


HyperCube



HyperCube Pictorial User's Guide

Contents

About HyperCube	6
Software	6
Loading Images	7
Files, Data, Images and Display	7
Intrinsic File Types	8
General File Types	8
SRF (Sun Raster Files) and TARGA	9
ERDAS Imagine®	10
ESRI Shape Files	12
HDF Files	13
JPEG Files	14
LASF/ASCII LIDAR Files	15
LASF Point Cloud Viewer	19
NITF Files	24
PNG Files	25
TIFF Files	26
SDF Files	27
Raw	35
Multiband	35
Image Format Dialogs	36
Overviews	40
Saving Overviews	41
Header Files	42
Keyword Format	43
Image to Reference Coordinates	45
Geo/UTM Reference Grids	46
Create Geographic Rational Polynomial Coefficients	47
Geographic Coordinates <-----> UTM Coordinates	47
Rational Polynomial Coefficients (RPC) Files	48
Wavelength File	49
Statistics	50
Image Operations	52
Zoom	52
Selection Rectangle	52
Copy, Paste & Duplicate	52
Gray Mapping	53
Color Mapping	55
Data Mapping	56

Flips & Rotations	58
Shifts	59
Filtering (Spatial Domain)	60
Magnify	62
Info Window	63
Image Conversions	64
Components to Color	64
Image to Various Components	64
Image to Principal Components	65
Indexed to Color	66
Color to Indexed	66
Gray Scale/Color Tables	66
Replace Color	67
Image to Mask	68
Histogram Equalization	69
Cube Color Linked Image	71
Color to Classmap	71
Attach/Remove Data	71
Image Functions	72
Statistical Measures	72
Histograms	76
Pixel Exclusion	78
Scatter Plot	79
Line Profile Plot	80
Image Arithmetic	81
Overlaying Images	90
Adjacency	91
Fourier Operations	95
Mosaic Points	97
Mosaic References	100
Warp	102
Plot Scaling/Selection	104
Multiband Functions	105
Band List	105
Annotation List	107
Band Scrolling	108
Flicker/Superimpose	109
Constructing/Modifying a Cube	111
Color Composite	113
Spectral Plots	114
Selection Point Colors	116
Spectral Libraries	117
Format	117
Plotting Signatures	119
Superimposing Signatures	120

Spectral Classifications	121
General Algorithms	121
Specific Algorithms	122
Classify Function	126
Classify Options	129
More Options	133
Image Products	135
F(Signatures)	138
ROC Curves	140
Library Searches	143
Dynamic Displays and Spectral Plots	145
Attach Class Map	145
Class Map Editor	147
Spectral Calibration	150
Saving Images	153
Printing	155
Utilities	156
Reformatting a File	156
Wrap JPEG in NITF	158
GeoTIFF, JPEG to NITF	159
Buckeye Files to NITF	159
Converting a Library to a Cube File	160
ASCII File to Library	161
Converting a Library to an ASCII File	162
Concatenate Images	163
Concatenate Cubes	163
Concatenate Files	164
Changing a Window Title	165
Apply Transformation	166
Merge LAS Files	167
ASCII LIDAR to LASF	167
Applications	168
Shaded Relief	168
Contour	169
Stereo Mate	170
Stereo Compilation	171
Perspective	179
Radiance	182
Lens Distortion	185
Miscellaneous	188
Icons	188
Error Messages	188
Capturing the Screen and Listing Pixels	189
Saving and Copying Scrollable Lists	189
Loading/Saving Text Files	189

HyperCube Menus	190
File	191
Edit	192
Applications	194
Image	195
Functions	197
Utilities.....	198
Windows.....	198
Index	199

About HyperCube

HyperCube is a Macintosh and Windows application program specifically directed to the analysis and display of multi and hyperspectral imagery. This includes the static and dynamic display of the image cube and the generation of spectral classifications using both imagery and spectral libraries. In addition, **HyperCube** contains functions to filter, warp, mosaic, reformat, calibrate, combine, photogrammetrically project and to perform arithmetic on imagery and data.

The current version number of **HyperCube** is shown in the lower right hand corner of the splash screen (on page 1 of this document) by selecting **About...** in the Apple menu or the **Help** menu in the Windows version.

There are numerous example data sets available at the HyperCube website. The usage of each is described in the Quick_Start.pdf document.

Software

There are currently two Mac (Macintosh) versions of **HyperCube**. One version is for the older Power PC processor and the other is for the current Intel based architecture. There are also two Windows versions: a 32 bit and a 64 bit executable. They are both Windows® 95 through Windows 7 compatible. Note: The Mac Power PC version is no longer maintained with updates and remains at version 10.80.

The Mac and Windows programs are 99.9% identical with two differences: the Windows version does not support the Mac **Pict** image format and the Mac version does not have [SDF](#) (Sample Data Files) full wave form Lidar processing. The later requires use of dlls. Although this documentation was prepared on a Mac version of Adobe InDesign® using Mac OS X examples, the Windows GUI matches it one-for-one. Any slight (very slight) difference between the two systems is delineated in the text (in particular see the “Load Selection” button in section [Band List](#)). Also, see “Windows version” in the index.

The Windows version uses the **Control (Ctrl)** key instead of the Macintosh special command key. So, for example, to perform a [paste](#) operation press the **Ctrl** and **V** keys simultaneously. The Mac **Option** key is replaced with the **Alt** key in Windows.

Loading Images

Files, Data, Images and Display

At the present time almost every window that is displayed on a computer monitor consists of 8 bit gray or 24 bit color (or possibly a few bits more). **HyperCube** can handle data up to 32 bits per pixel with a spatial extent up to at least 20,000 by 20,000 pixels. How data is converted from a source file to a displayed image window is described below.

Pixel depth:

8 bit - display directly as gray or pseudo color.

24 bit color - display as true color.

16 bit signed or unsigned - keep 16 bits in memory as “attached data” and remap by various methods to 8 bits for display (the image load dialogs contain options for remapping),
See section [Data Mapping](#) for other ways to remap after loading).

32 bit float or long - keep full depth in memory and convert for display.

64 bit and higher - not presently supported in HyperCube.

Spatial size:

Image sizes greater than 2048 pixels or 2048 lines can be loaded as a scaled overview image (see **Related functions** below) and used to outline areas that are loaded at full resolution. Image subsets may also be loaded.

Multiband Images:

Each band that is loaded into memory includes the attached source data for that band. However, the number of bands loaded into memory does not affect spectral classifications including most statistical operations. Therefore, it is not necessary to load an entire multiband image in order to exploit it (see sections [Header Files](#) and [Band List](#)).


Related functions:

[HDF](#), [JPEG](#), [LASF](#), [SDF](#), [NITF](#), [PNG](#), [TIFF](#), [Raw](#) and [Multiband](#) images,
[Attach Data](#),
[Overviews](#),
[Gray Mapping](#), [Color Mapping](#), [Data Mapping](#),
[Statistical Measures](#),
[Image Arithmetic](#),
[Classify Function](#).

Note, to see the underlying data of a display pixel choose menu **Windows -> Show Info** which in addition to listing the gray or color display pixel will also show any data value. [Duplicating](#) an image also makes a copy of the attached data. Some operations, such as menu Applications -> [Shaded Relief](#), also copy the source data and attached it to the newly constructed image.

Intrinsic File Types


The intrinsic image file format for the Mac and PC is **TIFF**. HyperCube also automatically recognizes Pict, **TIFF**, **JPEG**, **PNG** and .SUN (SRF) file types in the menu File->Open... (shown at the right) provided their Finder file types (not the file name suffix) are properly set. If the list does not show the file you are expecting then the image must be opened using menu File -> Open As... with you specifying the file type. If Edit -> Options -> Detailed Messages is checked then the header contents of the various file types will also be listed and is a good diagnostic tool.




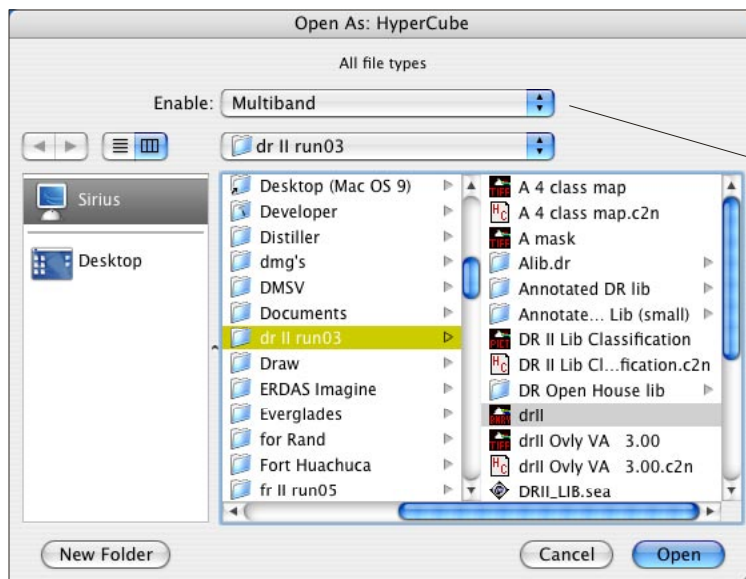
The Windows version of HyperCube uses File -> Open... (shown below in General File Types) to load all types of image formats except Pict which is not supported. TIFF is the general file type for Windows. Use the Files of type: popup menu to select which type.

General File Types

Any image supported by HyperCube can be opened by Open As... in the Mac version and just Open... in the Windows version as shown by the dialogs below and the following page.

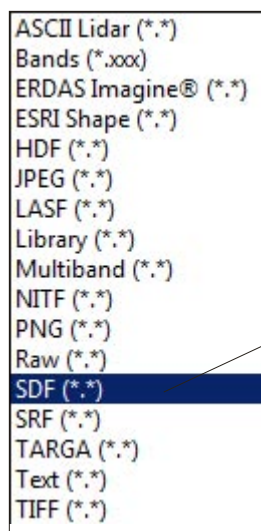


Windows standard Open... file dialog.



Currently supported file types.


Macintosh OS X standard Open As... file dialog.



Note: The Window version of the file-type popup menu, shown here, also includes the [SDF file type](#).

SRF (Sun Raster Files) and TARGA

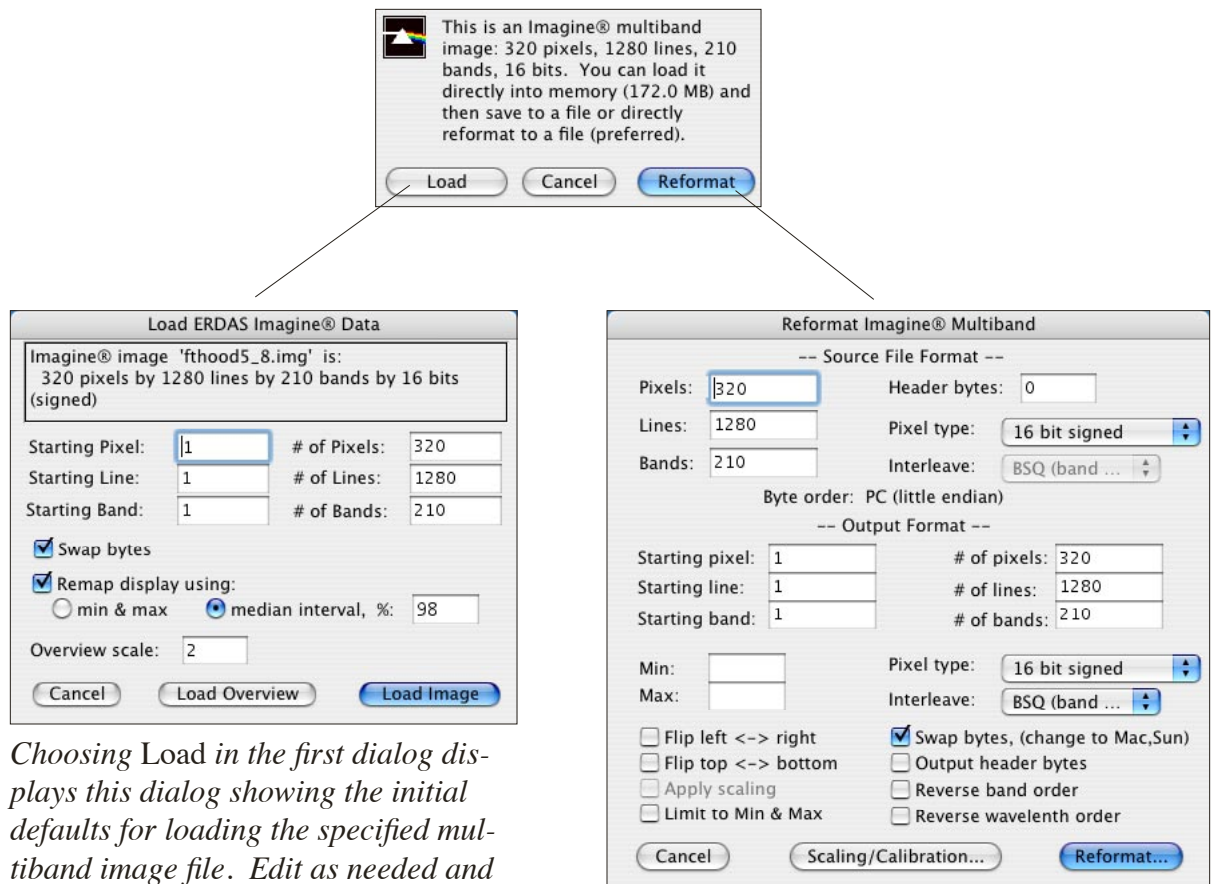
Sun raster (SRF) and TARGA formatted images can be loaded by selecting **Open As...** from the **File** menu and then choosing the appropriate item in the **Format** popup menu. No other dialog is required as these images contain sizing information in their header. HyperCube should be able to read 8 bit and 24 bit compressed and uncompressed Sun images. HyperCube can also save images in the Sun format but not TARGA.



ERDAS Imagine®

Certain types of ERDAS Imagine® *.img files can be read (imported in some cases) into HyperCube by selecting menu **File -> Open As...** and choosing **ERDAS Imagine®** as the file type. The supported pixels depths are: 8 bit unsigned, 16 bit unsigned, 16 bit signed, 32 bit signed and 32 bit float gray scale single and multiband images. Also, 8 bit pseudo and 24 bit true color images can be loaded. Once loaded the images may be saved in any format including Imagine®.

An Imagine® multiband image should be reformatted into the HyperCube multiband format with corresponding *.hdr and *.wvl files. Loading a multiband image brings up the following dialog immediately after the **Open As** selection. Choosing either **Load** or **Reformat** displays the two subsequent dialog paths. Spectral classifications can not be directly performed until an Imagine® multiband image has been converted to HyperCube multiband.



Choosing Load in the first dialog displays this dialog showing the initial defaults for loading the specified multiband image file. Edit as needed and select Load Overview or Load Image. The Mac version of HyperCube will have the Swap bytes check box on while in the Windows version it will be off.

Choosing Reformat in the first dialog above results in HyperCubes's general Utilities -> Reformat File dialog. See section [Reformatting a File](#).

If **Load Overview** is selected in the **Load** dialog then a scaled overview of the selected band is loaded. Any portion of the overview image can be designated and saved, including automatically reformatting the subset (see section [Saving Overviews](#)).

The Hierarchical File Architecture (HFA) of a given ERDAS Imagine® image file can be displayed by first toggling on menu Edit -> Options -> Detailed Messages and then loading the file. It is quite extensive and is displayed in several text windows containing the HFA nodes. Shown below are the results after loading a single 16 bit gray scale image. Note that a 210 multiband image has over 2000 nodes.

```

dma_dted.img HFA

Sirius:ERDAS Imagine:dma_dted.img

headerPtr: 64
version: 1
freeList: 1713
rootEntryPtr: 82
entryHeaderLength: 128
dictionaryPtr: 530373

Number of Ehfa entries: 12
Number of Edms_State entries: 1

image width: 512
image height: 512
layer enum: thematic
pixel enum: u16
tile pixels: 64
tile lines: 64
imageBands: 1
rgb order: 0,0,0

Ehfa_Entry 1 (82)
next node: 0
previous node: 0
parent node: 0
child node: 2632
data ptr: 0
data size: 0
name: root
type: root
mod time: 1337982976

Ehfa_Entry 2 (2632)
next node: 210
previous node: 0
parent node: 82
child node: 0
data ptr: 2760
data size: 4
name: IMGFormatInfo
type: IMGFormatInfo831
mod time: 1337982976

```

A partial listing of the 12 HFA nodes in the loaded image file.

```

dma_dted.img dictionary

dictionary (2618 bytes):
{1:lversion,1:lfreeList,1:lrootEntryPtr,1:
sentryHeaderLength,1:lDictionaryPtr,}Ehfa
aFile,{1:lnext,1:lprev,1:lparent,1:lchild,1:l
data,1:l dataSize,64:aname,32:atype,1
:tmmodTime,}Ehfa_Entry,{16:clabel,1:lhead
rPtr,}Ehfa_HeaderTag,{1:lfreeList,1:lfree
Size,}Ehfa_FreeListNode,{1:lsize,1:lptr,}
Ehfa_Data,{1:lwidth,1:lheight,1:e3:themat
ic,athematic,fft of real-valued
data,layerType,1:e13:u1,u2,u4,u8,s8,u16,s
16,u32,s32,f32,f64,c64,c128,pixelType,1:l
blockWidth,1:lblockHeight,}Eimg_Layer,{1:l
width,1:lheight,1:e3:thematic,athematic,
fft of real-valued
data,layerType,1:e13:u1,u2,u4,u8,s8,u16,s
16,u32,s32,f32,f64,c64,c128,pixelType,1:l
blockWidth,1:lblockHeight,}Eimg_Layer_Sub
Sample,{1:e2:traster,vector,type,1:lDictio
naryPtr,}Ehfa_Layer,{1:sfileCode,1:lOffse
t,1:lsize,1:e2:false,true,logvalid,1:e2:n
o compression,ESRI GRID
compression,compressionType,}Edms_Virtua
lBlockInfo,{1:lmin,1:lmax,}Edms_FreeIDList
,{1:lnumvirtualblocks,1:lnumobjectspblo
ck,1:lnextobjectnum,1:e2:no
compression,RLC
compression,compressionType,0:poEdms_Virt
ualBlockInfo,blockinfo,0:poEdms_FreeIDList,
freeList,1:tmmodTime,}Edms_State,{0:pcst
ring,}Emif_String,{1:toEmif_String,algorit
hm,0:poEmif_String,nameList,}Eimg_RRDName
sList,{1:toEmif_String,projection,1:toEmif
_String,units,}Eimg_MapInformation,{1:toEmi
f_String,dependent,}Eimg_DependentFile,{1
:toEmif_String,imageLayerName,}Eimg_Depend
entLayerName,{1:lnumrows,1:lnumcolumns,1:
e13:EGDAL_TYPE_U1,EGDAL_TYPE_U2,EGDAL_TYPE_U
4,EGDAL_TYPE_U8,EGDAL_TYPE_S8,EGDAL_TYPE_U16
,EGDAL_TYPE_S16,EGDAL_TYPE_U32,EGDAL_TYPE_S3
2,EGDAL_TYPE_F32,EGDAL_TYPE_F64,EGDAL_TYPE_C
64,EGDAL_TYPE_C128,datatype,1:e4:EGDAL_SCAL
AR_OBJECT,EGDAL_TABLE_OBJECT,EGDAL_MATRIX_O
BJECT,EGDAL_RASTER_OBJECT,objecttype,}Egda
_BaseData,{1:*
```

Corresponding dictionary defining the structure of each possible node.

```

dma_dted.img tile o...

Edms_State:
numvirtualblocks: 64
numobjectspblock: 4096
nextobjectnum: 0
compressionType: 0
Edms_VirtualBlockInfo:
tile offsets:
2768
10960
19152
27344
35536
43728
51920
60112
68304
76496
84688
92880
101072
109264
117456
125648
133840
142032
150224
158416
166608
174800
182992
191184
199376
207568
215760
223952
232144
240336
248528
256720
264912
273104
281296
289488
297680

```

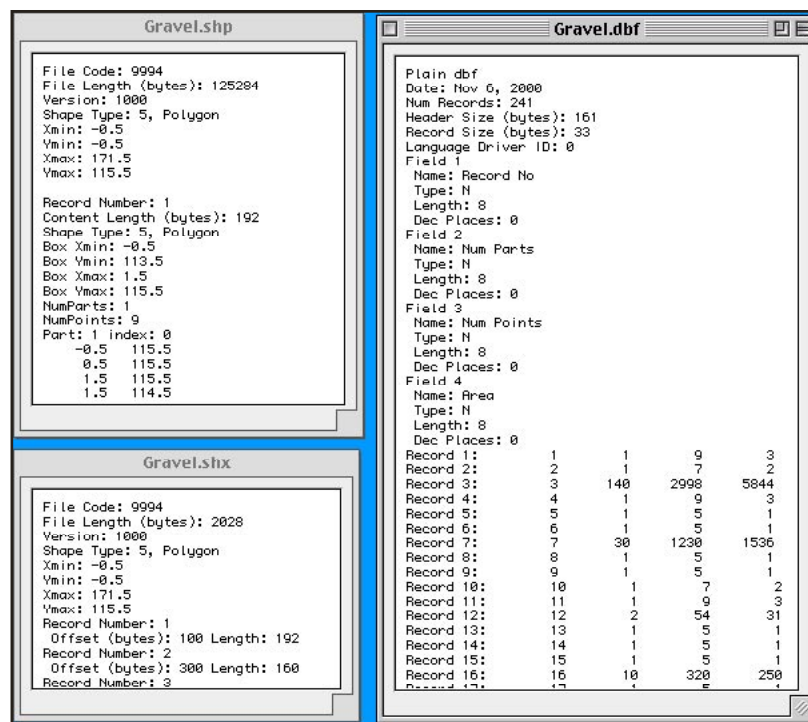
A listing of offsets showing the location of each image data block within the file.

