in the image being represented as some combination of 5 of the recovered materials (e.g., 70% Material 3, 25% Material 17, 4% Material 10, 0.6% Material 12, 0.4% Material 6).

Analysis Panel

- Number of Principle Components – The number of principal components to show in the Analysis panel (default 10)
- Apply to - Choose whether to perform the PCA on the reflectance image (default) or the irradiance image.

6.2.1.3 Output Options

The Output Options are used to configure how certain files are saved by Scyven.

The Make Workspace Portable option is used to toggle the way Workspace files are saved. When checked, all Workspace files will contain all Image source files in addition to their regular contents. Please see section 4.1.4 for more information on Workspaces.

The HDR/ENVI Image format options are used to choose the data type that is used when saving HDR/ENVI files. You can choose from 8 bit unsigned, 16 bit unsigned, and 32 bit floating point data types.

The HSZ Image Panel contains saving options for HSZ files. You can choose to save HSZ files using NURBS Interpolation for the Material Spectra and Illumination. The settings for NURBS are:

- Polynomial Degree – The degree of the polynomial curve function used in the NURBS encoding;
- Max Number of Iterations – The number of iterations used in the NURBS curve fitting algorithm;
- Knots Threshold – The max number of sample points on the curve;
- Alpha – The Alpha value used in the NURBS encoding.

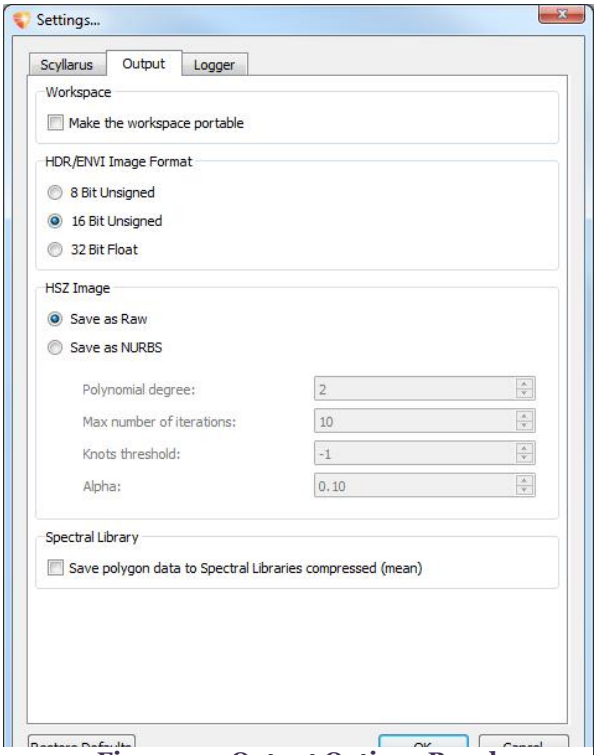

Figure 20: Output Options Panel

The Spectral Library option allows you to choose whether spectral libraries are saved in a compressed format. Normally, every pixel within a label is saved as an entry in the Spectral Library (for use in classification algorithms such as Support Vector Machines) but each label can be aggregated to a mean value (meaning one entry is stored per polygon) using this option.

6.2.1.4 Logger Options

The Logger Options Panel contains a slider that controls how much diagnostic information is output to the Processing Logs window.

6.2.2 Run Processing

Clicking the following buttons will run the Scyllarus Processing:

- Recover Illuminant
- Recover Reflectance (will run Recover Illumination as well)
- Run Material Discovery (will run previous two stages).
- Run Classification (will run Recover Illuminant and Recover Reflectance as well)
- Run Analysis (will run Recover Illuminant and Recover Reflectance as well)

Each item or button will begin the execution of the Scyllarus processing pipeline on the selected Image. This process will take less than a minute for most images. Once the processing is complete, the processed results can be viewed.

The processed results include the Illumination Power (displayed in the Plot panel), the Reflectance Cube, Shading Map and Specularity Map (displayed in the Image Components Browser) and the Material Discovery Results (displayed in the Material Discovery Browser)

6.2.3 Export

The Export item allows you to export Illumination, and Material Reflectances as CSV files. Selecting either option will open a save file dialog prompting you to choose a location and filename, clicking save will write the data to disk.

6.2.4 Processing Logs

The Processing Log Output window (accessed from the View menu) is used to view the Scyllarus processing pipeline diagnostic logs. Depending on the log-level you will see varying amounts of information. If you are experiencing issues processing your image, it might be useful to look at the log outputs, as they might give you an indication as to what settings you will need to adjust to get better results. Logs can be saved to a text file. If you need to submit a bug report, please send the processing log associated with the bug too.

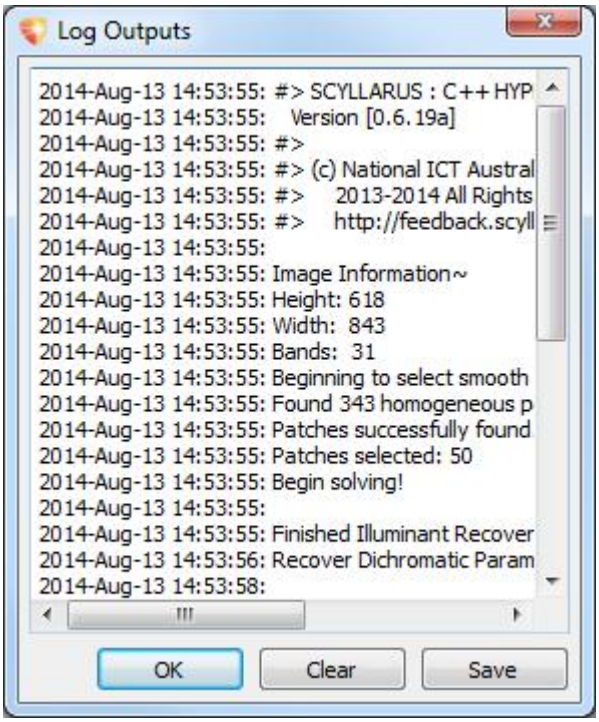

Figure 21: Log output

6.2.5 Spectral Plots

The Spectral Plots allow you to inspect pixel-wise information about the image you're viewing. You can view them in the Spectral Plots panel, and each plot can be popped out for more detail. The four plots are:

- Illumination – If you have performed Illumination Estimation, or loaded a file that has been previously processed, this will show the Illumination function. The function is interpreted as each data point describing the Illumination power for that particular band.
- Pixel Irradiance – This plot shows the raw values of the original loaded image.
- Material Reflectance – This plot shows the Spectral Profiles of each of the currently selected materials from the Image Materials Browser.
- Pixel Reflectance – This plot shows the reflectance values from the Reflectance Image.

The two plots, Illumination and Material Reflectance, show items that are static (to manipulate the Material Reflectance, use the Material Discovery Browser). The other two plots show data that you choose. To add data to the plots, you can click or select areas within the Image Display Window while the plot window is open. Clicking will add one pixel to the plot, while dragging will add a number of thinner lines representing the region selected to the plot.

- You can zoom in on the image when selecting pixels or regions to plot to more easily plot the data you are interested in.
- The plot window can be used to visually compare the spectra of two or more parts of the image, as well as compare sections of the image to the recovered material spectra.

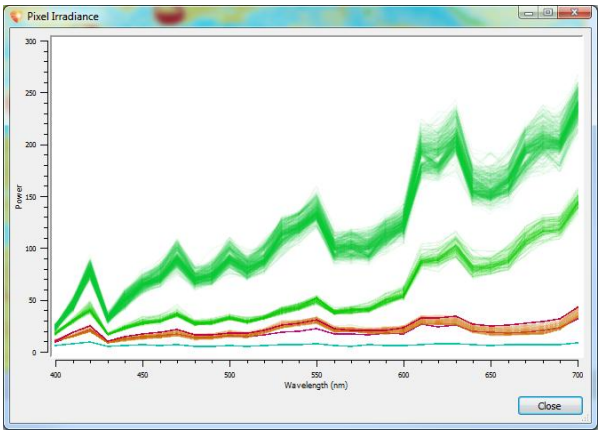

Figure 23: Spectral Plots Pop-out

6.2.6 Live plotting

The live plotting feature allows you to interactively view spectral signatures within an image. To enable it, find the option in the View menu. Once enabled, whenever you move your mouse across the image view a live plot will appear on the available plots within the plot window and on any pop-out versions. The live plot will update instantaneously as you move your mouse and show the spectral profile of the pixel that lies under the current mouse position.