Any loaded image can be saved in the Imagine® *.img format via menu Save As... and selecting ERDAS Imagine® as the type. Note: HyperCube does not generate the histogram or statistics of HFA nodes when saving, only those nodes needed for loading and displaying.



ESRI Shape Files

HyperCube can load and display certain types of ESRI (Environmental Systems Research Institute, Inc.) Shape Files. In particular are those produced by HyperCube via the [Class Map Editor Export](#) function and those produced by menu Applications -> [Contour](#). Shape files consist of a main file that should have a .shp extension along with the index file (.shx extension) and a dbase file (.dbf extension). The later two file's contents can be displayed in a scrollable text window by selecting File- > Open As..., choosing ESRI Shape and selecting the desired file extension type. Depending upon their size (number of records) multiple output windows may be generated. For convenience, the main shape file itself does not have to have the .shp extension whereas the other two do require their proper extension.

Shown below is the result of opening the three shape files that were exported by the [Utilities -> Class Map Editor](#). Each text window contains a parsing and because of the possible large size of shape files not all of the records may be shown. Also, the shape plot sizing is arbitrary unless it was exported from HyperCube.



The three shape files and the corresponding plot that were produced via the Utilities->Class Map Editor Export function. Only a portion of each file is shown in the text windows. See section [Class Map Editor](#) for more details.




*Choosing an *.shp file results in the above dialog so that you may skip the records output.*

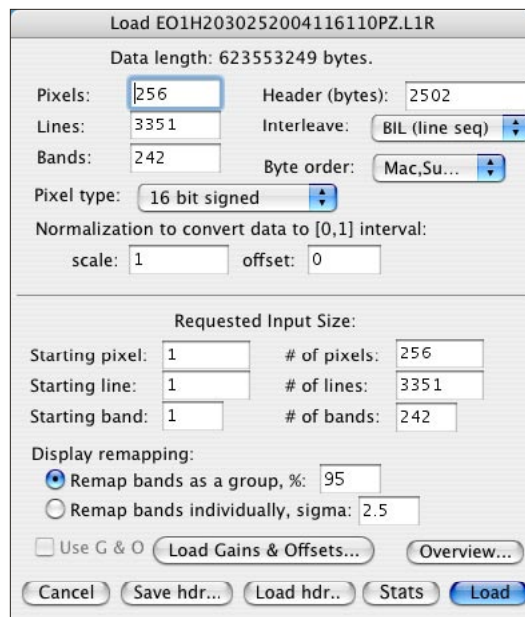
HDF Files

The Hierarchical Data Format (HDF) was developed at the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign. Most of the image/data formats described in the NCSA HDF Specification and Developer's Guide, Version 4.1r5, November 2001 can be loaded and saved in HyperCube. The strong exceptions are: the complex data type is not supported and the number of dimensions (rank) must be less than or equal to three for Numeric and Scientific Data Groups.

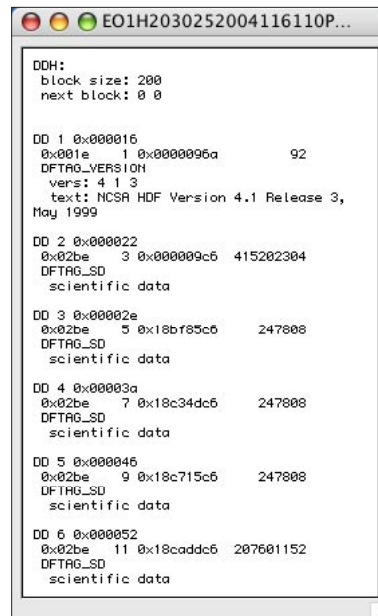
Load an HDF image by selecting menu **File -> Open As...** and choosing **HDF** as the file type. This results in the following dialog listing all of the single images and Scientific Data sets residing in the file. The list displays the characteristics of each embedded image. Select 1 or more images (use shift click for a contiguous region and command click for discontinuous selections) and then **Load**. If an image is multiband (number of bands greater than 1) then a subsequent dialog appears exactly as described in section [Image Format Dialogs](#) allowing you to load any portion or overview of the multiband image.



A listing of the 5 images within the HDF file. The first and last images are multiband with pixel depths of 16 and 8 bits per pixels respectively. Selecting the first image brings up the dialog shown on the right (see section [Image Format Dialogs](#)).



*If menu item **Edit -> Options -> Detailed Messages** is toggled on (checked) then a listing of all of the recognized tags in the header portion of the HDF file is generated. This is a tiny fraction of the tags contained in the image example shown above.*



JPEG Files


Baseline JPEG files may be loaded and saved by HyperCube. Specifically, these are defined by ISO 10918 as:

- DCT - base process
- 8 and 16 bit samples per component (a little beyond baseline)
- Sequential
- Huffman coding
- 1,2,3 or 4 components
- Interleaved and non-interleaved scans

To load a JPEG image file select **File -> Open As...**, **JPEG** and choose the file. If the image is not in JPEG format then you will get an alert message. Once loaded, a JPEG image is treated like any other loaded image and can be saved in any of the image formats.

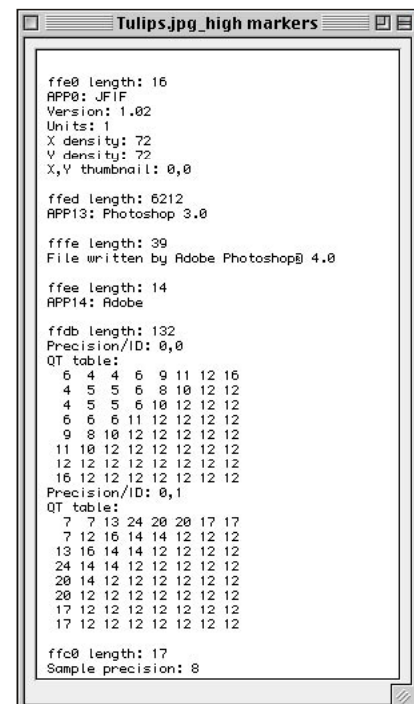
Saving an image, either gray, pseudo or true color, is done via **File- > Save As...**, **JPEG** when the image is the active window. This displays a subsequent dialog (after the standard save file dialog) that allows you to choose the level of compression as shown below. Also shown on the right is a text file dump of the various JPEG markers for the loaded image which appears if **Edit -> Options->Detailed Messages** has been checked before the image is loaded.

Note: Because JPEG is a lossy format you cannot save a classmap in JPEG format. The color correspondence will be lost, use NITF, PICT, SRF or TIFF.



Three levels of compression available to you when saving an image in JPEG format.

A portion of the JPEG markers defining the contents a loaded JPEG image. This text window appears automatically if Edit->Preferences->Detailed Messages is checked.



LASF/ASCII LIDAR Files

Files of type LASF (e.g., LIDAR point cloud data) can be loaded using menu **Open As** and choosing LASF as the **Format**. The file is converted from its individual points to a raster image that can be resaved in LASF or any other HyperCube format. Any combination of elevation, intensity, or RGB values can be extracted (see section [LASF Point Cloud Viewer](#) for displaying the actual points). In addition, LASF files that have been converted into ASCII columnar text files, where columns represent X,Y,Z intensities and RGB can be similarly processed by using **Open As... ASCII LIDAR** (this description follows the current LASF details). The elevation and intensity images can be output as an anaglyph.

Point cloud data doesn't have a uniform spacing so the resulting raster image will contain some "holes" and some overstrikes (multiple points mapped to the same pixel and line). The larger the scale (expansion) the more voids and fewer overstrikes. The **Returns** popup menu permits keeping a specific elevation in overstrikes (e.g., the lowest elevation mapped to a pixel/line). The dialogs, shown below, specify which return values to rasterize, the output scale and an option to interpolate single pixel holes by using neighboring raster values.

Generally, the scenario for loading LASF files is to load an [overview](#) and then select a rectangular area to be loaded at full resolution. The resulting image will automatically be given the [reference](#) coordinate system defined by the points' X and Y values.

See section **Utilities** ([Merge LAS Files](#)) for a method to merge two or more LAS files having common UTM zones into a single LAS file.

This file does not contain RGB values. Only elevation and intensity are available.

See section: [LASF Point Cloud Viewer](#).

Simple raster holes can be interpolated.

See following pages for an example.

Exclude zero only affects the display.

Image file details.

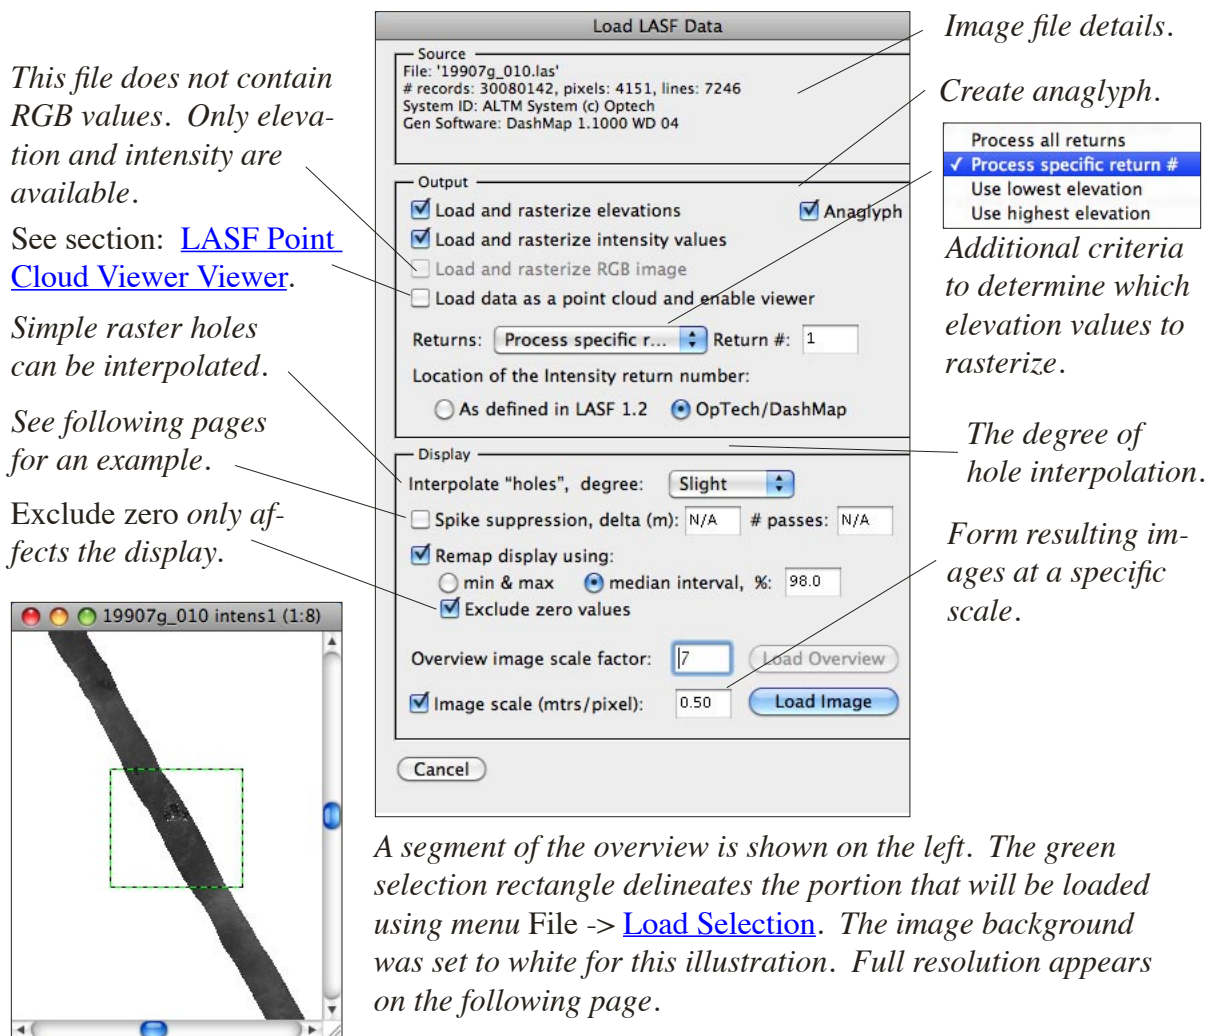
Create anaglyph.

Additional criteria to determine which elevation values to rasterize.

The degree of hole interpolation.

Form resulting images at a specific scale.

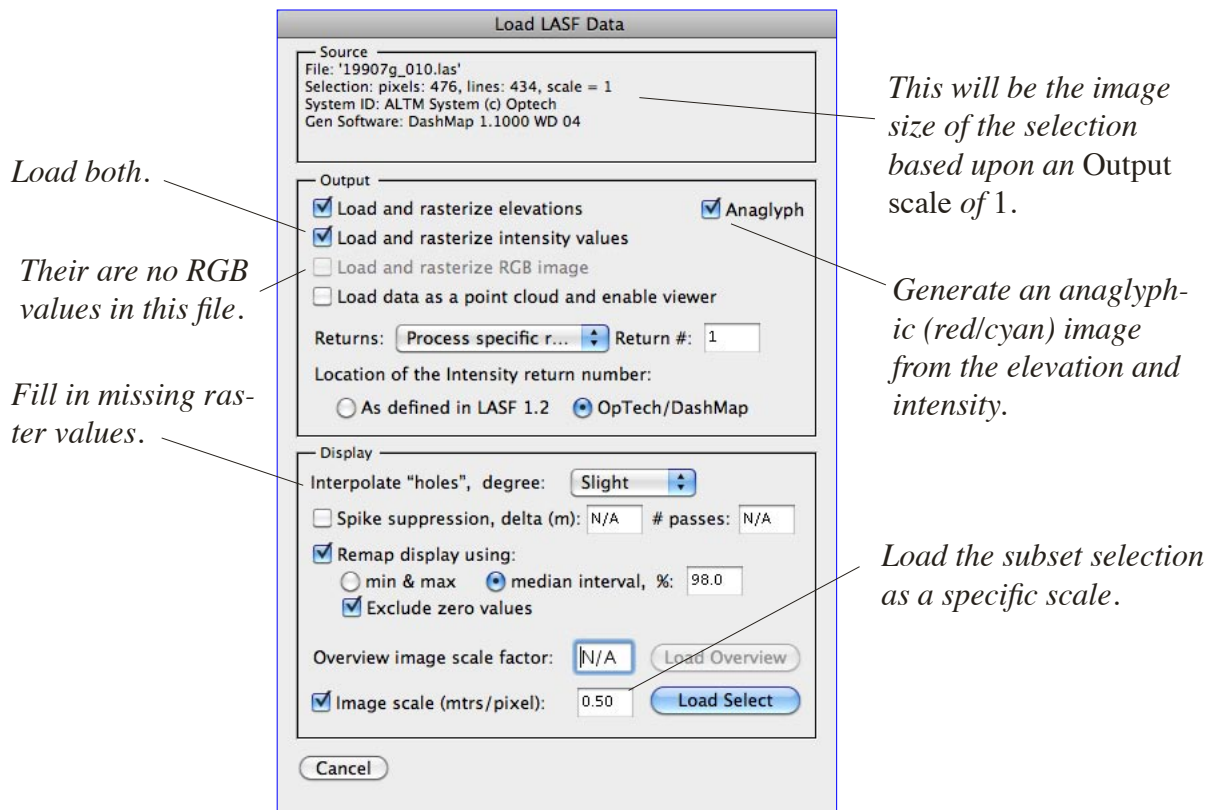
*A segment of the overview is shown on the left. The green selection rectangle delineates the portion that will be loaded using menu **File -> Load Selection**. The image background was set to white for this illustration. Full resolution appears on the following page.*




The 'Load LASF Data' dialog box is shown with the following settings:

- Source:** File: '19907g_010.las', # records: 30080142, pixels: 4151, lines: 7246, System ID: ALTM System (c) Optech, Gen Software: DashMap 1.1000 WD 04.
- Output:** ☒ Load and rasterize elevations, ☒ Load and rasterize intensity values, ☐ Load and rasterize RGB image, ☐ Load data as a point cloud and enable viewer, ☒ Anaglyph.
- Returns:** Process specific r... Return #: 1.
- Location of the Intensity return number:** ☐ As defined in LASF 1.2, ☒ OpTech/DashMap.
- Display:** Interpolate "holes", degree: Slight, ☐ Spike suppression, delta (m): N/A, # passes: N/A, ☒ Remap display using: ☐ min & max, ☒ median interval, %: 98.0, ☒ Exclude zero values.
- Overview image scale factor:** 17, Load Overview.
- Image scale (mtrs/pixel):** 0.50, Load Image.


The preview window shows a dark, diagonal line on a white background with a green selection rectangle.



This dialog appears after selecting an area on the overview and choosing menu File -> [Load Selection](#). Both the return intensity and Z values are selected as well as an anaglyph and the option to interpolate any zero (missing) rasterized values.




The LIDAR Return number: 1 intensity.




The corresponding LIDAR Z values.

The dialog to the right was used to load and rasterize the two image subsets shown below which were subsequently generated from menu Applications -> [Shaded Relief](#). Note that Use lowest elevation has been selected in order to minimize vegetation effects.

A spike is a rasterized elevation that exceeds each of its neighbors by the value set in delta (m). # Passes ranges from 1 to 12 and determines the pattern of surrounding elevations that form the comparison.

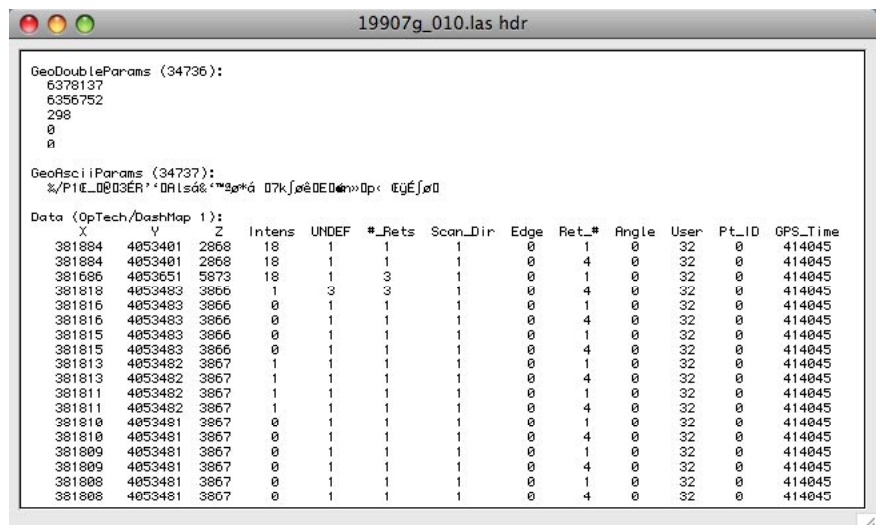
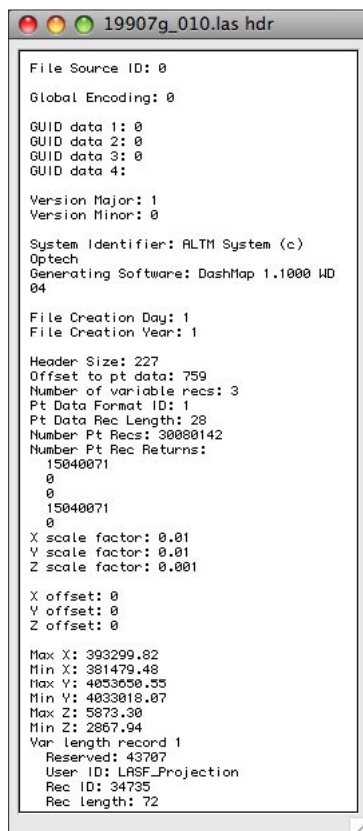


A [shaded relief](#) subset of an LASF file in which the Returns popup menu was set to Use lowest elevation. The black dots indicate shadows cast by vegetation spikes.





A [shaded relief](#) of the same area after selecting Spike suppression and 12 passes in the dialog shown above.

LIDAR data is sometimes stored in a text (ASCII) format consisting of three or more columns. However, since there is no detailed header, processing requires doing an initial parsing of the file to determine its X and Y boundaries. This requirement along with ASCII to binary decoding results in much longer execution times than having the file in LASF. Plus, the converted raster image may be in a local X,Y coordinate system. However, this can be corrected by using menu [Edit -> Options -> Image to Ref Coords](#). Shown on the lower half of the following page is the dialog after opening a file as ASCII Lidar.



If menu Edit -> Options -> Detailed Messages is selected before opening a LAS file then header information and approximately the first 30 point cloud values will be listed in a text window (shown here broken into two parts for display purposes).

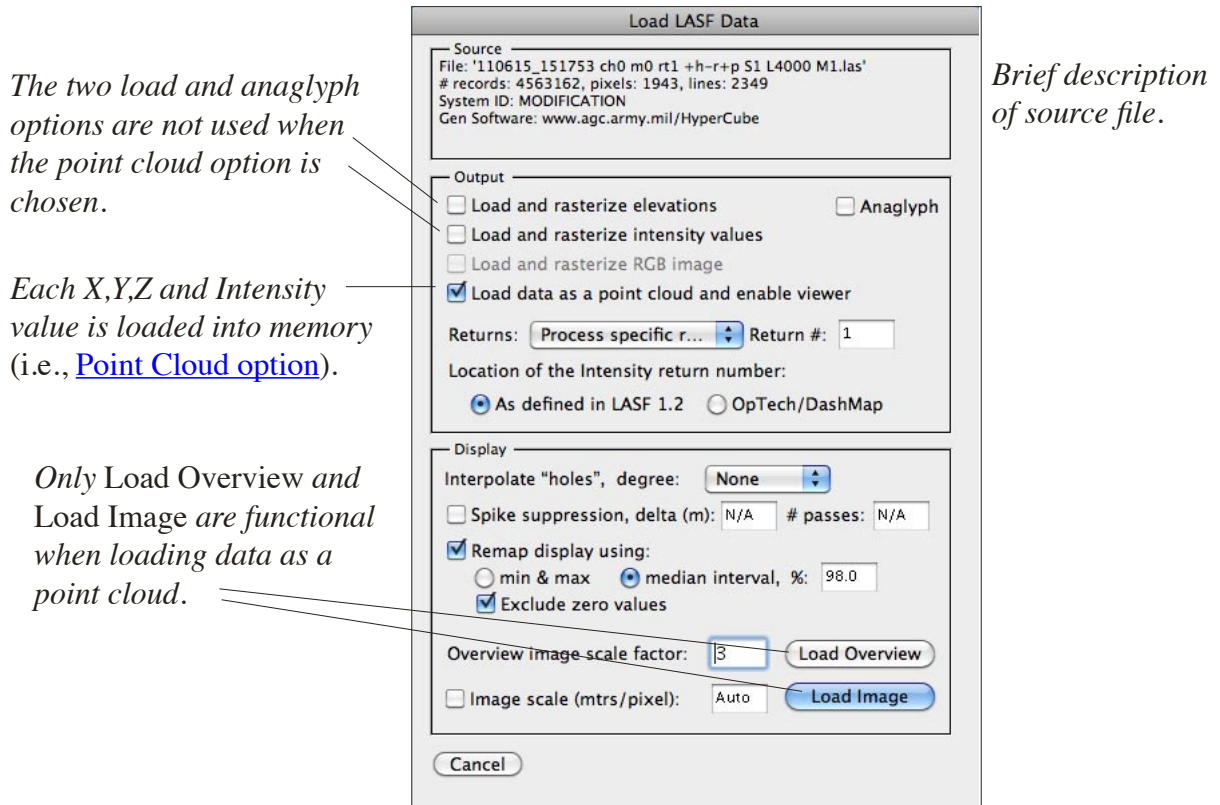


This text window appears after parsing the file and displays information about the contents of the file.

The ASCII LIDAR dialog above is very similar to the [Load LASF](#) dialog. The information box at the top lists the first 4 records of the file and in this example shows that the first record lists the order of the data which consisting of 7 columns: Northing, easting, elevation, intensity and red, green and blue values.. The dialog contains the Starting record edit field to allow skipping initial records. If you load the overview and select menu [File -> Load Selection](#) another dialog will appear exactly like the one above permitting you to modify some of the previously selected options (this works in the same manner as described in [Load LASF](#)). Also, see section [ASCII LIDAR to LASF](#).

LASF Point Cloud Viewer

LASF data may be loaded as a point cloud, instead of a rasterized image, and viewed as points projected upon a planar surface from a user chosen view point and direction. The load dialog is the same as used for loading LASF (section [LASF/ASCII LIDAR Files](#)). By selecting checkbox: **Load data as point cloud...** the data is extracted and stored as individual X,Y,Z and Intensity values. A reduced resolution overview image window is created along with a full resolution subset of the central portion of the overview image. Accompanying these is an interactive dialog that defines the position and viewing angle of the projected point cloud upon the full resolution view window (see following page).



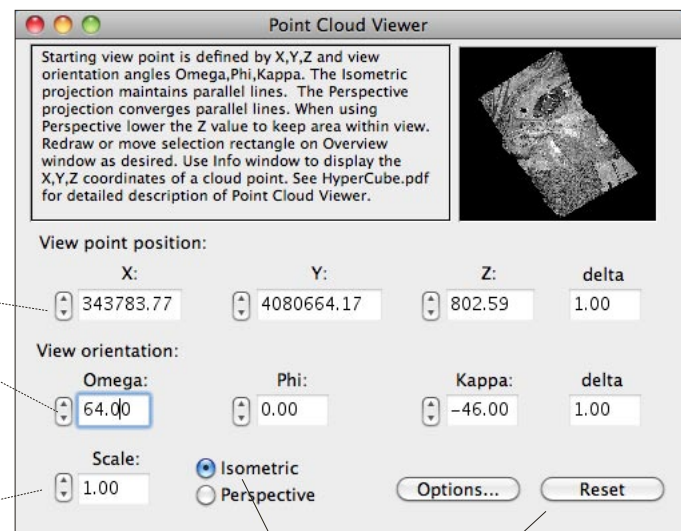
When loading a point cloud consideration must be given to the amount of data that the program retains in memory. Therefore it's best to first load the LASF without selecting the point cloud option using **Load Overview** as described in section [LASF/ASCII Lidar Files](#). Then, draw a selection rectangle around the area of interest and choose menu **File -> Load Selection**. The load dialog (above) will reappear and this time select the point cloud option followed by **Load Image**. Three new windows will appear: a point cloud **overview** image (separate from the LASF load overview) depicting the chosen selection rectangle, a detailed **view** of the point cloud and an interactive **dialog**. The **overview** image is the data corresponding to menu **File -> Load Selection** function and is the portion of the total point cloud that is retained in memory. It also contains a selection rectangle and on the initial point cloud load this area is projected onto the **view** image using the values in the [Point Cloud Viewer dialog](#). You may then freely move or resize the point cloud **overview** selection or use any of the interactive dialog controls to change the projection upon the **view** image. Examples of the point cloud viewer appear on the next pages along with two types of profile plots.

The Point Cloud Viewer is one of three windows that appears after choosing Load data as point cloud... on the [Load LASF Data](#) dialog. Use any of the controls interactively or type a value into the edit field and press enter/return.

View point position in coordinates defined by the LASF data.

View point direction using omega, phi, kappa as rotations about the X,Y and Z axes respectively.


Magnify/de-magnify the view image.



Restore all values to the initial load.


Specifies the type of projection from the point cloud onto the View image.

The Point Cloud Overview image, left, with a selection rectangle of the area that is projected onto the View image shown below using the values in the above dialog.




The projection of the point cloud values that are delineated by the selection rectangle in the overview image. The projection type is Isometric and Interpolation has been chosen in the [Options](#) dialog (described on the next page).



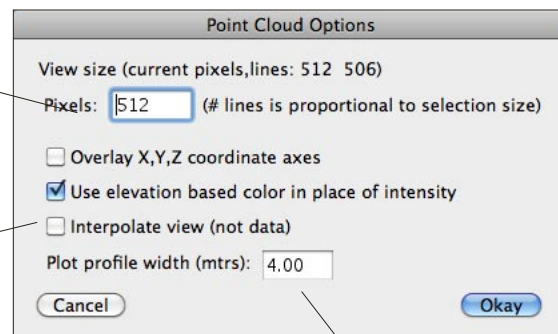


The selected point cloud area projected through a viewing angle of $\omega = 73$ degrees is displayed below.




The width of the View image can be edited. The height is determined by the aspect ratio of the Overview selection rectangle.




Interpolation is only a cosmetic modification of the View image. The underlying data is unchanged.




Planar width for [profiles](#).





The entire point cloud overview area (using no selection rectangle) has been projected through an angle of $\omega = 60$ degrees and subjected to elevation colorization in the Point Cloud Options dialog (above). Red represents the highest and blue the lowest elevations.



The above two images represent the same view ($\omega = 60$ degrees) with different projections. The left is using Isometric geometry and the right one via Perspective geometry. The former preserves parallel lines while the later does not. Each can be used for measuring points (see [Info](#)) and generating profiles (see).



The center image is an Isometric projection with $\omega = 66$ degrees and $\kappa = 160$ degrees with a profile drawn as shown. At the left is the corresponding [profile plot](#). This plot displays the elevation of each point cloud value lying on the line. On the next page is an example of a profile plot that represents a planar slice through the point cloud and perpendicular to the ground.



This is an additional plot produced via menu [Functions -> Plot -> Profile](#) if the source image is a Point Cloud View window (above). Only the starting and ending coordinates are used to define a plane perpendicular to the X,Y coordinate plane. All point cloud coordinates lying within a given width of the plane are used to create a scatter plot of elevation verses distance along the profile line.

NITF Files

NITF (National Imagery Transmission Format) images are loaded via menu **File -> Open As...** and then choosing the **NITF** item in the **Format** popup menu. HyperCube can read versions 2.0 and 2.1 NITF data, gray, pseudo, true color and multiband images. It can also extract embedded JPEG (DCT) and JPEG (2000) images (see [Wrap JPEG in NITF](#) for inserting JPEG within NITF). Any graphics or text within the NITF file is ignored. If the image is multiband, larger than 2048 (pixels or lines) or the pixel depth is greater than 8 bits (e.g., short or float) you are presented with the dialog shown below. This allows you to load either an overview of the image (necessary if very large) or the entire image. Overview images are then used to select full resolution subsets that are subsequently loaded using menu **File -> Load Selection** (see section [Overviews](#)). If there is insufficient memory to load then you will be alerted with a message stating how much additional memory is needed. Saving a 16 or 32 bit image as NITF saves the data as short or float.

If the NITF image is multiband you may select the starting band and the number of bands to load and if **Overview** is chosen, the overview band and scale as well. To exploit a NITF multiband image you must first save the loaded image cube. However, there are possible memory constraints doing this and the best technique is to use menu [Utilities -> Reformat Cube File](#) to convert the NITF file into an image cube file.

A description of the image.

What part of the image to load.

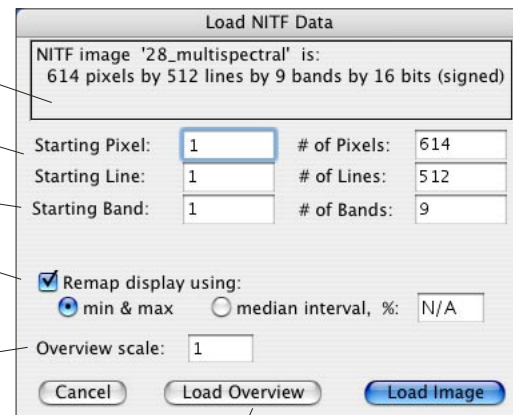
Initially, set to the image size.

If multiband which bands to load.

How to remap the data for display.

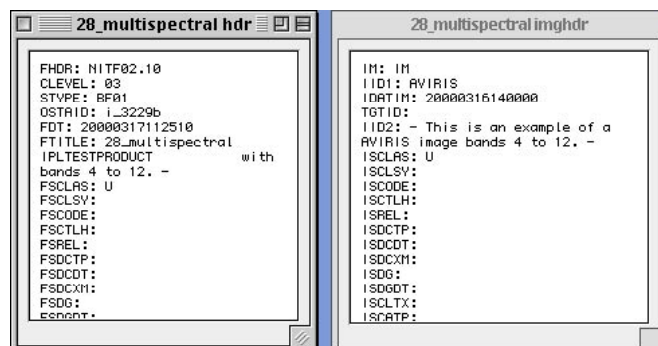
Also, see section [Pixel Exclusion](#).

For the given size image, the scale automatically computed to display the whole image within the monitor.



Select either an overview or the image at full resolution. You will be notified if insufficient memory.

*If menu **Edit -> Options -> Detailed Messages** is checked then each NITF image load will produce windows listing the main header and image header contents. Also, if the image header defined a coordinate system, use the Info Window to display them.*




PNG Files

A subset of Portable Network Graphics (PNG) files can be loaded and saved. The subset consists of 8 and 16 bits per pixel gray scale and 8 and 16 bits per color images. Images of other bit lengths (e.g., 1,2,4 bits) can not be loaded or saved. PNG files containing an alpha channel can be loaded but the transparency (alpha byte) is ignored. If an image contains a palette, the palette will be discarded after loading and the image will be converted to direct gray bytes or a true color image depending upon the number of components. Image interlacing is also not supported.

Currently, HyperCube does not generate an optimized compressed file when saving a PNG image. It uses the standard fixed Huffman tables and the output file will be slightly larger than an equivalent TIFF file. However, the file is correct and lossless.

Shown below is the result of loading a PNG file after setting menu Edit -> Options -> Detailed Messages (applicable for any type of file loading) in order to list the PNG header which in this case has been synopsized.



Uncompressed blocks: 0
Fixed blocks: 0
Dynamic blocks: 74
Filter types:
 none(0): 0
 sub(1): 10
 up(2): 0
 average(3): 121
 paeth(4): 778

IHDR

Width: 918
Height: 909
Bit depth: 8
Color type: 2
Compression method: 0
Filter method: 0
Interlace method: 0
CRC: 1920493681

RGB triplets: true
Alpha channel: false
Bytes/pixel: 3
De-filtering interval: 3

Chunks:	name	size	dataStart	CRC
iCCP		8482	41	1158558444
pHYs		9	8535	2024095606
IDAT		8192	8556	3518965669
IDAT		8192	16760	3907572812
IDAT		8192	24964	2565274342
IDAT		8192	33168	3030120714
IDAT		8192	41372	3541105835

The details at the left describe the composition of the compression used and the above detail (a segment of the actual list) shows the type of image and some of the PNG 'chunks'.

TIFF Files

TIFF (Tagged Image File Format) images are loaded via menu File -> Open As.... HyperCube can read 8 bit and higher baseline TIFF 6.0 images plus 16 and 32 bit tiled format images including LZW compression. Embedded JPEG encoding is not included. JPEG DCT images can be read by themselves (section [JPEG Files](#)). If the image is larger than 2048 (pixels or lines) or the pixel depth is greater than 8 bits (e.g., short, float or 16 bits per color) you are presented with the dialog shown below. This allows you to load either an overview of the image (necessary if very large) or the entire image. Overview images are then used to select full resolution subsets that are subsequently loaded using menu File -> Load Selection (section [Overviews](#)). Saving a 16 or 32 bit image as TIFF saves the data as short or float in strip format even if the source format was tiled.

If the TIFF image contains Geo tags (i.e., it's a GeoTIFF) then any coordinates will be extracted and used as the image reference (section [Image to Reference Coordinates](#)). A single tie-point GeoTIFF can be converted to a referenced image in geographics (lat/long) with the geometry defined by [RPC](#)'s. This also works if the GeoTIFF is in UTM coordinates. Saving an image that contains [reference coordinates](#) as TIFF saves it as GeoTIFF and if saved as NITF with included RPC. If an [RPC](#) text file in the same directory has the same image base name concatenated with [_rpc.txt](#) then it will be automatically read and used to form the geo reference (see [Stereo Compilation](#)).

If a file with extension .tfw exists having the same base name as the TIFF file and resides in the same directory it will be read and parsed as if it was an ESRI world geo-reference. The geo-reference contained in it will be applied to the TIFF image.

Image description.

What part of the image to load.

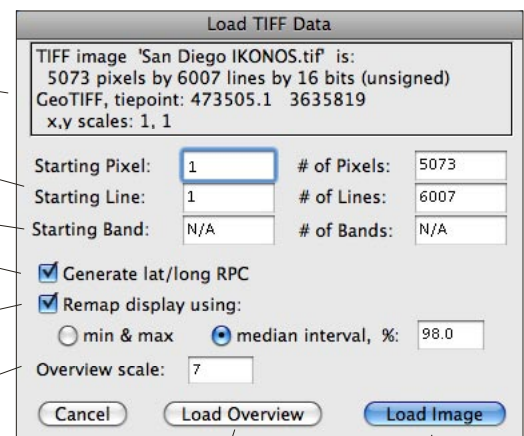
Initially, set to the image size.

Not used for TIFF, see NITF.

This GeoTIFF image has a single UTM tie-point and can be converted to lat/long. Also, see [GeoTIFF, JPEG to NITF](#).


How to remap the data for display. Also, see section [Pixel Exclusion](#).