7 Classification and Analysis

Classification and Analysis can be performed easily on Images within Scyven using the following Tools:

- Spectral Angle Mapping
- Support Vector Machines
- Linear Spectral Unmixing
- Principal Component Analysis

Each analysis tool has different features and merits, and can be used to reveal a wide range of information within Hyperspectral Image Data.

7.1 Spectral Angle Mapping

Spectral Angle Mapping (SAM) is a Classification technique. SAM will classify pixels in an image to a given Spectral Library by calculating the Spectral Angles between each pixel and spectra, and assigns the closest match as the pixel's class.

To perform Spectral Angle Mapping classification, you will need a Spectral Library that is appropriate for the image you are trying to classify. The algorithm cannot determine whether a match is 'good' or not, it will merely provide the best match for each pixel it classifies. Therefore, it is fundamental to getting good results that the library you use be appropriate for the image.

To see how to load a Spectral Library into Scyven, see section 8.1. Once you have loaded a Spectral Library, you can run the Spectral Angle Mapping Classification by selecting 'Spectral Angle Mapper from the pull-down in the menu bar, and then pressing the 'Run Classification' button to the right of the pull-down.

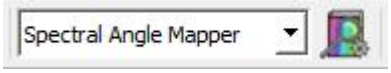

Figure 24: Classification Tools

The results will be shown in the Classifcation Panel (see section 2.2.6). When viewing the results, you can toggle the transparency of the classification layer. Each pixel is classified with it's closest matching spectra from the given library. The more transparent a pixel is, the less closely it matches with the class it has been labeled with (however this is the closest match among the given spectra that could be found). Conversely, a solid colour indicates a near 100% match.

7.2 Linear Spectral Unmixing

Linear Spectral Unmixing is a classification technique. It performs pixel-wise classification using a Spectral Library by solving to compose each pixel as a linear combination of some of the provided Spectra in the Spectral Library.

The comments from the Spectral Angle Mapping section regarding the appropriateness of the Spectral Library used still apply to this classification method. To run Linear Spectral Unmixing, Load a Spectral Library into Scyven (see section 8.1) and select 'Linear Spectral Unmixing' from the pull-down in the menu bar, and then press the 'Run Classification' button to the right of the pull-down.

The results will be shown in the Classifcation Panel (see section 2.2.6). When viewing the results, you can toggle the transparency of the classification layer. Each pixel is classified with it's closest matching spectra from the given library. The more transparent a pixel is, the less closely it matches with the class it has been labelled with (however this is the closest match among the given spectra that could be found). Conversely, a solid colour indicates a near 100% match.

7.3 Support Vector Machines

Support Vector Machines (SVM) is a classification technique. It performs pixel-wise classification using a Spectral Library. Unlike Spectral Angle Mapping and Linear Spectral Unmixing, SVM can take advantage of multiple data points per class, so using 'raw' Spectral Libraries (where polygons are not saved as mean but as all values) will enhance classification results. SVM will build a classification model based on the input data, and then classify each pixel against the input spectra to give a class label for each.

As with the last two methods, the Library used must be appropriate to the image to get meaningful results. To run Support Vector Machines, Load a Spectral Library into Scyven (see section 8.1) and select 'Support Vector Machines' from the pull-down in the menu bar, and then press the 'Run Classification' button to the right of the pull-down.

The results will be shown in the Classifcation Panel (see section 2.2.6). Each pixel is classified with it's closest matching spectra from the given library. SVM only gives a binary result (it assigns a class, with no probability/certainty) so unlike SAM and LSU, there is no transparency option.

7.4 Principal Component Analysis

Principal Component Analysis (PCA) is an unsupervised (no additional input) analysis technique that can be used to highlight areas of high variance within an image. The PCA algorithm will process the Reflectance Cube (default) and present the first few Principal Components recovered. The Components are sorted by their 'quality' so the most useful results should be found near the top of the list.

To perform Principal Component Analysis on an image, click the 'Run Analysis' button.