For the given size image, the scale automatically computed to display the whole image within the monitor.



The 'Load TIFF Data' dialog box displays information for the file 'San Diego IKONOS.tif'. It shows the image dimensions (5073 pixels by 6007 lines by 16 bits unsigned) and GeoTIFF tiepoint (473505.1, 3635819) with x,y scales of 1, 1. Below this, there are input fields for 'Starting Pixel' (1), 'Starting Line' (1), and 'Starting Band' (N/A), along with their respective counts: '# of Pixels' (5073), '# of Lines' (6007), and '# of Bands' (N/A). There are two checked options: 'Generate lat/long RPC' and 'Remap display using:'. Under 'Remap display using', the 'min & max' radio button is selected, and the 'median interval, %' is set to 98.0. The 'Overview scale' is set to 7. At the bottom, there are three buttons: 'Cancel', 'Load Overview', and 'Load Image'.

Select either an overview or the full image .



If menu Edit -> Options -> Detailed Messages is checked then each TIFF image load will produce a window listing the main TIFF tags and any Geo tags (left and right are same window, scrolled differently).



SDF Files

This section uses the Riegl Laser Measurement System's (www.riegl.com) sdfifc32.dll or sdfifc64.dll and is only applicable in the Windows version of HypeCube. Both dlls are included in the HyperCube.zip download folder. SDF (Sample Data Files) files consist of records of Full Wave Form (FWF) Lidar. Each record contains a sequence of samples of the return intensity as a function of time. HyperCube uses the dll to extract the samples and additional timing information in order to reconstruct the location of the samples. It accomplishes this with the help of an auxiliary file, [SBET](#), (Smooth Best Estimate Trajectory) that gives the location and orientation of the Riegl Lidar scanner as a function of GPS time. This time is in sync with the GPS timing embedded in each FWF record. The SBET file should be in the same directory as the SDF file and be named "sbet_Mission 1.out" and the Riegl sdfifc dll should be in the same directory as the HyperCube executable, if not, then dialogs will appear to search for them. A sample sdf file, Duncan Knob.zip, is available on the web page.

Processing SDF files starts with menu File -> Load... SDF. The file records are first parsed (progress indicator) followed by the Load SDF Data dialog shown below separated into four sections: **Source**, **Parameters**, **Coordinate Permutations** and **Outputs**. Familiarity with SDF data and FWF is assumed.

Load SDF Data

Source
File: '110615_151753.sdf'
Records: 5779450, records/line: 1168, scan lines: 4948
Chnl 0 min, max range (m): -75.6982 351.685
Chnl 1 min, max range (m): -75.889 2.55181e+011
PRR: 353 kHz, LRR:198 Hz

Samples/channel: 160 160 0 24 0 0 0 0
Channel counts: 575447 610618 0 583486 0 0 0 0
Sosl_count Hist: [0]500 [1]4241 [2]11056 [3]538520 [4]20707
[5]8962 [6]0 [7]0

Parameters
☒ List SBET ☒ List SDF GPS offset (sec): 15.00
☒ Plot flight path
Start line: 1 # of scan lines: 2000
Process channel number: 0
Multiple-Time-Around Zone: 1
Range and amplitude determination: Multiple returns Param: 5
☒ Plot heading, roll, pitch ☒ Plot deltas: N, E, range
☒ List a sampling of processed SDF records
☒ Include range plots
☒ Include specific return count range plots
Return count: 5

Coordinate Permutations: (H)eading, (R)oll, (P)itch
Each checkbox selection produces an output.
☒ H, R, P ☐ H, P, R ☐ R, H, P
☐ P, H, R ☐ R, P, H ☐ P, R, H
Select sign direction. ± does SBET and negation (2 outputs).
☒ SBET H ☐ SBET R ☒ SBET P
☐ Negate H ☒ Negate R ☐ Negate P
☐ ± H ☐ ± R ☐ ± P

Outputs
☒ Intensity image ☒ Elevation image ☒ Anaglyph
☒ LAS file ☐ FWF cube file ☐ Geo cube file
☒ Interpolate ☐ Mask image ☒ Rtrn # image
☐ List times
☐ Image scale (mtrs/pixel): Auto
☒ Overlay markers, value: 128

Cancel Process

This dialog appears after SDF file parsing. The example shows a file consisting of 5779450 records, 1168 records/scan line and 4948 lines. Only scan lines 1 through 2000 will be processed. The SBET file Roll direction is set to Negate and four types of outputs will be generated: Intensity, Elevation, Anaglyph and an LAS file.

Source (extracted during parsing):

File: The name of the SDF file.

Records: Number records in file, records/line and scan lines.

Chnl 0 min, max range (m): A very rough estimate of channel 0 range (not used as is).

Chnl 1 min, max range (m): Same for channel 1.

PRR and LRR: The calculated Pulse Repetition and Line Repetition Rates respectively

Samples/channel: A count of the number of return sample for each channel

Channel counts: How many time a specific channel appears within all parsed records.

Sobl count Hist: A histogram of the number of Sobl blocks per record,

e.g., the values shown on previous page, Load SDF Data dialog, are interpreted as:

500 records with zero blocks, 4241 records with 1 block, 11056 with 2 blocks,...

Note: Parsing does not examine all records, only a subset of at most 1500000.

Parameters (set by user):

List SBET - If checked then create a text list of approximately 10 SBET records.

List SDF - List the contents of the SDF file header.

GPS offset - The number of nanoseconds needed to sync SDF with SBET values.

Plot flight path - Display a scatter plot of the flight path (Northing verses Easting).

Start line: The starting scan line at which to extract data.

of scan lines: How many scan lines to process.

Process channel number: Which return channel's data to extract.

Multiple-time-Around Zone: Interval between successive laser pulses. A measure of how many pulses are in flight at the same time (used to match the outgoing pulse with its corresponding return pulse).

Range and amplitude determination (pull down):

Greatest amplitude

N'th temporal return

Mean amplitude and range

Mean amplitude and greatest amplitude range

Multiple returns (prohibits Cube outputs)

Generally, the maximum distinct number of returns that the software can detect is 5.

Param:

none

1,2,3,...

none

none

1,2,3,...

Plot heading, roll, pitch - If checked then display each as a function of scan line #.

Plot deltas: N, E, range - Display Northing, Easting and range corrections.

List a sampling of processed SDF records - A text listing of selected SDF records.

If checked then the following options are available:

Include range plots: Graph of range (mtrs) verses sample for a specific record

Include specific return count range plots: - Also generate range plots of records that contain "Return count" number of returns.


Coordinate Permutations: (H)eading, (R)oll, (P)itch

Each of the six checkboxes defines the order of the three rotations of the platform and is used by the program to create a rotation matrix. For each checked box a complete set of **Outputs** is generated. Choosing more than one is a big help in determining, via visual inspection of the outputs, the proper order without restarting the load process. Each set of outputs will have the rotation order embedded in the name. The most likely order for SBET is shown in the Load SDF Data dialog, e.g., H, R, P.

In addition to the rotation order, there can also be an uncertainty in the interpretation of the proper sign direction of each rotation. The various radio buttons are used in a similar fashion as the rotation order described above. The first row of buttons uses the rotation angle signs as given in the SBET file. The second allows you to negate the SBET value for that angle (i.e., if the SBET angle is positive then make it negative and if the SBET value is negative, make it positive). The last row of buttons allows both: every combination of positive and negative signs.

Note: The rationale for the permutations and sign directions is due to the inconsistencies of their usage and definitions in the SDF and SBET files. This simplifies determining which are correct. If each permutation checkbox is set and all three plus/minus radio buttons in the last row are chosen then there will be 48 sets of outputs and can require several hours of run time.

An example of the output image (file) naming convention:



Outputs (selected by checkbox)

Intensity image - The rasterized return intensity values consisting of ‘records/line’ number of pixels and ‘# of scan lines’ lines (see **Parameters** section) in geo-referenced coords.

Elevation image - The rasterized elevations matching the same size as the Intensity image also in a geo-reference coordinate system.

Anaglyph - A red and cyan anaglyphic generated from the Intensity and Elevation images both of which must be selected as outputs.

LAS file - A point cloud file of the extracted elevation and intensity values. This file can be loaded using menu [File -> Load LAS](#) and, additionally, viewed with the [LASF Point Cloud Viewer](#). Best used with ‘Multiple returns’ chosen.

FWF cube file - Generate a cube file consisting of each SDF record such that the nano-second spaced samples are stored in BIP (Band Interleaved by Pixel) format. This option is not available if the Multiple returns option in section **Parameters** is selected. The data is not geo-referenced. It is stored in the order of the records. See Geo cube file next for a geo-referenced version.

Geo cube file - This option uses the FWF cube file as the source and repositions each BIP cube vector to a corrected geo-position using nearest neighbor sampling. Each band will exactly match the geometry of the Intensity and Elevation images. The ‘FWF cube file’ option must be selected to enable this output.


Interpolate - When selected, a small amount of interpolation is applied to the Intensity and Elevation images to fill holes. A hole is filled if it is surrounded by at least 5 out of 8 possible non-null neighbors.

Mask image - A same size image generated during interpolation to mark the non-null intensity positions. Can be remapped for viewing to check the sparsity of the Intensity image.


Rtrn # image - An image matching the geometry of the Intensity and Elevation image where each corresponding pixel is the number of returns extracted from the SDF for that geo-referenced position. This is a very low value image (0 - 5,6 or 7) and can be directly viewed by using [Histogram equalization](#) or by using [Image Arithmetic](#) to remap the pixel values. The image shows where the most return ‘activity’ occurs (generally in trees).

Image scale ... - The software computes a nominal scale for the output images. If this checkbox is selected then a specific scale (mtrs/pixel) can be entered. Note that as the value is decreased then the density of the number of holes in the Intensity and Elevation images decreases.


Overlay markers ... - A [value](#) traversal line is overlaid upon the Intensity image delineating the intersection of the collection platform normal and the ground (diagnostic tool).






The elevation image resulting from the values and options similar (4500 lines) to those in the Load SDF Data dialog.



The matching intensity image. Both images have interpolation applied.



A false color composite image in which color is used as a representation of the shape of each vector of the generated FWF Geo cube file. In this example, each vector of the BIP cube is composed of 160 samples of the return pulse. The vectors were then transformed using the Discrete Cosine Transform (DCT) producing 160 coefficients. The image was constructed using DCT coefficient 1 (actually coefficient 0 in transform numbering) for the blue component, coefficient 2 for the green component and coefficient 3 for the red component. This is just one of a myriad number of possible combinations. Other methods such those in [spectral classifications](#) can be used to extract meaningful structure from the FWF geo cube.



From top to bottom these plots respectively show the scanner platform's Heading, Roll and Pitch in degrees as a function of the scan line. They are extracted from the SBET file and are used to generate the orientation matrix. Other non-image outputs are available by selecting the appropriate checkboxes in the Load SDF Data dialog.

SBET values:

Total records: 574602

record: 1

GPS time: 313034.001350142

Lat & Long: 36.978483 -82.585390 36 58 42.54N 82 35 7.40W

UTM: 358893.51 4093659.92 17

roll, pitch: 0.335330 -3.826155

heading: -140.018631

altitude: 978.621011

velocityEW: -42.725775

velocityNS: 36.939076

velocityVert: 1.160774

speed: 56.479971

record: 57461

GPS time: 313321.300678197

Lat & Long: 36.875784 -82.736870 36 52 32.82N 82 44 12.72W

UTM: 345202.41 4082501.00 17

roll, pitch: 3.122390 -2.501643

heading: -139.684906

altitude: 696.695452

velocityEW: -42.178927

velocityNS: 38.580642

velocityVert: -1.648243

speed: 57.162293

record: 114921

GPS time: 313608.600093058

Lat & Long: 36.852470 -82.742890 36 51 8.89N 82 44 34.40W

UTM: 344618.63 4079924.23 17

roll, pitch: -1.450092 2.559776

heading: 94.492762

altitude: 669.791675

velocityEW: -2.559079

velocityNS: -34.562419

velocityVert: 1.884352


speed: 34.657029

A sampling of the records in the Smoothed Best Estimate Trajectory (SBET) file. A GPS time is embedded in each SDF FWF record. This time is then used to extrapolate into the SBET file's time to determine the position and orientation of the platform for that record.

Raw

The **Raw** image format permits you to specify the dimensions and format of general image files and the resulting size of what will be loaded into an image window. The only difference between a **Raw** format and **Multiband** format (explained later) is that the **Raw** format results in loading a single image plane. Opening a **Raw** image always results in displaying an 8 bit image with the underlying source data in memory. You have considerable control on how this file data is converted to 8 bits for display. The **Raw** image file to be opened is selected by choosing **Open As...** from the **File** menu and then selecting **Raw** from the **Format** popup menu. Since any file can be considered as a **Raw** image, the open dialog displays all file types. You must know which files represent meaningful images.

Any operations that process the image (e.g. [Fourier](#), [Image Arithmetic](#), [Line Profile Plots](#), [Mosaic](#), [Warp](#)) always use the actual underlying data instead of the 8 bit display values. Moving the cursor over a data image displays the data value when the **Info** window is open. Histograms and Line Profile Plot operate on the data. See the **Applications** section for examples of data usage.




The Open As... dialog showing the selection of a Raw format file.

Multiband

A multiband image differs from a raw image in that it consists of 2 or more registered bands each having the same spatial dimensions. **Multiband** images are selected from the same custom open file dialog as **Raw** shown above but with **Multiband** chosen from the **Format** popup menu when you select **Open As...** from the **File** menu. Raw and multiband images are defined by a format dialog (see below) that appears immediately after the **Open** button is clicked. This dialog lists all of the information necessary to correctly read and interpret the image file and to load the image into a window for display. The initial content of the dialog is read from a simple ASCII text header file having the same name as the image file but with a **.hdr** extension (see section [Header Files](#)). If this file doesn't exist (i.e. has yet to be created) then the dialog fields are empty. In either case you can edit the dialog fields and save the proper values to create/update the ***.hdr** file. Subsequent opening of the image or image.hdr file will display the image definition in the format dialog.

Image Format Dialogs

An image format dialog showing the definition of an 16 bit 512 by 512 raw image ready to load the entire image with remapping. The Pixel type is unsigned short and Byte order indicates big endian. The image will be converted to 8 bits for display but will also retain the 16 bit underlying data that will be used in all operations.



The upper portion of the dialog window defines the parameters of the image as it appears on disk. Only the data length is determined by HyperCube, all other information must come from you or the .hdr file. However, HyperCube does check the input for consistency.

Pixels: The number of pixels (columns) per line (row).

Lines: The number of lines (rows).

Header: The number of bytes, if any, to skip at the beginning of the file.

Pixel type: A popup menu specifying the format of each pixel

- 8 bit unsigned
- 16 bit unsigned
- 16 bit signed
- 32 bit floating
- 32 bit signed long

Byte order: A popup that indicates whether the most significant data byte is low order in memory (little endian) or high order (big endian). Only meaningful for 16 bit and larger data. If gray values seem incorrect, try a different order.

- PC (little endian)
- Mac,Sun (big endian)
- Unknown.

The lower portion of the dialog window is where you specify what rectangular subset of the image file is to be loaded into memory. There are also options that determine how the pixels are to be converted to 8 bits for display (the full pixel depth always remains in memory). Pixel conversion is explained more fully in the later multiband example.

Specifically, the requested size fields have the following effect:

- Starting pixel:** The first pixel within a line to load and corresponds to the left edge of the resulting image display window. Pixel numbering begins with 1, not zero.
- Starting line:** The first line of the image file to start loading and corresponds to the top line in the image display window. Line numbering begins at 1.
- # of pixels:** How many pixels per line should be loaded.
- # of lines:** How many lines of the image are to be loaded.
- Auto Remap:** If this check box is selected then the image pixels are automatically converted from their pixel type to 8 bits based upon the distribution of their histogram and upon the percentage popup menu. Generally, this is meaningful only for 16 bit and wider data.
- %:** If **Auto Remap** is checked then this popup specifies what percentage of the distribution of the histogram about the median is mapped to the display interval [0,255]. See the statistical display window shown in the multiband dialog example.
- Save hdr...** This button brings up the standard Mac save file dialog and allows you to save the *.hdr file defining this image. Subsequently, opening this image will refill the dialog fields from the *.hdr file.
- Load hdr...** Displays an Open file dialog so that you can load a possibly different header file. There are no restrictions on the file's name.
- Stats:** When clicked, this button reads the image file and generates various statistics within the rectangular area that you want to load. The results are displayed graphically in a new window. See section [Statistics](#).
- Cancel:** The customary button for aborting the image load.
- Load:** Loads the image automatically computing statistics as necessary.
- Overview:** This button loads a reduced resolution view of the source image. The initial scaling is determined by the display screen size and the **Starting pixel**, **Starting line**, **# of pixels**, and **# of lines**. You can select a different scaling. See section [Overviews](#).

The image format dialog window for a hyperspectral image is shown below. It includes all of the raw image fields plus additional to define band depth and the relationship of the file pixel values to an absolute [0,1] range for comparison with a reflectance library or when classifying one image cube against a different image cube (see section [Classify Function](#)).

The cube image format dialog showing the definition of a 210 band image with 16 bit pixels. Auto remapping of a 95% interval about the median will be used to load bands 60 through 69. The pixel values were previously calibrated to lie between 0 and 1000 will be converted to a [0,1] range by scale.

The screenshot shows the 'Load drill' dialog box with the following settings:

- Data length: 129024000 bytes.
- Pixels: 320
- Lines: 960
- Bands: 210
- Header (bytes): 0
- Interleave: BIL (line seq)
- Byte order: Mac,Su...
- Pixel type: 16 bit unsigned
- Normalization to convert data to [0,1] interval:
 - scale: 0.001
 - offset: 0
- Requested Input Size:
 - Starting pixel: 8
 - # of pixels: 308
 - Starting line: 1
 - # of lines: 960
 - Starting band: 60
 - # of bands: 10
- Display remapping:
 - ☒ Remap bands as a group, %: 95
 - ☐ Remap bands individually, sigma: 2.5
- ☐ Use G & O
- Buttons: Cancel, Save hdr..., Load hdr..., Stats, Load

The additional fields are defined as:

Bands: Total number of bands in the image.

Interleave: A popup menu indicating the type of data arrangement:
Band Line 1 through N for band 1, Line 1 through N for band 2,...
Line Line 1 for all bands, Line 2 for all bands,...
Pixel Pixel 1 for all bands, Pixel 2 for all bands,...

scale & offset: These are the transformation coefficients that convert a file pixel to the range [0,1]. This has nothing to do with display only spectral library matching and classifications with another cube. In the example above the file pixels lie between 0 and 1000 so the transformation is .001 and 0. Also, see **Load Gains & Offsets**.

Starting band: Which band from the image will be the first display band. Band numbering begins at 1.

of bands: How many sequential bands beginning at the starting band should be loaded.

Load hdr... Displays an Open file dialog so that you can load a possibly different header file. There are no restrictions on the file's name.

Remap bands individ: Each band's statistics is used to rescale it for presentation. Choosing this option produces a display cube having a uniform brightness and contrast. This only alters the display and does not affect any subsequent analytical operations. Also, see **Remap bands as a group** which preserves the relative band to band brightness differences.

... sigma: For each band's mean and standard deviation, pixels within $[\text{mean} - \text{sigma} * \text{sd}, \text{mean} + \text{sigma} * \text{sd}]$ are mapped to [0,255].

Remap bands as a group: The composite statistics for all of the bands that are to be loaded are computed and applied globally. The result is a display cube that represents the intensity variation of the image cube file.