After the PCA algorithm has completed, you can view the results in the Analysis Browser (see section 2.2.7). The slider can be used to change the scale of the results to better pick out the features found. The following pictures illustrate this.

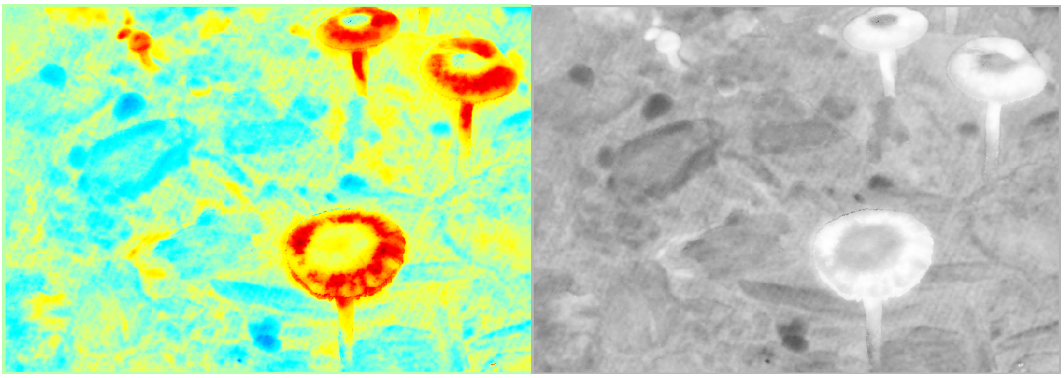

Figure 25: PC 2 for 'mushrooms.hdr' shown in colour and grayscale.

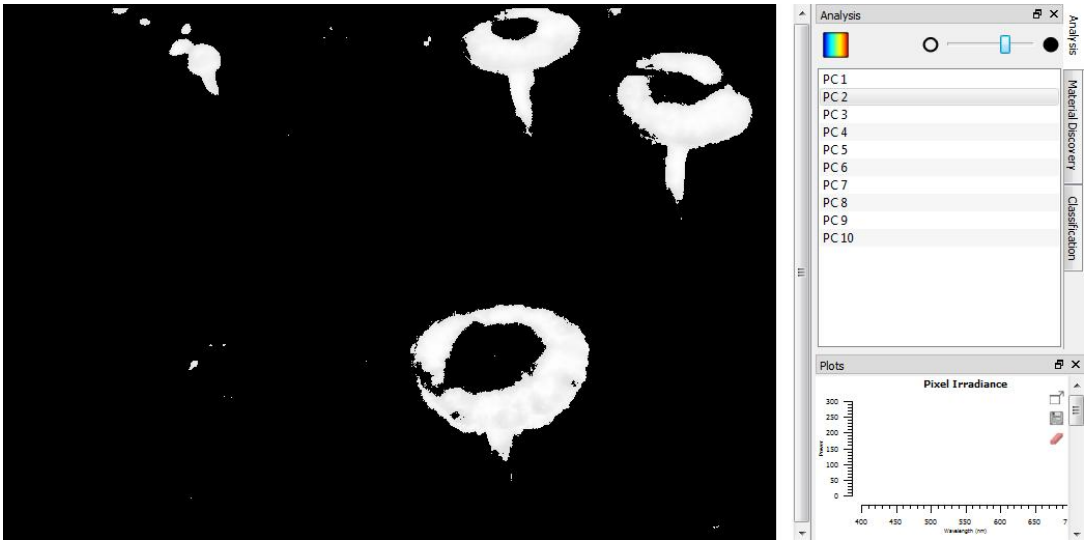

Figure 26: PC 2 of the image 'mushrooms.hdr' with the scale slider set as shown.

Figure 26: PC 2 of the image 'mushrooms.hdr' with the scale slider set as shown. Illustrates how the Slider can be useful to isolate the features that have been found with PCA.

8 Tools Menu

8.1 Spectral Library Manager

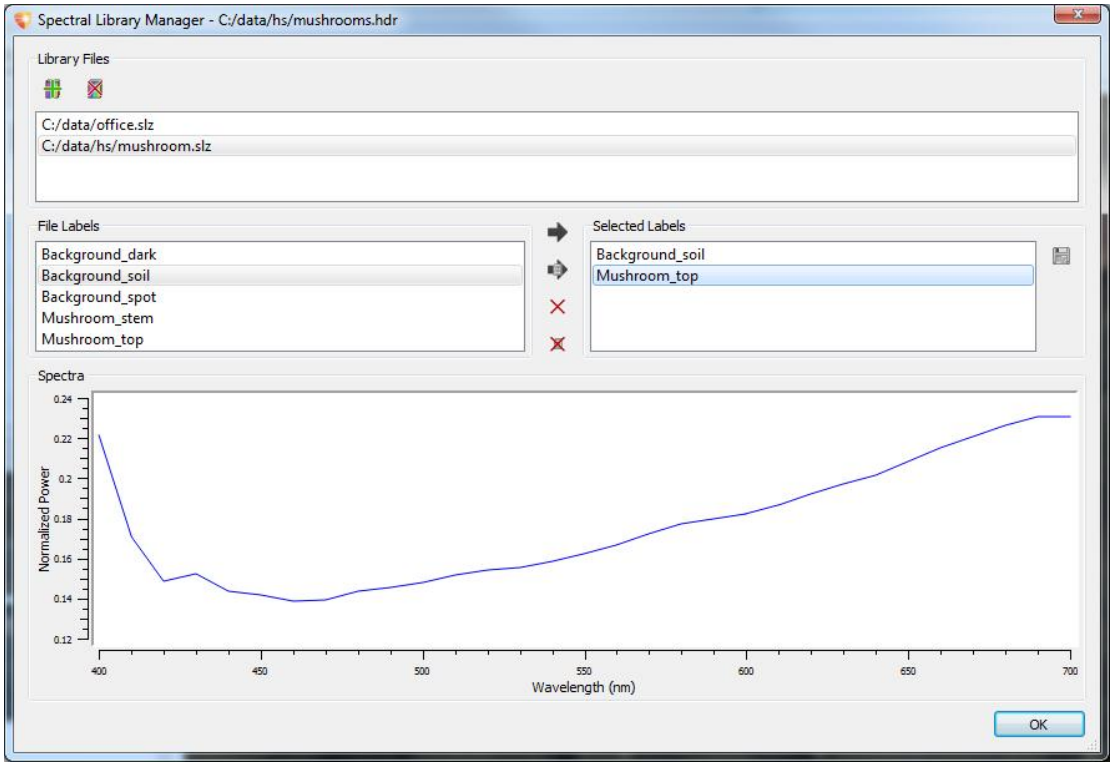

Figure 27: Spectral Library Manager

The Spectral Library Manager is used to load Spectral Libraries into Scyven for use with the following features:

- End Member Indexing
- Spectral Angle Mapping
- Linear Spectral Unmixing
- Support Vector Machines

SLZ files contain lists of 'endmembers' or 'spectral signatures'. Each has a label and spectra associated with it. There are two ways that Spectral Signatures are stored – 'Raw', where there are many signatures for each label, and 'Mean', where there is just one signature for each label. For classification the Spectral Library can be considered as a list of classes, with each class having a label, and a centroid (the spectral profile associated with that label).

The Spectral Library manager allows you to load one or more Spectral Libraries into Scyven, and view the contained labels and corresponding data. Each Image that is loaded can be

assigned a subset of spectra that will be used for classification on that image. Once you click OK, the Spectral Library Manager sets the Spectral Library that will subsequently be used for classification on the currently selected image.

To load a Spectral Library, click the green + button. This will open a file dialog. To remove a Spectral Library, select it from the list and press the red -.

Once a library is loaded, you can view the contained data by selecting the library from the list and browsing through the contained spectra in the 'File Labels' list.

The Selected Labels list contains the labels that have been 'attached' to the current image. The four buttons in the middle of the two lists allow you to move items into the Selected Labels list to add them to the current image. The Copy One button (top) moves the currently selected spectra from the File Labels to the Selected Labels pane. The Copy All button (second from top) will move all the spectra in the current file to the image list. There are also buttons to delete one and delete all.

You can save the 'Selected Labels' list as a new SLZ file by pressing the save button to the right of the pane.