Load Gains & Offsets... Allows you to open a text file containing a list of gains and offset pairs that will be applied every time cube file data is accessed and affects all analytical operations such as plots and classifications. The number of gain & offset pairs must equal the number of cube file bands, not just the number of bands loaded. After a gains & offsets has been loaded the the data is internally converted to floating even if the source file data is integer and the gains and offsets were also integers. Therefore, subsequent operations will be a little slower than if no gains and offsets were in effect. The scale and offset fields that are used in **Normalization to convert data to [0,1] interval** are applied after any gains & offsets, so be sure to change this transformation to match the new range. After loading a G & O file the **Remap Bands** check box will be on. This is necessary due to the new floating point range. **Remap** can be checked off but the resultant display image cube will probably be very unsatisfactory. In any case, this will not affect the plots or classifications. See sections [Spectral Calibration](#), [Reformatting a File](#) for more on gains & offsets.

Use G & O: Automatically enabled and checked on after a gains and offsets file is loaded but may be checked off to discard the gains and offsets. To determine later whether a G & O file is being used by a cube open the **Info** window, (**Windows -> Show Info**), and examine the **Kind** field. It will show **Kind: CUBE w/G&O** if one is being used.



The remaining cube file dialog fields are identical to those of the raw format.

Overviews



The File- > Open As menu items: [Multiband](#), [Raw](#) and certain [ERDAS Imagine®](#), [NITF](#) and [TIFF](#) dialogs each contain an [Overview](#) button. Its purpose is to generate an image of reduced scaling that displays the entire or user specified subset of the image at a scale so that scrolling is not needed. A dialog is presented showing the nominal scale factor and the overview band which you may edit.

The Requested Input Size, the bottom portion of the dialog, specifies the area of the image used to generate the overview. The remapping options are used as normal to produce a decent appearing image. When generating a multiband overview, dialog item [Starting band](#) is the default image for the overview.

Once an overview image has been generated you can draw a rectangle (see section [Selection Rectangle](#)) on it that defines the portion of the source image to load by choosing menu File -> Load Selection. Doing so brings up the dialog defining the image along with the Starting pixel, Starting line, # of pixels, and # of lines representing the image subset delineated by the selection rectangle. If no selection rectangle has been drawn on the overview then the fields will be set to the spatial extent of the overview. For multiband images, the original Starting band and # of bands are not changed. You may edit any value before clicking the Load button.





Choosing Overview from the dialog at the upper left generated a 1:2 scale of the image shown above. A selection rectangle was drawn and File->Load Selection was used to display the dialog at the lower left. Note the new values that delineate the image shown below when the Load button was clicked.



Saving Overviews

An overview image serves as a convenient method of automatically extracting the source file and saving it in a possibly different file format. This applies to both a selection rectangle and the entire image. Specifically, NITF, TIFF and Raw image overviews can be saved in any of those 3 file types. Cube image overviews can only be saved in cube image format. Save is always performed on the source image file depth. Any reference coordinates associated with the source image overview are correctly propagated to the saved image file (see sections [Image to Reference Coordinates](#) and [RPC Files](#)). If the overview is color then the output is restricted to NITF and TIFF.

Shown below is a typical dialog that appears after selecting menu File->Save As when saving an overview image. In this case the overview has a selection rectangle defined. Also, see sections: [ERDAS Imagine®](#), [NITF Files](#), [TIFF Files](#) and [Overviews](#) for detail concerning the generation of overviews.



A true color 1:6 overview image with a selection rectangle.

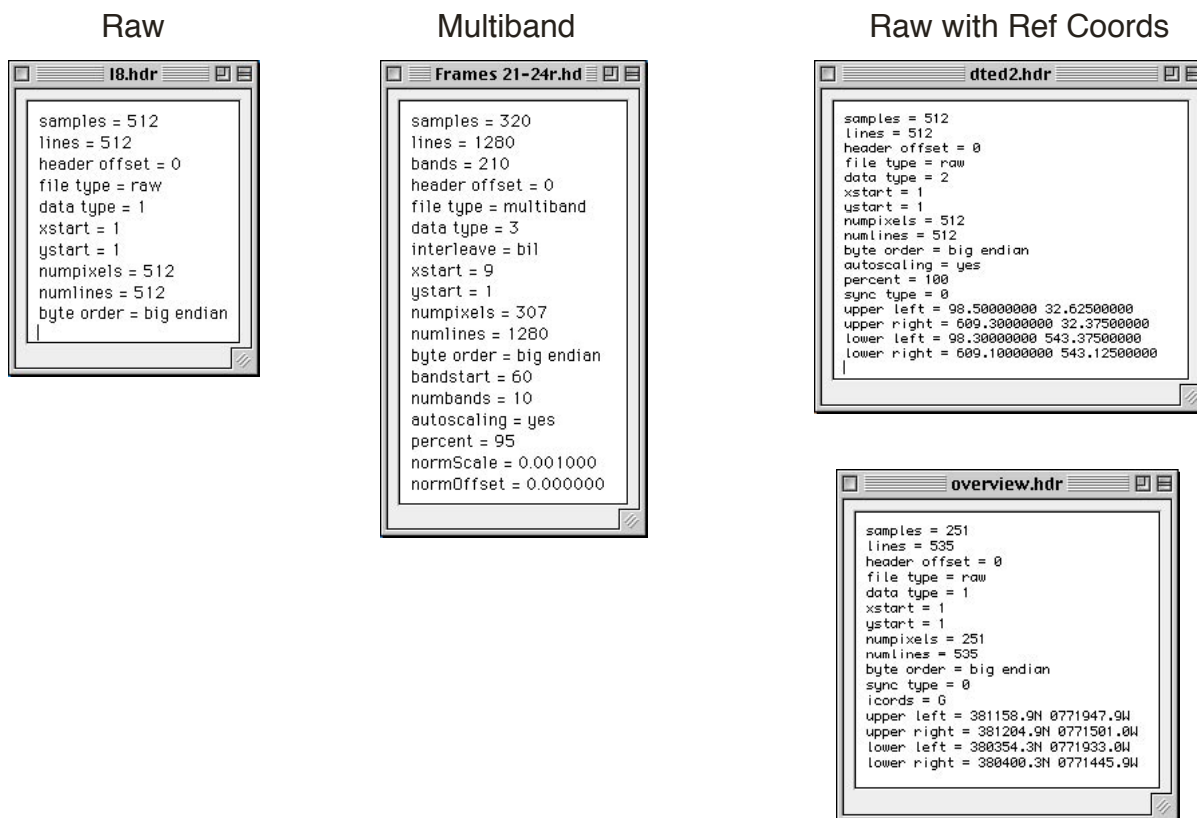


The reloaded image resulting from clicking the Full Res button in the above dialog.

Header Files

As stated previously, each image file can have a unique ASCII text header file that contains all the image characteristics needed to load and display it. This file has the same name as the image but with a .hdr extension and must reside within the same folder as the image file. The two header files shown here (opened as Text files) correspond to raw and multiband examples respectively. They consist of key word assignments, one pair per line in no specific order (see section [Keyword Format](#)). A missing assignment defaults to “no” or zero. Any value can be overridden when brought up in the image format dialog and optionally re-saved. Although header files can be generated using HyperCube (see section [Loading/Saving Text Files](#)) and saved as “Text”, it’s easier to fill in the dialog fields and click the Save hdr... button.

HyperCube can read certain ENVI® header files. Not every format is supported. The program first looks for the beginning ENVI® keyword and if found parses only the information necessary to conform to its own header parameters including wavelengths and bad band definitions. You are alerted if an ENVI® header file is being read. Saving (Save hdr...) information extracted from an ENVI® file writes a new HyperCube file not an ENVI® header file.



Various examples of header files. The last 4 lines of the two right examples list the corner coords of the source image in an external reference system where each keyword specifies an X,Y pair (see next section, Keyword Format). Any combination of 1,2,3 or all 4 keywords may be given. Selecting menu Windows -> [Show Info](#) will display the computed reference coordinate as a function of the window cursor position. Section Image to Reference Coordinates describes how to edit these values.

Keyword Format

Keywords used in the header file are defined below. Each keyword is of the form:
keyword = value. None of the words are case sensitive.

<u>Keyword</u>	<u>Value</u>
samples	number of source image pixels per line
lines	number of source image lines
bands	number of source image bands, if multiband or sync
file type	multiband raw
data type	1 - unsigned byte 2 - unsigned 16 bit short 3 - signed 16 bit short 4 - 32 bit float 5 - 32 bit signed long
interleave	bsq - band sequential bil - line sequential bip - pixel sequential
xstart	starting pixel within the source image to load
ystart	starting line within the source image to load
numpixels	number of pixels to load beginning at xstart
numlines	number of lines to load beginning at ystart
byteorder	little endian - PC big endian - Mac, Sun
bandstart	starting band to load for display purposes
numbands	number of bands to display starting at bandstart
autoscaling	yes - remap image for display purposes
percent	% of remapping statistics to use for display purposes
bandnormalize	yes - remap each displayed band using its own statistics
sigma	number of standard deviations used by bandnormalize

normScale	scale factor to convert pixel values to [0,1] range
normOffset	offset to convert pixel values to [0,1] range
sync type	tc - true color synthetic cube % - classify overlay synthetic cube representing percentage # - classify overlay synthetic cube representing counts d - classify overlay synthetic cube representing thresholds
icords	G - reference coords represent latitude and longitude formatted as ddmms.sZ for latitude (y) where Z is N or S followed by dddmms.sZ for longitude(x) where Z is E or W D - decimal degrees latitude (y) and longitude (x) Note the order: latitude value followed by the longitude value. U - UTM coords. Each reference is a triplet defined as: (Easting, Northing, Zone) in the WGS84 datum. A negative Zone indicates the Southern hemisphere. Also, icords is not required to define a reference. If icords is missing then the reference system is general.
upper left	x,y reference coord pair [triplet] corresponding to pixel, line [1,1]
upper right	x,y pair [triplet] corresponding to pixel, line [# image pixels,1]
lower left	x,y pair [triplet] corresponding to pixel, line [1,# image lines]
lower right	x,y pair [triplet] corresponding to ... [# image pixels,# image lines]
color specific	red, green and blue wavelength values corresponding to those in menu Image -> Cube Color Composite -> Definitions
color intervals	red, green,blue wavelength intervals corresponding to those in menu Image -> Cube Color Composite -> Definitions
color equalize	0 - use color percent value in Cube Color Composite 1 - use histogram equalization for Cube Color Composite
color percent	remapping percentage corresponding to the value in menu Image -> Cube Color Composite -> Definitions keyword not present if color equalize = 1
ann1,...,annN	free text annotation associated with each (1,...,N) band that is displayed in the band scroll area (see Annotation List). Note: a space instead of '=' is used to separate the key word from the text. For example: <u>ann26 Band 26 contains vertical striping</u> Upper case letters are maintained.

Image to Reference Coordinates

You can specify the relation between image coordinates and a reference coordinate system by including keywords such as `upper left`, `upper right`..., described in sections [Header Files](#) and [Keyword Format](#) or at any time an image window is active by choosing menu `Edit -> Options -> Image to Ref Coords`. This displays a dialog showing the current pairings of the image corner coordinates and the corresponding reference coordinates. If no relation has been defined then the default mapping is the identity and the reference values will match the image values. Note that the dialog is completely general and you may specify other than the given corner coordinate pairings, e.g., a known image interior point and the matching reference point. From 1 (simple translation) to 4 (rational poly) pairs of coordinates can be specified. The **Enable** check boxes allow you to exclude any pairings. Saving an image after defining an image to reference relation saves the equivalent corner coordinates to reference mapping in the image's header file. The option for formatted latitude and longitude is specific: `ddmmss.sy` where `y` is either `N` or `S` for latitude and `dddmmss.sz` where `z` is `E` or `W` for longitude. Some general reference values don't correspond to geographic or UTM values and you will get a warning message. Use a negative Zone to indicate the Southern Hemisphere. The **Verify** button checks the validity of the current values against the chosen reference system. **Reset** restores values to those when the Dialog was opened. Radio button **None** will remove any existing image to reference.

Note: Changing the image origin via `Edit -> Options -> [0,0] Image Origin` will disable all **Image to Reference** correspondences in effect and will have to be re-established. Also, some image formats such as NITF and TIFF have reference coordinates in their headers which will be automatically extracted (see sections [NITF Files](#) and [TIFF Files](#)).

Image to Reference

Enter 1 to 4 image and reference coordinate pairs to define a mapping between the image and reference coordinate systems. Only Enabled points are used to compute the mapping and type of transformation (offset, rotation, affine, poly).

Define Reference value units as:

- ☐ None, remove any current mapping
- ☐ General, any decimal values
- ☐ Longitude (x) and Latitude (y), +-ddd.d...
- ☐ Long (x), dddmmss.s...[E,W] and Lat (y), dddmmss.s...[N,S]
- ☒ UTM: Easting Northing +-Zone in WGS84 Datum

Single point reference x and y pixel scales, meters: 1 1

Enable	Image x	y	Reference x	Reference y	Zone
1 <input checked="" type="checkbox"/>	1	1	313986.59	4290293.96	18
2 <input checked="" type="checkbox"/>	1088	1	313470.24	4290223.03	18
3 <input checked="" type="checkbox"/>	1	816	313932.49	4290679.73	18
4 <input checked="" type="checkbox"/>	1088	816	313418.67	4290609.71	18

Cancel Reset Verify Okay

Image to Reference

Enter 1 to 4 image and reference coordinate pairs to define a mapping between the image and reference coordinate systems. Only Enabled points are used to compute the mapping and type of transformation (offset, rotation, affine, poly).

Define Reference value units as:

- ☐ None, remove any current mapping
- ☐ General, any decimal values
- ☐ Longitude (x) and Latitude (y), +-ddd.d...
- ☒ Long (x), dddmmss.s...[E,W] and Lat (y), dddmmss.s...[N,S]
- ☐ UTM: Easting Northing +-Zone in WGS84 Datum

Single point reference x and y pixel scales, meters: 1 1

Enable	Image x	y	Reference x	Reference y	Zone
1 <input checked="" type="checkbox"/>	1	1	0770825.07W	384430.30N	
2 <input checked="" type="checkbox"/>	1088	1	0770846.38W	384427.61N	
3 <input checked="" type="checkbox"/>	1	816	0770827.67W	384442.76N	
4 <input checked="" type="checkbox"/>	1088	816	0770848.89W	384440.10N	

Cancel Reset Verify Okay

The dialog window on the left shows a reference system expressed in UTM coordinates while the dialog to the right shows how the values have automatically changed after selecting the Geographic units as the reference. No changes take effect until the **Okay** button is pressed and all units are within the expected ranges of the chosen reference system. The type of units will be displayed in the [Info Window](#) as a function of cursor position within the image. The `x` and `y` pixel scales are used when there is only one reference pair, typically, the upper left. Either scale may be negative but not zero.

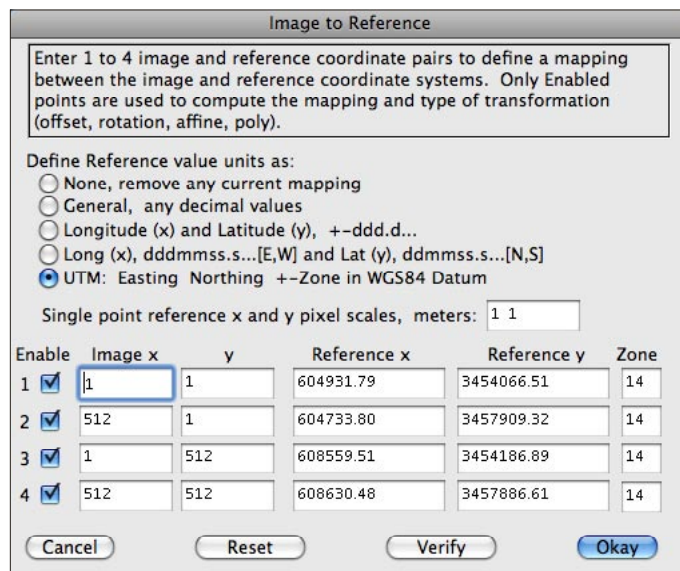

Geo/UTM Reference Grids

Any image that has a geographic reference (either via a geo tif or a corresponding [RPC file](#)) can be displayed with an overlaying Latitude, Longitude or Northing, Easting grid. Menu Edit -> Options -> Geo/UTM Ref Grids... displays the following dialog.

Enable/Disable a grid over all applicable currently open image windows.

Only the current image.


Overlay grid color. Remains in effect until changed.



The [Image to Reference](#) dialog shown on the left displays the coordinate reference that was derived from an [RPC file](#) corresponding to the image show below.




Mini axes show the direction of increasing Northing and Easting.

If the geo-reference was derived from a [RPC file](#) then the direction of increasing height is also indicated by the green line segments in each corner of the image's window.





Create Geographic Rational Polynomial Coefficients

Any image that has a geographic [reference](#) either in Latitude/Longitude or UTM can be modified to contain an equivalent set of RPCs. If the image already has [RPCs](#) a warning message will be presented in order to proceed and overwrite the existing one. Selecting menu **Edit -> Options -> Create Geo RPC...** with the target image active displays the dialogs shown below. If the image is unreferenced then menu **Edit -> Options -> [Image to Ref Coords...](#)** can be used to add one.



Geographic Coordinates <-----> UTM Coordinates

Any image that has a geographic reference in Latitude/Longitude can be transformed to a UTM coordinate system and vice versa (i.e., UTM to Geographics). If the image contains Rational Polynomial Coefficients ([RPC](#)) these will also be converted as shown in the dialogs below. If the image has no reference then menu **Edit -> Options -> [Image to Ref Coords...](#)** can be used to add one.



Rational Polynomial Coefficients (RPC) Files

When an image file is opened the program automatically searches for a file in the same directory having the same base name concatenated with ‘_rpc.txt’, e.g., the file ‘terrain.tif’ and a matching RPC file named ‘terrain_rpc.txt’. In addition, a RPC file can be manually associated with the current active image via menu Edit -> Options -> Attach RPC File.

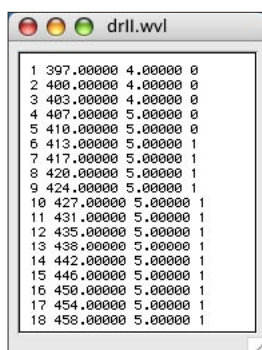
Shown below is one of the RPC text files that was used in section [Stereo Compilation](#) to produce the various outputs. The actual file is one continuous list, not two columns.

```
lat_off: 3.45599170374970e+006
lat_scale: 2.59681012770393e-004
long_off: 6.06685418582732e+005
long_scale: 2.56109265392068e-004
height_off: 1.50000000000000e+002
height_scale: 3.33333333333333e-003
line_off: 0.00000000000000e+000
line_scale: 2.60416666666667e-004
samp_off: 0.00000000000000e+000
samp_scale: 2.60416666666667e-004
line_num_coeff_1: -3.41929000000000e-002
line_num_coeff_2: 1.98476000000000e+000
line_num_coeff_3: 3.25206000000000e-002
line_num_coeff_4: -6.64100000000000e-003
line_num_coeff_5: -1.35979000000000e-001
line_num_coeff_6: 2.28285000000000e-001
line_num_coeff_7: 4.19479000000000e-003
line_num_coeff_8: 7.92820000000000e-002
line_num_coeff_9: -2.20915000000000e-003
line_num_coeff_10: -7.60500000000000e-004
line_num_coeff_11: -3.08305000000000e-002
line_num_coeff_12: 2.38494000000000e-003
line_num_coeff_13: 8.21080000000000e-003
line_num_coeff_14: 2.60906000000000e-002
line_num_coeff_15: -1.05586000000000e-002
line_num_coeff_16: 1.26039000000000e-004
line_num_coeff_17: 4.38989000000000e-004
line_num_coeff_18: 1.79415000000000e-002
line_num_coeff_19: -4.82501000000000e-004
line_num_coeff_20: -3.55988000000000e-005
line_den_coeff_1: 1.0
line_den_coeff_2: 1.87610000000000e-002
line_den_coeff_3: -2.76296000000000e-002
line_den_coeff_4: 1.95251000000000e+000
line_den_coeff_5: 1.05762000000000e-002
line_den_coeff_6: 7.99540000000000e-002
line_den_coeff_7: -2.76620000000000e-003
line_den_coeff_8: 2.24295000000000e-001
line_den_coeff_9: -1.14709000000000e-003
line_den_coeff_10: -1.35064000000000e-001
line_den_coeff_11: 1.21464000000000e-003
line_den_coeff_12: 1.79438000000000e-002
line_den_coeff_13: -3.22646000000000e-005
line_den_coeff_14: -1.05638000000000e-002
line_den_coeff_15: -2.96325000000000e-004
line_den_coeff_16: 2.32113000000000e-003
line_den_coeff_17: 8.35030000000000e-003
line_den_coeff_18: 2.54771000000000e-002
line_den_coeff_19: -2.47739000000000e-004
line_den_coeff_20: -3.06604000000000e-002
samp_num_coeff_1: 1.10369000000000e-004
samp_num_coeff_2: 2.57084000000000e-004
samp_num_coeff_3: -4.35454000000000e-004
samp_num_coeff_4: 5.74180000000000e-004
samp_num_coeff_5: 1.53128000000000e-004
samp_num_coeff_6: -1.08740000000000e-004
samp_num_coeff_7: 1.68296000000000e-004
samp_num_coeff_8: 7.38240000000000e-005
samp_num_coeff_9: -9.20920000000000e-005
samp_num_coeff_10: -3.72540000000000e-005
samp_num_coeff_11: 1.85246000000000e-003
samp_num_coeff_12: -3.19758000000000e-005
samp_num_coeff_13: -5.15168000000000e-004
samp_num_coeff_14: -1.55547000000000e-003
samp_num_coeff_15: 2.61763000000000e-004
samp_num_coeff_16: 2.56452000000000e-004
samp_num_coeff_17: 2.69936000000000e-003
samp_num_coeff_18: -4.41552000000000e-004
samp_num_coeff_19: -1.50709000000000e-003
samp_num_coeff_20: -1.50616000000000e-003
samp_den_coeff_1: 1.0
samp_den_coeff_2: 2.97709000000000e-004
samp_den_coeff_3: -4.25453000000000e-004
samp_den_coeff_4: 6.71890000000000e-004
samp_den_coeff_5: 1.45216000000000e-004
samp_den_coeff_6: -1.72159000000000e-004
samp_den_coeff_7: 1.32578000000000e-004
samp_den_coeff_8: -1.90052000000000e-005
samp_den_coeff_9: -3.41082000000000e-005
samp_den_coeff_10: -8.98320000000000e-005
samp_den_coeff_11: 1.85153000000000e-003
samp_den_coeff_12: -2.63163000000000e-005
samp_den_coeff_13: -5.19085000000000e-004
samp_den_coeff_14: -1.54901000000000e-003
samp_den_coeff_15: 2.57449000000000e-004
samp_den_coeff_16: 2.60572000000000e-004
samp_den_coeff_17: 2.69862000000000e-003
samp_den_coeff_18: -4.31742000000000e-004
samp_den_coeff_19: -1.51713000000000e-003
samp_den_coeff_20: -1.50164000000000e-003
```

Wavelength File

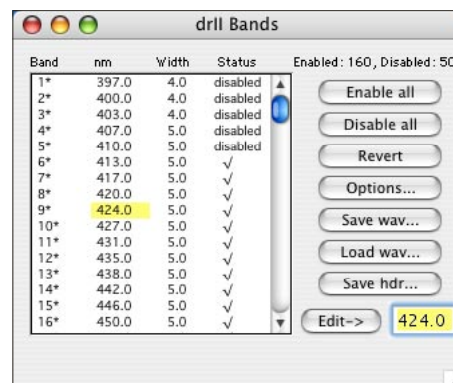
In order to perform analytic operations on a multiband image, it is necessary to know the wavelength associated with each image band. Similar to the *.hdr file there is an ASCII text file that specifies this information. The wavelength file has the same name as the multiband image file but with a .wvl extension (previous versions of HyperCube used wav extension which are still read correctly if the .wvl extension is not found). For an N band image there are N rows of values. HyperCube can parse 3 different formats: a row consisting of just the wavelength, a row consisting of a band number and wavelength and a row consisting of band number, wavelength, band width and status. Each element in a row is separated by a tab or space with the row terminated by a return or new line. The file can be generated using a spreadsheet or word processor and is completely editable in HyperCube. Below is the text *.wvl file for a 210 band image and the way it appears in the program for editing. See section [Band List](#) for a definitions of the various options.

HyperCube can also extract the wavelengths as well as other image parameters from an RSI ENVI® header file (see section [Header Files](#)).



1	397.00000	4.00000	0	
2	400.00000	4.00000	0	
3	403.00000	4.00000	0	
4	407.00000	5.00000	0	
5	410.00000	5.00000	0	
6	413.00000	5.00000	1	
7	417.00000	5.00000	1	
8	420.00000	5.00000	1	
9	424.00000	5.00000	1	
10	427.00000	5.00000	1	
11	431.00000	5.00000	1	
12	435.00000	5.00000	1	
13	438.00000	5.00000	1	
14	442.00000	5.00000	1	
15	446.00000	5.00000	1	
16	450.00000	5.00000	1	
17	454.00000	5.00000	1	
18	458.00000	5.00000	1	

*One of the three possible ASCII *.wvl files defining wavelengths recognized by HyperCube. This format also includes the band widths and status of each band.*




Band	nm	Width	Status
1*	397.0	4.0	disabled
2*	400.0	4.0	disabled
3*	403.0	4.0	disabled
4*	407.0	5.0	disabled
5*	410.0	5.0	disabled
6*	413.0	5.0	✓
7*	417.0	5.0	✓
8*	420.0	5.0	✓
9*	424.0	5.0	✓
10*	427.0	5.0	✓
11*	431.0	5.0	✓
12*	435.0	5.0	✓
13*	438.0	5.0	✓
14*	442.0	5.0	✓
15*	446.0	5.0	✓
16*	450.0	5.0	✓

The HyperCube band list after selecting menu Windows -> Show Band List while cube image drII is active.


If the *.wvl (or the older *.wav) file doesn't exist for the multiband image then band numbers will be substituted for wavelength values and you will get a warning message during the image loading process directing you to menu Windows -> Show Band List where several options are available for creating/modifying one. See section [Band List](#) for more detail.

Statistics

Shown below and on the following page are the results of clicking the **Stats** button when loading a **Raw** or **Multiband** image. Two histograms are generated: a basic data plot that can be resized and magnified and a remapping plot defining how the data is converted for display. Also, see section [Histograms](#).



The histogram of an image with 16 bit unsigned pixels along with various statistics.



The yellow highlight covers the entire area to match the 100% remapping specified in the load dialog. In this case pixel value 883 will be mapped to zero and value 1534 will be mapped to 255 for display.

The numbers shown for Min, Max, and so on are the values for the pixels in the image file within the rectangular area chosen by the user for subsequent loading. In this case they are confined to 16 bit unsigned [0,65535] pixels. In the above example the remapping limits extend over the full range of the input data because the percent default is at 100. Another example using hyperspectral data appears on the following page.


The statistics window is important for visualization because based upon its values the image file pixels are remapped to the 8 bit [0,255] display range. Explicitly, the image file interval:

[pixel_median - %/2, pixel_median + %/2] is mapped to [0,255].


Values < pixel_median - %/2 are set to zero.

Values > pixel_median + %/2 are set to 255.

If the resulting loaded image appears too bland (washed out) then the auto remapping percentage should be made lower using the % edit field. If the image appears too contrasty then the percentage should be made larger.



The complete histogram and statistics of a 210 band hyperspectral image having 16 bit unsigned pixel values.




The yellow highlighting delineates the 95% area that will be remapped to [0,255] for display. Pixel values below 16 will be black and those above 490 will be white.

Below are display snippets of the example multiband data after being loaded with 3 different auto remapping percentages.



100% Interval



95% Interval



70% Interval

Regardless of how the imagery is remapped for display purposes, all analytic operations such as plotting, spectral classification and image arithmetic can operate upon the image file as well as the converted display values. So although saturation and underflow typically occurs when displaying multiband images all of the arithmetic is performed on the full range of the file data.

If Remap Bands Individually is chosen in the image format dialog then each band will be correctly gray mapped but at the expense of maintaining their relative brightnesses. Again, this will not affect any classifications or plotting. If menu Edit -> Options -> Detailed Messages is checked on then two scrollable lists will be output showing the individual statistics and remapping coefficients for each band.

Image Operations

Zoom

Image zooming is selected via menu **Image -> Zoom**. It is automatically enabled whenever the active window is an image window. Each **Zoom In** magnifies an image by 2 until it reaches 16 times (or less for very large images) then it is disabled. Each **Zoom Out** demagnifies by 2 until the magnification is unity. The current zoom factor is displayed in each window's title bar. Zooming is performed by pixel replication and is a "virtual" image in that a saved zoomed image is still un-zoomed. To generate a "real" magnified image use [Magnify](#) in the **Image** menu. Printing or copying a zoomed window does produce a zoomed output. The keyboard equivalents for **Zoom In** and **Zoom Out** are command + and - respectively.

Selection Rectangle



© 1935 Paramount

Many image operations can be constrained to a subset of an image by drawing a selection rectangle within the image window. This is done by positioning the cursor to one corner of the area, pressing the mouse button and dragging to the opposite corner and releasing the mouse button. When completed the area will be delineated by a rotating marquee at the boundary as shown here. A single mouse click (double click in the Windows version) anywhere within the image window or pressing the **delete** key or choosing **Clear** from the **Edit** menu will deselect the rectangle. The rectangle may be repositioned by pressing and holding down the mouse while the cursor is within the selection boundary and dragging to a new position and then releasing the mouse. The cursor changes to a multiple arrow while dragging. Each image window can have at most one selection rectangle. It's possible to unintentionally hide a selection if a window is re-sized. Scrolling a window will bring the selection area back into view.

Copy, Paste & Duplicate

HyperCube allows copying and pasting entire images or selection rectangles between the Mac clipboard and its own image windows as well as other application's windows. Only Pict type data is recognized. This is a standard clipboard data type used by every application. In order to perform a copy draw a selection rectangle about the image area to be copied, then choose **Copy** from the **Edit** menu or the keyboard equivalent, command C. If there is no selection rectangle active or the window contains a plot or list or some other non-image then the entire window is copied to the clipboard. To paste something already in the clipboard, click anywhere within the destination window to activate it and choose **Paste** from the **Edit** menu or its keyboard equivalent, command V. **Paste** will not be enabled unless there is Pict data or color pixel value in the clipboard. The pasted data appears in the center of the viewable area of the front most image window with a selection rectangle around it. It does not paste into an existing selection and if there was one it just disappears. If the target image is too small to hold the paste then the clipboard data is re-sized to fit within the image boundaries. The pasted data can be moved by clicking and holding down the mouse within the selection rectangle and dragging it to a new position. The cursor changes to a hand during the

repositioning. A single click anywhere within the image window embeds the paste. Pressing the delete key or choosing menu **Edit->Clear** erases the selection without pasting it.


If **File -> New...** is chosen and there is something in the clipboard to paste then the default size of the new window will be set equal to the size of the image in the clipboard. Performing a paste will copy the clipboard to the new window.

A copy of the currently active image window or selection rectangle can be made by choosing the **Edit -> Duplicate** menu. This option is available for almost every type of window. The duplicated image will have the same name as the original followed by the word "Copy". Some HyperCube operations, such as constructing a cube, automatically make image duplicates. Also, see section [Saving and Copying Scrollable List](#).

Any time that the cursor is over an image window, regardless if the window is front most, the current pixel value may be copied to the clipboard by also pressing the option key, e.g. command+option+C combination (**Ctrl+Alt+C** in [Windows](#) version). This value may subsequently be used as the selected color in many HyperCube functions, (see sections [Image to Mask](#), [Adjacency](#), [Selection Point Colors](#) and [Class Map Editor](#)).

Gray Mapping

The contrast and brightness of the active gray scale or color image window may be interactively changed by selecting menu **Image -> Gray Map**. This opens a window containing controls to modify the pixels in the image and a graphical plot of the current transformation and resulting histogram. The **GrayMap** window applies only to the image window that




The effect of modifying the brightness within a selection rectangle.

was active (front most) when **Gray Map...** was selected. However, every image window can have its own gray mapping operator. For **Linear** mapping the **Contrast** and **Brightness** sliders control the slope and y intercept of the graphical plot respectively. The horizontal plot axis represents the input pixels [0,255] and the vertical axis represents the remapped pixel values, also in the [0,255] range. Initially, the slope is 1 and the intercept is 0 signifying that the input is identically mapped to the output. The contrast range is [-9,9] and the brightness range is [-56,56]. Negative slopes produce complementary pixels. Above is an example of gray mapping applied to just a selection rectangle within an image. If there is no selection rectangle then the entire image is gray mapped. Clicking the **Apply** button embeds the gray mapping into the selection area, or the entire image if no selection.

There is one exception: an 8 bit color indexed image is remapped entirely regardless of any selection rectangle. If a **GrayMap** window becomes hidden behind another window you can bring it to the front by selecting its name from the **Windows** menu. The original image and the identity gray mapping can be recovered by clicking the **Reset** button. A histogram dynamically displays the entire image's histogram as a function of the current transformation starting with the unmapped histogram of the image (also, see section [Data Mapping](#)).

Rubberband gray mapping doesn't use the slider controls. It allows you to position the cursor within the plot area while holding down the mouse button. The graph then tracks the cursor and consists of two slopes with the cursor at their break point. Rubberband mapping is very useful in bringing out shadow detail without completely washing-out the surrounding area. An example of this is shown below.




Rubberband gray mapping being applied to a selection area. The sliders are not used and are grayed out.

Color Mapping

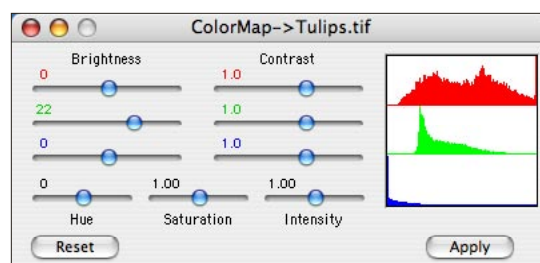
You can modify the color balance of an image by selecting menu Image -> Color Map. It is similar to gray level mapping ([Gray Map](#)). Sliders are used to independently modify the red, green and blue brightness and contrast of an image and its hue (color tint), saturation (purity) and intensity (overall brightness) while showing the effect on the 3 histograms. Note that subtracting a color is the same as adding its complement and vice versa: adding is subtracting the complement. The controls and buttons perform the same as gray mapping with the same exception: **Apply** modifies an entire 8 bit indexed color image regardless of any selection rectangle. The **Reset** button restores the original image and sets the controls to their null positions.

An image may have only one operator attached to it at a time. Specifically, [Color Map](#), [Gray Map](#), [Data Map](#), [Static](#) and [Dynamic](#) filtering, [Fourier](#) and the [Class Map Editor](#) are all mutually exclusive


Also, see section [Info Window](#) which describes displaying the red, green, blue, hue, saturation and intensity values associated with the cursor position.



Color balance emphasizing green within a selection rectangle.



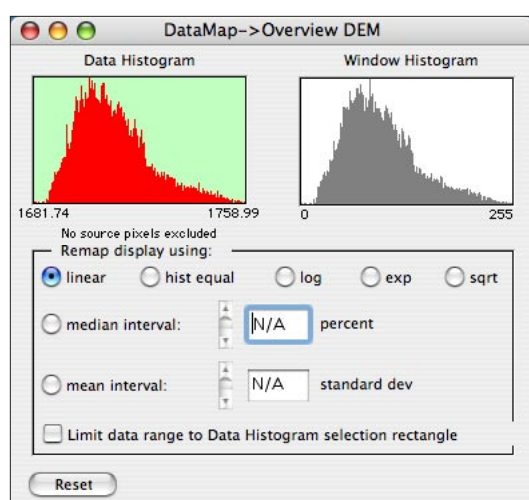
Effect of modifying the hue, saturation and intensity within a selection rectangle.



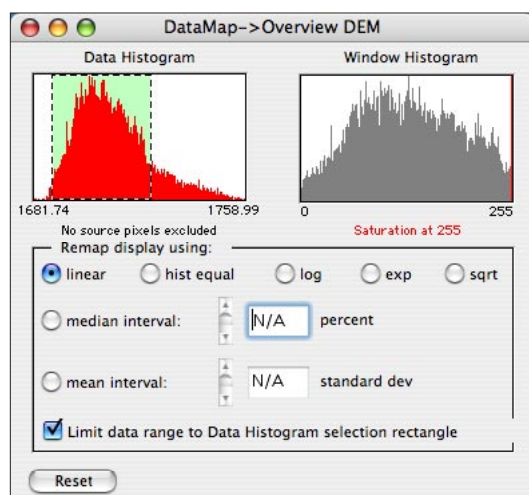

Data Mapping

Windows containing images that represent more than 8 bit pixels are automatically mapped from their full pixel depth (e.g., 16 bit signed and unsigned, 32 bit float and signed long) to 8 bits for display when they are loaded or created (the underlying data is kept in memory at full pixel depth, see section [Files, Data, Images and Display](#)). Menu item **Image -> Data Mapping...** allows you to interactively remap the source data (in memory) to 8 bits. It does not change the actual underlying data, only the way it is displayed. This menu item is only enabled for gray scale images. It does apply any pixel exclusion that may have been set via menu [Edit -> Options -> Pixel Exclusion](#).


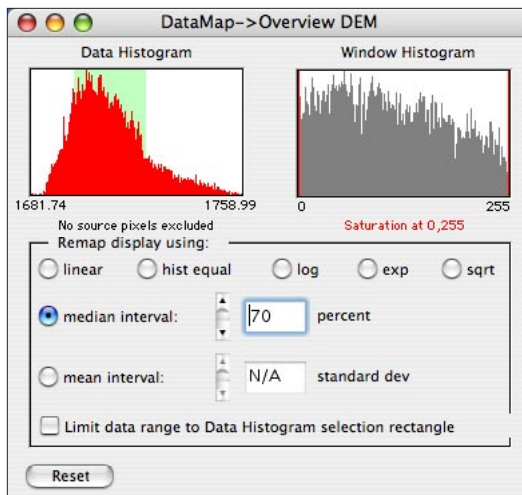
If the source image contains a [selection rectangle](#) then the remapping is restricted to that spatial region. A selection rectangle can also be drawn within the **Data Histogram** area of the dialog to limit the mapping to a range of pixel values, see examples below.




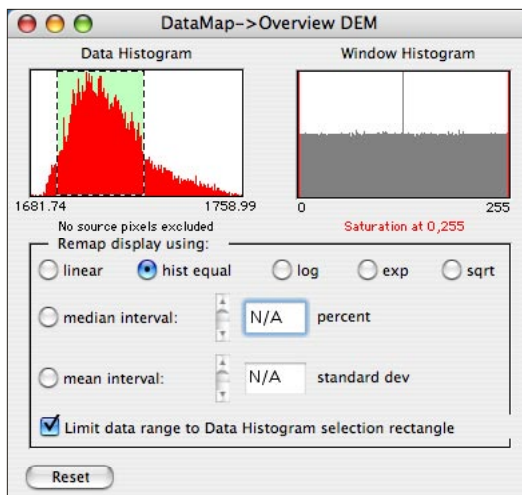
Linear remapping using full histogram.



The linear remap applied to an image selection in combination with a selection drawn over the Data Histogram dialog that further limits the linear remapping to pixel values within the green shaded histogram area. The red vertical line in the Window Histogram plot indicates the resulting saturated values (pure white) in the displayed image.