Pressing OK will finalise the changes and attach the list of spectra to the current image for use in classification.

8.2 Pixel Inspector

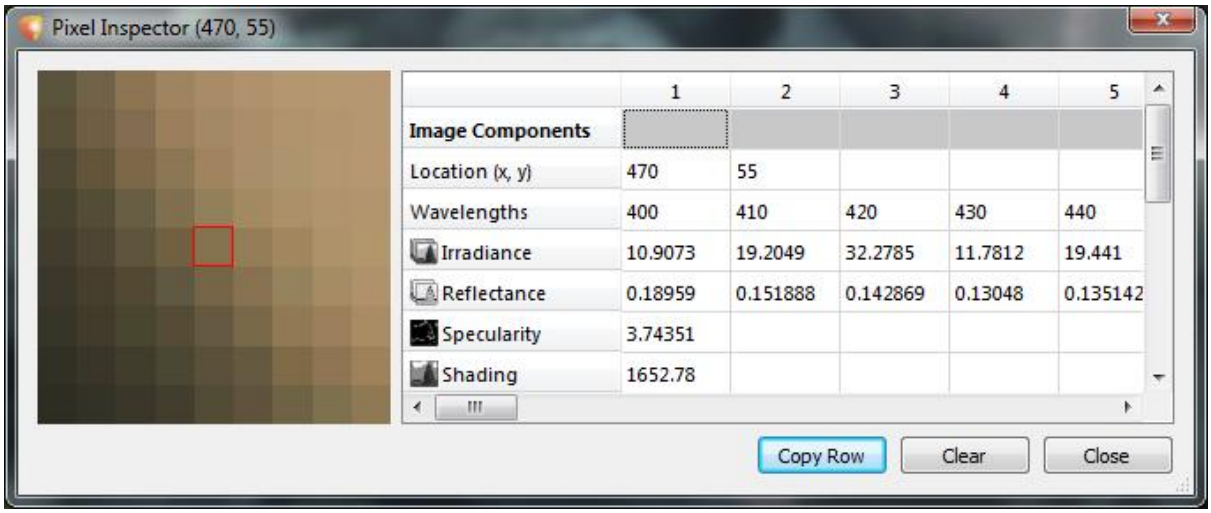

Figure 28: Pixel Inspector

The Pixel Inspector is used to see the numerical data that is associated with each pixel in an image. There are several sets of data available to view with the Pixel Inspector, with certain

rows only being available after the corresponding processing has taken place (Wavelengths are the same for each pixel and given as a reference):

- Irradiance values – the irradiance for the selected pixel at each band (not normalised)
- Reflectance values – the reflectance for the selected pixel at each band (normalised)
- Specularity – the specularity value for the selected pixel (not normalised)
- Shading – the shading value for the selected pixel (not normalised)
- Material Discovery 'Material' – the labels for the clusters associated with this pixel
- Material Discovery 'Abundance' – the probability for the above label at this pixel
- Classification 'Class' – the classes associated with this pixel
- Classification 'Abundance/Angle' – the probability (LSU) or angle (SAM) for the above label at this pixel
- PCA: Value – the values for the principal components calculated at this pixel.

To display data about a pixel, open the Pixel Inspector, and then click anywhere in the Image Display Window to select a pixel to display. You will also see a highly zoomed image centred about the pixel selected.

You can copy data from the Pixel Inspector to other applications via the Operating System clipboard. Select the row you would like to copy and press the 'Copy Row' button to copy it to the clipboard. The data is copied as tab separated values (suitable for pasting into Excel, Matlab, etc.)

8.3 Polygon Labels

Polygon Labels are used to select certain areas of Images in order to make annotations and export data from those regions.

8.3.1 Create Polygon

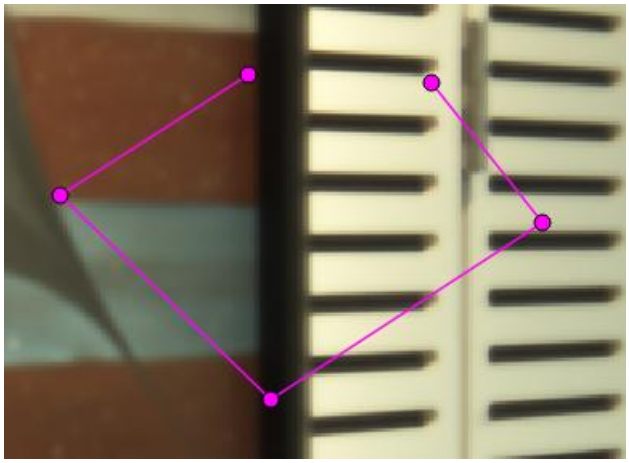

Figure 29: Drawing a polygon

To create a polygon, select the Create Polygon item from the tools menu or press the create polygon button. A polygon is drawn by clicking on the outer corners of the shape you would like to create in series. Once the polygon shape is complete, clicking again on the first point will close off the polygon and complete the drawing. Or alternatively, you can double click on the last point to close the polygon. A dialog will appear asking you to name the new polygon, or to add it to an existing group of polygons.

8.3.2 Export Polygon Data

There are several options for exporting Polygon Data. The 'Pixel Irradiance' menu item and the 'Pixel Reflectance' menu item output raw pixel data. Choosing either one will launch a file save dialog asking for a location and filename. The format used is CSV, with each polygon having an image file name, a label name, and the (X, Y) co-ordinates of each pixel contained within the polygon and the corresponding value for that point.

The Spectral Library option will create a new Spectral Library (SLZ) file from the polygon data. This can be then used to perform Classification by loading it into Scyven with the Spectral Library Manager.

8.3.3 Add Polygons to Spectral Library Manager

Once polygons are created, they can be added to the Spectral Library Manager for classification or be selectively stored in SLZ files. To add the currently selected Polygons from the Polygon Browser to the Spectral Library Manager, either select the Tools->Add Selected Polygons To Spectral Library option, or press the corresponding button in the Polygon Browser. Once added, Polygons will show up in the Spectral Library Manager under the 'Polygons' heading.

9 View Menu

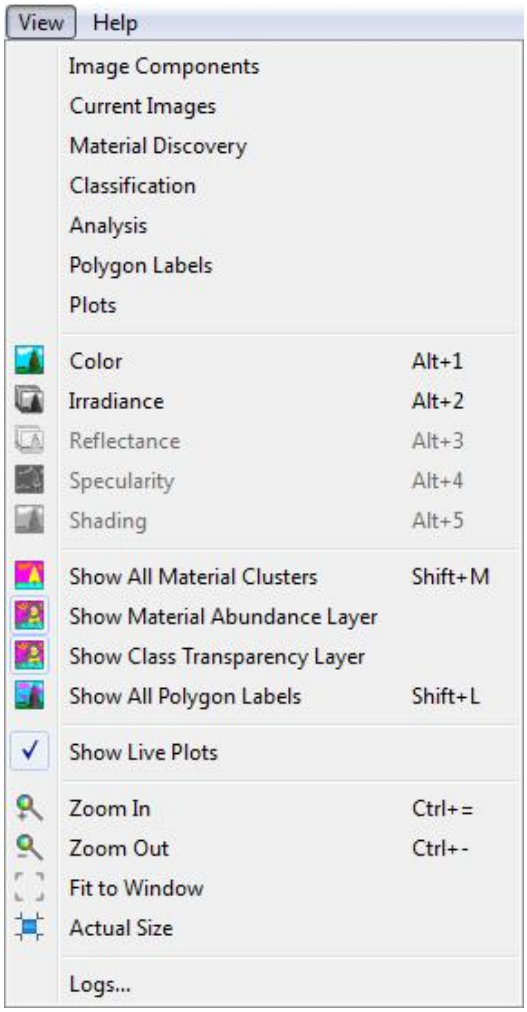

Figure 30: The View Menu

The View menu contains items that configure the Scyven Window.

The top panel contains items that will re-show the browsers. The second panel contains Duplicates of the buttons in the Image Components Browser, as well as a list of quick keyboard shortcuts.

The third panel contains controls for toggling Materials and Polygons, and the final panel contains the Image Display Window zoom controls.

10 Bug Reporting and Feedback

If you find a bug in the software, or would like to send feedback / comments, please use the form at http://feedback.scyven.com