Same image area as the previous examples using median interval. The median of the pixel values within the image selection is computed and an interval about the median from 35% below to 35% above it forms the remapping interval. Pixels below the interval are set to zero and those above the interval are set to 255. This causes saturation at both ends of the displayed image and the resulting pair of red lines in the Window Histogram plot.



As above with histogram equalization remapping along with a Data Histogram selection. This mapping also stretches out the pixel range and produces saturated display values.


Flips & Rotations

Each loaded image can be flipped and rotated in orientation. As shown below there are 3 kinds of flips and 3 rotations. The **Left->Right** flip changes the order of the pixels within a line and the **Top -> Bottom** flip changes the order of the lines. **Band Order** is only enabled for multiband images and re-sequences loaded bands but not their wavelengths. Therefore, the last band becomes the first band and the first band becomes the last band. Flips and rotations are global operations and ignore any selection rectangles.



Flip, Rotate and Shift appear in the Image menu.



The **Rotate** menu item is similar to **Flip**. Rotations may be 90, 180 or an arbitrary number of degrees. 90 and 180 operate on an entire image (including multiband) while **General** can rotate a selection rectangle as well as the whole image and uses the currently selected interpolation method (see section [Magnify](#)). Note that rotating a non-square image results in interchanging the number of pixels and lines. Also, flips and rotations disassociate the new image geometry from the original image file geometry. Therefore, library matches and classifications of multiband images will be erroneous as well as spectral plots of their files.



The **Reformat Cube File** function in the **Utilities** menu is used to change the internal structure of the multiband image file itself not just what is loaded in memory. It also regenerates the *.[hdr](#) and *.[wav](#) files as necessary.

Shifts

The Image -> Shift menu item operates on whole images and allows both integral and fractional shifts. All of the integral shifts are cyclic so that pixels shifted off one end appear on the opposite end so they are invertible. The fractional part of a shift is not invertible. It is done in the frequency domain by applying a linear phase shift to the [Fourier](#) transform of the line or column. Therefore, a 1.6 pixel left shift followed by a 1.6 pixel right shift will not exactly reproduce the original image. **Left** and **Right** shifts can be applied to **All** lines or just the **Even** lines. They are useful in correcting interleaving errors in captured video images. Below is an example of such an image before and after shifting only the even lines. Shifts applied to a cube only affect the currently displayed band. **Overwrite** is a toggled flag. If checked (on) then the shift overwrites the source image or cube band, else the result of the shift generates a new separate image. A warning is given when overwriting.



Magnified captured video frame before and after Shift Right Even lines.


Filtering (Spatial Domain)

HyperCube provides several methods for filtering an image in both the spatial (pixel) domain and the frequency ([Fourier](#)) domain. The later is described in the [Image Functions](#) section. Spatial filtering convolves the image pixels with various user selected kernels and the software provides two methods: a 5x5 interactive kernel and a 5x5 static kernel. The interactive filter contains a slider for altering the kernel elements while the static filter uses edit fields and operates as a one-shot. You can save the 5x5 static coefficients and read saved 5x5 coefficients including those saved by Photoshop™.

Each type of convolver is an operator to the specific image window that was active when menu **Image -> Filter** was selected. Closing the image also closes the filter. If a selection rectangle is present then the filter is constrained to that area. The filtering does not become permanent until the **Apply** button is clicked. However, the original image can still be re-stored with the **Reset** button as long as the filter is still active.

Filtering is an image operator and an image may have at most one operator at a time, e.g. if [Gray Map](#) is active on an image then it must be closed before any filtering can be applied. Both the dynamic and static filters respond to [Pixel Exclusion](#) when it is specifically set to exclude zero.


Below is an example using the **Static** filter. The **Blur <--> Sharp** slider generates a continuous low pass Gaussian through high pass Laplacian filter effect as it is moved from the far left to the far right.



© 1935 Paramount

Blurred area

Sharpened area.




The static convolver displaying the sharpening kernel that was applied to the right portion of the example image. The left portion is the result of a previously applied Blur. Each region was done by using a selection rectangle followed by the Apply button.

The 5x5 static convolver, shown below, allows 25 separate coefficients. All of the coefficients must be integers. The **Scale** value is divided into the coefficients and the **Offset** value is added to the final resultant pixel. The popup menu contains several predefined kernels which may be modified. The **Preview** button shows the effect of the filtering without embedding the result into the image. However, just like the 5x5 dynamic convolver, the **Reset** button restores the image and sets the kernel to the identity as long as it's still active. Since there are so many degrees of freedom it's easy to generate meaningless pixels and experimentation helps. Any set of coefficients may be saved via the **Save...** button and retrieved with the **Load...** button. The coefficients are compatible with those produced by Photoshop™ and HyperCube will recognize the later. A newly loaded or saved kernel will occupy the **User** menu position. The **Sobel XY Gradient** operation actually involves two separate convolutions, an x gradient and a y gradient that are combined. Therefore, when the combined xy Sobel is selected the kernel elements shown are not representative although the combined gradient is correctly executed. Also, the **Sobel 3** operations produce a finer result whereas the **Sobel 5** operations average over additional pixels. Experiment to find the desired result.



The image displays the result of apply the Sobel 3 XY Gradient operation upon a selection rectangle. Each of the red, green, and blue components is done separately and recombined as a true color image.




- Identity
- User
- Laplacian 1
- Laplacian 2
- Laplacian 3
- Gaussian 1
- Gaussian 2
- Gaussian 3
- Sobel 5 X Gradient
- ✓ Sobel 5 Y Gradient
- Sobel 5 XY Gradient**
- Sobel 3 X Gradient
- Sobel 3 Y Gradient
- Sobel 3 XY Gradient


Magnify

A entire image or a selection rectangle subset can be magnified producing a new independent image. The magnification range is mostly unlimited as long as the resultant image is at least 32 by 32 in size and there exists enough RAM. The type of interpolation can be replication, bilinear or cubic resampling. Menu Image -> Magnify toggles between Image and [Selection](#) depending on whether or not a the image has a selection rectangle. A submenu, shown below, lists several magnifications and permits entering an arbitrary factor that is

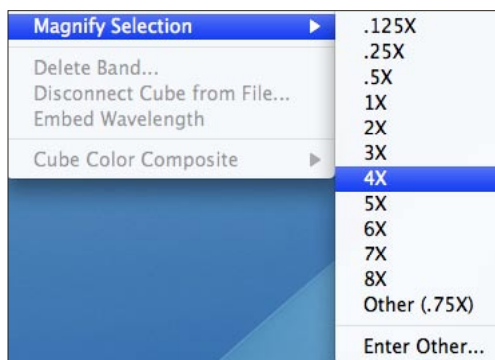
initially set to .75x. Menu, Edit -> Options -> Mag/Rotate Interp determines the type of interpolation. It is initially set to Replicate. Included below are examples of a 4x magnification using each type of interpolation. Replication is best for viewing the actual pixels while bilinear and cubic produce photographic like enlargements. Also, see [Warp](#).



A 4x magnification applied to an image subset utilizing all three interpolators. The original is shown at 2x for clarity.



Original selection



Function Magnify popup submenu.

Replicate



Bilinear




Cubic




Info Window

HyperCube contains one floating window, **Info**, that is always available and dynamically displays information about any of the other HyperCube windows. The kinds of information depend upon which window the cursor is in. When the program starts, **Info** is hidden and can be shown at anytime by selecting menu **Windows -> Show Info**. To hide it click its close box. To print it, use the keyboard equivalent command **P** instead of the mouse so that the cursor doesn't move off its current position and change the **Info** values.


The information displayed include the type of window, and the position and current gray or color values of the pixel beneath the cursor as well as image [reference](#) coordinates in various systems. The origin of all windows is [1,1] at the upper left unless [0,0] Image Origin is selected from menu **Edit -> Options**. And when placed on a histogram plot the cursor shows the bin number. Placing the cursor on a contour line (see [Applications](#)) displays the [contour](#) value. Various **Info** window examples appear below.




Cursor over a [spectral](#) plot window.




Cursor over a [profile](#) plot showing the corresponding x,y,z image values for a point on the profile line.




The cursor is positioned over a [flicker](#) window. Values change according to the flicker rate. Image 1 has been offset using the arrow keys.



Cursor over a gray scale image. [Data](#) value matches pixel value.



A true color image pixel with the red, green and blue color components.


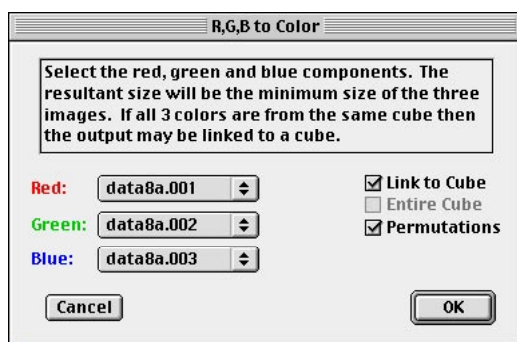



Over an [NITE](#) image with reference coords.

Image Conversions

Components to Color

A composite true color image can be formed by selecting Image -> Convert Image -> R,G,B to Color or H,S,I to Color (Hue, Saturation and Intensity to color). Any three gray scale images can be used, including bands of a multiband cube image, HyperCube displays a dialog window asking you to specify which image represents which of the three colors. The three components do not have to be the same size, the true color result will be their minimum size. If all three images are bands from the same image cube the Link to Cube check box is enabled and if selected the resultant color image will be logically connected to the parent cube image. This image can then be used to select points for subsequent spectral plotting and classification (see section [Classify](#)). If red, green and blue are all cube bands (not necessarily from the same cube) then Entire Cube may be selected and a synthetic color cube will be automatically generated in which each band consists of the corresponding components (bands) from the cubes. The number of bands in this cube will be the minimum of the number of bands currently loaded from the source cubes.



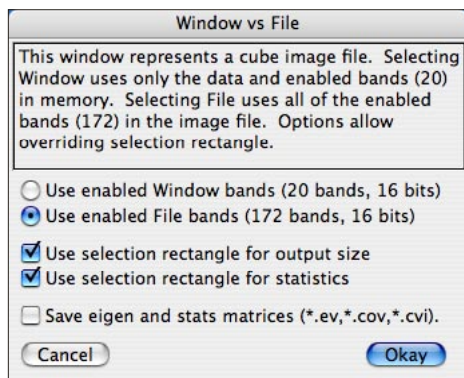
The left dialog shows the selection of three bands to be composited to form a true color image. The right dialog similarly forms a true color image treating the components as if they represented hue, saturation and intensity (also see [Applications->Stereo Mates](#)). Checking Permutations produces 6 color output images.

Image to Various Components





These are inverses of the previous functions. Any true color image window, i.e. those that show a depth of 24 or 32 in the Info window, can be separated into 3 same size gray scale images representing the red, green and blue or hue, saturation and intensity, components respectively. The original image is not affected. To perform the conversion activate the color image window and choose Image -> Convert Image -> Color to R,G,B or Color to H,S,I. Also, see [Attach/Remove Data](#) to reduce the color depth of an image.

Image to Principal Components


Image -> Convert Image -> Principal Components uses the Hotelling transform to convert a color or cube image into its principal components. Either the covariance or correlation matrix statistics can be used as the basis for the transform (and optionally saved). If the source is a cube then you will be presented with the following dialog allowing you to transform either only the [enabled](#) bands that are currently loaded into memory or the actual cube file enabled bands (if a file currently exists). The result of the transformation is an image cube. Selecting “File” produces a file as the output; selecting “Window” generates an in-memory image cube.



Selecting Window transforms only those enabled bands presently in memory and produces a cube in memory. Selecting File transforms the enabled bands of the image cube file and produces a new cube file. The new cube bands are the principal components and are ordered in decreasing eigenvalues. The Save... check box writes the various matrices to an ASCII file. See section [Apply Transformation](#) for format details. The two Use selection... options allow ignoring any selection rectangle.



Eigenvalues	
1	+7.406526e+05
2	+1.713476e+05
3	+1.710998e+04
4	+4.762276e+03
5	+1.898845e+03
6	+7.998950e+02
7	+6.457251e+02
8	+5.176797e+02
9	+3.355326e+02



Shown at the top are the first 3 principal components generated from 166 enabled bands of a 210 band image. The plot and listing of the eigenvalues are also output as well as an image representation of the eigenvectors. The later may be examined via a line profile plot, histogram or using the cursor with the Info window (see section [Apply Transformation](#) for format details).

Indexed to Color

An indexed, or pseudo, image is an 8 bit image similar to gray scale but one in which the pixel value is an index into a color table. Therefore it has at most 256 distinct colors. A gray scale image is a special case of an indexed image in that its color table has the same red, green and blue intensity for each of its indices and is the actual pixel value (HyperCube checks all indexed images to see if they are actually gray because this greatly simplifies a lot of arithmetic). If the currently active window contains an indexed image then you can select **Image -> Convert Image -> Indexed to Color**. This will generate a new true color image. If the original image was gray then the new image will also be gray with equal red, green and blue components at each pixel. Note that on the Macintosh a full true color image is 32 bits in depth, not 24 as might be expected. The extra byte is for special effects. HyperCube also handles the Macintosh 16 bit (thousands) color image: three 5 bit color components plus a special effect bit.

Special note, Mac OS X only: Eight bit pseudo color Pict images are converted to true color when loaded. This does not apply to Pict 8 bit gray images or any other file types.

Color to Indexed

This is the most difficult type of color image conversion because there is no one to one transformation. A typical 512 by 512 true color image could have over a quarter million colors which must be reduced to only 256 colors. HyperCube tries to find the best 256 but sometimes the results are not acceptable. Normally, true color to index should only be used when the source image has very few distinct colors or to save storage as in making a true color image fit on a floppy disk. Its necessity has become less valid with JPEG compression (see section [JPEG Files](#)). You can select the **Image -> Convert Image -> Color to Indexed** menu when the front most window contains a true color image.

Gray Scale/Color Tables


The active color image window can be converted to a new 8 or 16 bit gray scale image by selecting the **Image -> Convert Image -> Gray -> Equal/NTSC** submenus. The three color components are combined together with either equal or NTSC (National Television Standards Committee) weights to form a new image defined by:

$$\text{gray equal} = .33 * \text{red} + .33 * \text{green} + .33 * \text{blue}$$

$$\text{gray NTSC} = .30 * \text{red} + .59 * \text{green} + .11 * \text{blue}$$

Menu [Functions -> Arithmetic](#) can be used with arbitrary weighting.

Menu **Image -> Convert -> Color Table** shows several color mappings that may be applied to an 8 bit indexed image. The top one overwrites the image's color table with a linear gray ramp forcing an indexed color image into some gray appearing image. The second strip complements the existing table so that selecting it twice restores the original table. The remaining 3 menu options overwrite the existing color table and color code the image pixels so that pixel (index) zero represents the left most color and pixel 255 the right most color. Also see the **Classify** section describing output products.



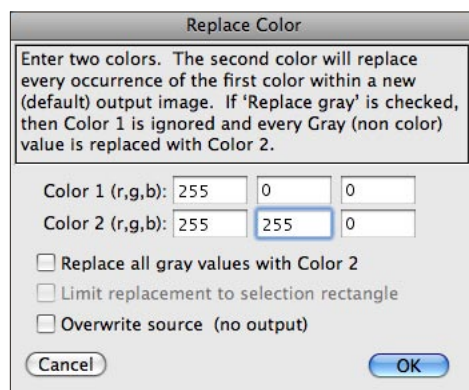
Replace Color

Any specific color or gray value within an image may be substituted with another value. The result is a same sized image with only that modification. If the image is a cube then a copy of the band face is used as the source. Choosing **Image -> Convert Image -> Replace Color...** brings up the following dialog with the last two colors or gray values, if any, that have been first selected by the **command+option+C** key combination or **Ctrl+Alt+C** in **Windows** version (also see section [Copy, Paste & Duplicate](#)). The first selected value will be **Color 1**, the color to search for, and the second as **Color 2**, the replacement color.


If **Replace all gray...** is checked then **Color 1** is ignored and every gray value (non-color) is replaced with **Color 2**. This is useful for converting a color overlay to a mask.

The **Limit replacement to selection rectangle** is an available option when the source image has a selection rectangle. If checked then changes are confined to the rectangle.


Overwrite source modifies the source image, it does not generate a new output.



The Replace Color dialog was used twice to convert the image at below left, a classification mask, to the image at its right. The dialog shows the first conversion, red to yellow. A second conversion using the output from the first replaced the black background with gray [128,128,128]. The intermediate yellow and black image is not shown. Replace all gray... can be used to convert a color overlay to a color mask. See section [Image Products](#).



Original spectral classification mask as output from Classify.

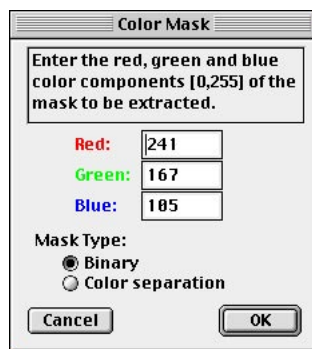


After two successive uses of Replace Color.


Image to Mask

An image can be converted to another image that corresponds to a specific gray level or color within the image. Two types of output are possible. The first (**Binary**) produces a same size image (mask) that consists of white (255) everywhere that specific color (gray level) appeared in the source image and black (0) everywhere else. The second type (**Color separation**) consists of the chosen color in place of white and black everywhere else, i.e. only the color selection itself remains. This is most useful where the source image does not contain numerous colors (gray levels) such as when separating a specific classification color from a color composite produced with the **Classify** function. The initial color (gray level) to be extracted is automatically preset to the last command-option C key combination used to copy a pixel color to the clipboard (see section [Copy, Paste & Duplicate](#)).


When an image is frontmost, choose Image -> Convert Image -> Image to Mask. This will bring up the following dialog window:




The respective red, green and blue components of the color to be extracted in generating the mask images in the examples below. Mask Type Binary produces white wherever the rgb pixel value is [241,167,185] while Color Separation outputs the actual pixel value. If the image to be converted is a gray level image then the dialog will instead solicit a single gray level with the same options.



A source true color image that was derived from a spectral classification. The cursor was positioned over a road and the command-option-C key combination pressed to store the pixel color into the clipboard prior to choosing Image->Convert->Image to Mask... from the menu.




Binary Mask Type consisting only of black (0) and white (255) values.



Color separation Mask Type consisting of black and the extracted color.

Histogram Equalization

Three functions within menu Image -> Convert Image -> Histogram Equalize are used to perform histogram equalizations of an image. In each case a new image is generated leaving the source image intact. Any gray or true color image may be the source. Each function attempts to construct an image whose histogram matches a particular distribution. Generally, this can not be done exactly. For example, a non-uniform 256 bin histogram (8 bit gray or a color component) can't be mapped to a uniform 256 bin histogram, there will be "holes". The methods used by the three techniques are as follows: **Uniform** attempts to remap the pixel values to produce an equal number of pixels in each of the histogram bins. If the source image is 8 bits/color then the red, green and blue pixels are treated as a truncated 18 bit pixel (the low order 2 bits of each color is ignored) whose histogram is used to remap the 24 bits uniformly. The second method, **Uniform (ea. color)**, performs the uniform equalization independently on each color (8 bit red, green and blue components) and then combines the results. The last method, **Match Image...**, is described on the following page. Note, if the source image has a selection rectangle then the equalization is based only on the selection rectangle statistics but is applied to the entire output image.



Shown below are examples using the first two methods.



Original true color image.






Resulting image after Uniform histogram.




Result after Uniform (ea. color).



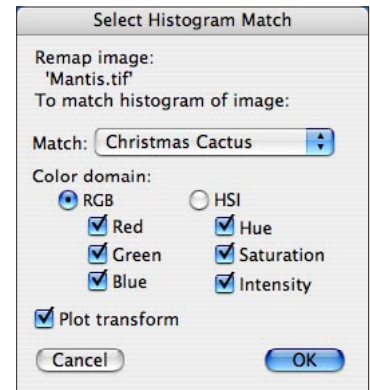
Left and right are the original and histogram equalized image along with their corresponding 8 bit histograms. Note that the resulting histogram cannot be made uniform.



The third method, Histogram Equalize -> Match Image, uses the histogram of a second image as the target histogram distribution, i.e., each color component of the source image is remapped so that its histogram matches the corresponding component histogram of the Match Image in the selected Color domain.




Shown at the right is the corresponding Dialog. The current active (front most) image is the Remap image to be equalized. Both the Color domain and the components to remap are selectable.



The red remapping table generated in Christmas Cactus -> Mantis with Plot transform checked.




Above and to the right are two original images. Each image has been alternately histogram matched to the other in the RGB Color domain with the resulting images shown below.



Cube Color Linked Image

The current, if any, linked color composite image can be switched to another existing color image as long as it is the same spatial size by selecting **Image -> Convert Image -> New Cube Link**. The following dialog then appears listing candidate images. Any previous link will be disconnected and any current selection points will be transferred to the new link.





Color to Classmap

A true color image can be transformed into a classmap in which the first 1024 colors (gray shades are excluded) are assigned classes 1 thru 1024 by selecting menu **Image -> Convert Image -> Color to Classmap**. The classmap can then be modified using all of the features in menu **Utilities -> Classmap Editor** (see section [Class Map Editor](#)). If the source image is pseudo color you will be prompted to first convert it to a true color image. In this case, there will be at most 256 classes. Since a typical color image usually contains far more than 1024 colors this conversion is most useful after masking and/or image arithmetic functions to reduce the number of colors.

Attach/Remove Data

Two disparate types of images can be combined. For example: a classmap (color coded features) and a matching elevation image (data values). This is accomplished by selecting the data image window to make it active and then choosing menu **Image -> Attach -> Data**. This brings up a dialog listing all of the candidate images (same spatial size). After data has been attached to an image the [Info](#) window will show the image color/gray values and the corresponding data value as the cursor is moved over the image (see section [Info Window](#)). When saving a merged image you are presented with a dialog stating that the source image has attached data and asking you if you want the data saved as well. If your answer is affirmative then the data is saved separately in **Raw** format (see section [Raw](#)) so that it and the source image can be reloaded and merged together again.

An image, including an entire cube, can be stripped of attached data by menu **Image -> Attach -> Remove Data**. This method can also be used to convert a 16 bits/color image to an 8 bits/color image (i.e., 48 bit color to 24 bit color)



The Attach Data dialog is shown to the left and the Remove Data warning to the right.






Image Functions

Statistical Measures

Functions -> Statistics contains two methods for computing the linear correlation between 2 or more gray scale images. The first of these, **Pair Correlations...**, allows you to select any two non-color images from popup menus and then calculates their correlation or covariance as a function of pixel and line offsets. The results are listed in a 7 by 7 matrix showing the computed values for each given offset. The **Options...** button permits changing the initial offset between the 2 images and the pixel and line shifts per iteration. If image 1 has a selection rectangle then the computation can be confined to the selection by checking the **Use Image 1 Selection Box**. The check box will only be enabled if there is a selection rectangle on image 1. **Pairs Correlations** can be used to determine how good the registration is between various bands of a multiband image (use menu [Image -> Convert -> Color -> R,G,B](#) to separate color components followed by **Pairs Correlations**). The **Image Shift** checkbox is used to compute where the maximum correlation occurs and via interpolation, what the fractional shift is. Each computation is generated by selecting the **Update**


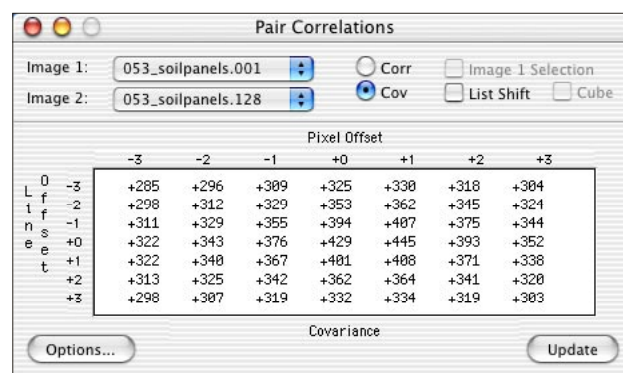
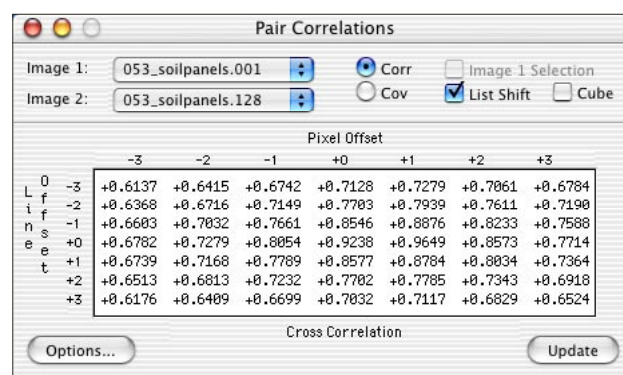


The linear cross correlation between two bands of a multiband data set using the default values for offset and delta is shown at the right. The corresponding covariance between the two bands is shown below it.

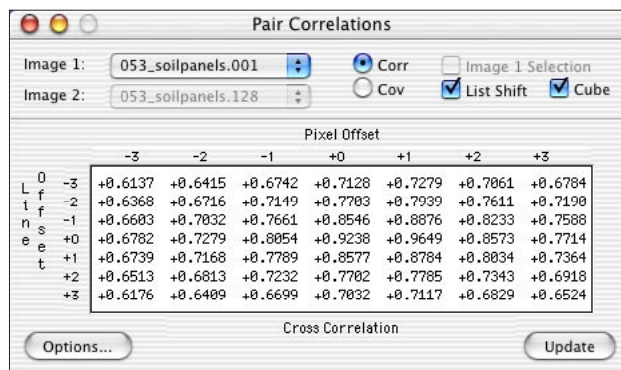


The computed pixel shift (shown above) between the images after checking Image Shift. This is in accordance with the position of where the maximum correlation value appears (i.e., between 0 and +1 pixel offset to compensate for the negative shift value).


The options dialog window specifies the initial pixel and line offset as well as their respective steps for each iteration.




If the Image 1 popup menu displays a band from a cube image then you can generate all of the relative shifts between that band and every band within the same cube. This is done by selecting the List Shift and Cube checkboxes which will ignore the Image 2 popup menu. Shown below is the Dialog and the 3 resulting outputs after the Update button is pressed.



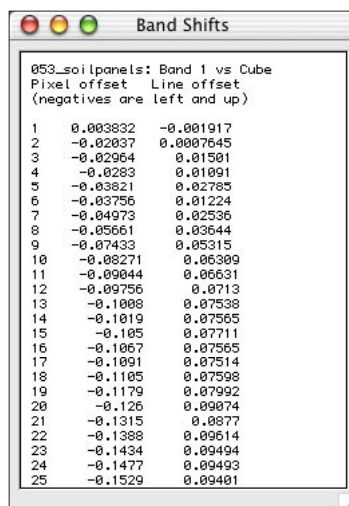
A band from an image cube is selected in Image 1 along with List Shift and Cube to produce the outputs shown below.



The relative pixel (horizontal) shift between band 1 and each band.

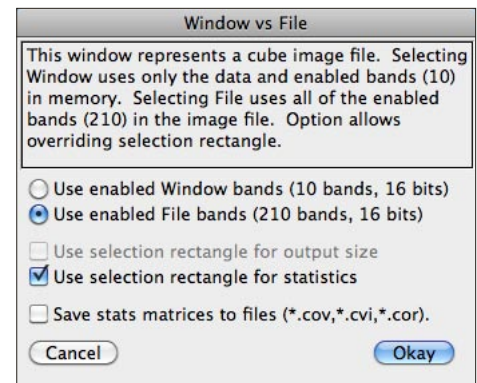


The relative line (vertical) shift between band 1 and each band.



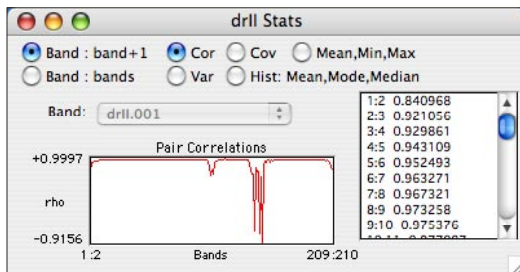
A partial list of each band's offset relative to band 1.

Menu Functions -> Statistics -> Band Correlations applies to multiband images. It calculates various combinations of linear correlation coefficient, covariance, variance, mean and histogram. The dialog at the right gives you the choice of using either the window image data already in memory or the file image data (if a file currently exists). The result is shown graphically with a scrollable list window. Clicking on an item in the list highlights it on the graph. The radio buttons determine the type of comparison and the statistics.

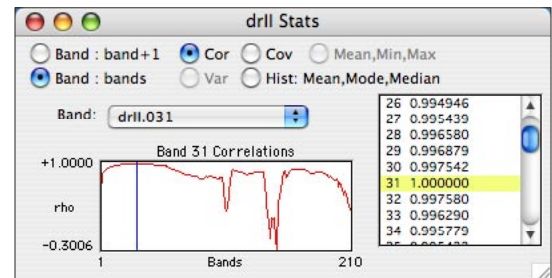


Sequential correlates each band with the next band, i.e. band 1 to band 2, band 2 to band 3 and so on. Specific allows you to select any band from the Band popup menu and compare it with all of the others. Shown below are the results of correlation applied to a 210 band image. The vertical marker in each graph corresponds to the highlighted values in the lists. The list values may be saved as text by choosing File Menu -> Save when the dialog window is front most or copied, (Edit -> Copy), to the clipboard and subsequently pasted as columnar data into a spreadsheet application such as Excel™. The correlation and covariance matrices may be saved to a file in ASCII format which can subsequently be opened in Excel™ (see [Apply Transformation](#) for details).

If menu Edit -> Options -> Detailed Messages is set then an image will be created where each pixel is the correlation [-1,1] of each cube vector to its neighbor vector.




The sequential band to band correlation plot showing the absence of correlation within the water absorption portion of the spectrum. Generated from file data.





Same multiband data with band 31 compared with all of the bands within the cube including itself. Clicking a list item highlights its position within the graph.

If 32 or more bands are used then both the covariance and correlation matrices generated in the statistics are converted into (savable) images that can be further analyzed by performing profile plots, histograms and examining values via the cursor and Info window.

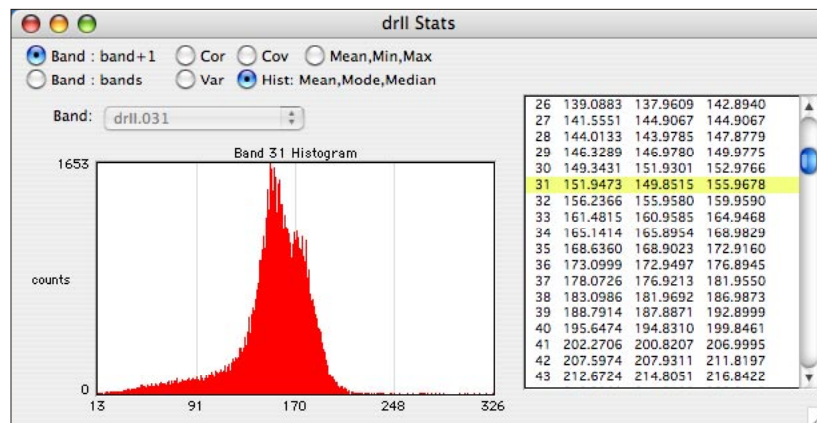


The matching Covariance and Correlation images corresponding to the statistics plots shown above. The image size is the number of bands and pixels represent floating point data: large values for Covariance and -1 to +1 values for Correlation.






The plot above displays the mean pixel value for each of 210 bands. The list shows the mean, minimum and maximum value for each band. Band 31 is selected in the list and the corresponding position within the plot is indicated by the blue vertical line



The pixel histogram for any of the 210 bands is displayed by selecting that option (Hist...) within the radio buttons and clicking on a list item which shows the mean, mode and median for that band.


Histograms

A histogram of the distribution of the pixel values for any single or multiband image can be generated for the front window by selecting menu **Functions -> Plots -> Histogram -> Data, Window or Masked**. The later is described on the following page. If a selection rectangle is present the histogram is confined to those pixels within the rectangle. If the band list for a multiband image has any bands disabled then those bands are ignored in computing the histogram. If the window is a cube then you will be presented with the dialog shown below allowing you to choose either a specific band in memory or the file data. Also, see section [Pixel Exclusion](#) which describes excluding certain pixel values.




A true color image produces three separate histograms: one for each of the three color components: red, green and blue. If the window is a color map overlay then the output is a list of the frequency count of the classes (see section [Image Products](#)).


Data generates the histogram of the underlying data and **Window** show the histogram of the remapped display gray values. The menu changes automatically depending upon the image content (e.g., pixel depth). To view a text summary with a cumulative distribution plot of the statistics, toggle **Detailed Messages** in menu **Edit -> Options** before histogramming.




Cube source selection dialog.




The total file histogram of all 160 enabled bands.



The underlying Data histogram of just band 60 with a cumulative distribution overlay.

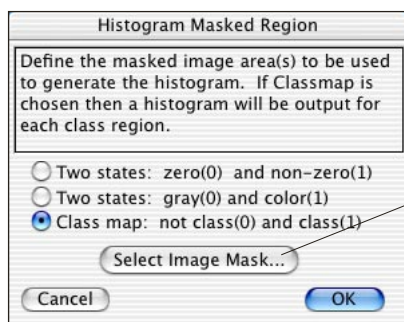


This summary window is displayed if Edit -> Options -> Detailed Messages is toggled on.

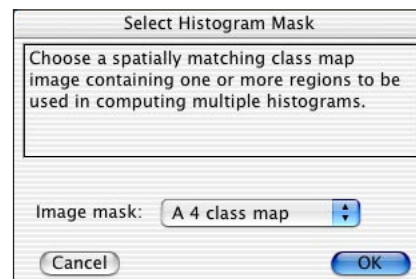


The display remapped gray shade histogram of band 60 also showing a cumulative distribution.


Menu Functions -> Plots -> Histogram -> Masked also generates a histogram but uses an intermediary image mask that defines the region(s) of pixels to include. The mask image must be the same spatial size as the target image, i.e., same number of pixels and lines, depth doesn't matter. Class map masks generate a histogram for each class.



A dialog allowing the choice between 3 types of mask images.



The resulting mask selection dialog for choice number 3 shown at left.



A four class overlay mask (partially shown) generated by a vector angle classification and used to generate the four histograms (see section [Classify](#) for a description of overlays and section Specific Algorithms for creating masks).




Image cube class 1 statistics.




Image cube class 2 statistics.




Image cube class 3 statistics.





Image cube class 4 statistics.

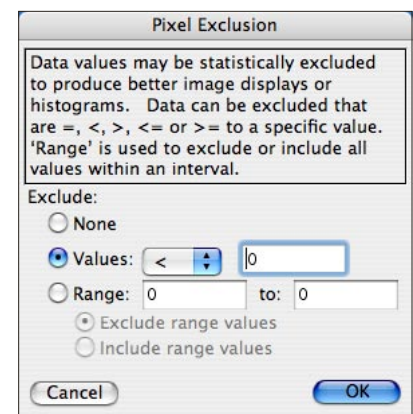
Pixel Exclusion

For display purposes and to yield better histogram plots you can select a value or range of values to be ignored (excluded from computation) when generating certain statistics. This option permits image pixels to be better mapped for display or to see details in a histogram that would be overshadowed by a too populous bin (e.g. a lot of zeros). All dialogs that contain a **Remap** checkbox obey the current global Pixel Exclusion definitions. Menu Applications -> [Shaded Relief](#) and [Contour](#) both contain a local option for convenience to specifically exclude zero which temporarily overrides global settings (see section [Applications](#)). Menu Image -> [Filter](#) (3x3 and 5x5) responds only to global pixel exclusion equal to zero. Excluded values do not affect analytical operations such as [Principal Components](#), [Classifications](#) and [Image Arithmetic](#), only what is displayed.

The following dialog is presented by selecting menu Edit -> Options -> Pixel Exclusion. Enabled exclusions are global, i.e., they are in effect until they are disabled.



The dialog shown at right will cause all negative pixels to be excluded when generating [Histograms](#) and determining the display remapping necessary when loading an image cube. This will also apply to the initial slope estimates in menu Applications -> [Shaded Relief](#) and to all dialogs containing a Remap checkbox. It has no effect on menu Functions -> [Statistics](#).




Scatter Plot


Choosing **Scatter...** from the **Functions -> Plot** submenu brings up the image selection dialog window shown below. After selecting two 8 bit images, which may differ in size, a two dimensional scatter plot is generated with the origin, point [0,0], at the lower left. A plot point is derived by using image one's 8 bit pixel value as the y coordinate and image two's pixel value as the x coordinate as the images are scanned in unison. Therefore the output image will always be 256 by 256 and mostly black. Each time there is a multiple hit, i.e. the inputs generate an x,y already plotted, the color (or gray value) for that point is modified from blue (black) towards red (white). If both selections are the same image then the output image consists of a 45 degree broken line of colored dots from the lower left to the upper right. After the output image is generated you may select other color tables from the menu string **Image -> Convert Image -> Color Table** to view different encoding schemes.

By selecting the scatter plot and moving the cursor over it the corresponding pixels in the two input images will be overlaid with red dots. To force their gray scale restoration choose the first color table (linear ramp) as described in the previous paragraph.


The initial coding method is determined by these check boxes and can be later modified by choosing a different color table.



All currently open 8 bit gray windows, regardless of size, are listed in the popup menus.




Red Tulips component.



Green Tulips component.

Points in this region correspond to picture elements where red is the dominant color.


The diagonal band maps to values of similar intensity in the two images.



The red and green component images of a true color image were used to generate this scatter plot image. The red image pixels are the vertical coordinate and the green pixels the horizontal coordinate. The red tulip flower tops form the cluster in the upper left. As the cursor is moved over the plot the inverse transformation shows which pixels in the source images contributed to this point.




Line Profile Plot

A linear plot of image values can be generated by first drawing a line and then selecting menu **Function -> Plot -> Profile**. The line is specified by positioning the cursor on one endpoint (start), depressing the control key (previous versions used the shift key) and clicking and dragging the mouse to the other point (end). Upon releasing the mouse a marque-like line will remain. Repeat the procedure to select a different line. Control-click (control-double-click in the Windows version) the mouse without moving it to clear the line or choose menu **Edit -> Clear**.



If an image is color, including pseudo color, then three separate profile plots will be generated for the red, green and blue components respectively. Plots are resizable and include the coordinates of the start and end points and the image scale in meters per pixel if the image is geo [referenced](#). Overviews show properly scaled coordinates. Vertical coordinates are expressed as either gray values ([0,255]) or in data values ranging from the minimum to the maximum data within the profile line (see section [Plot Scaling/Selection](#)). Also, see section [LASF Point Cloud Viewer](#) for a special type of line profile.

The profile window is an independent window that can be printed or copied to the clipboard, and saved. To save a profile as an image activate its window and select menu **Edit -> Duplicate**. Then save the duplicate as a normal image. The actual profile values can be saved as a text file via menu **File -> Save**. The file will consist of 1024 x,y,z triplets regardless of the length of the profile line. The x,y,z triplets are also copied to the clipboard when performing **Edit -> Copy**. Moving the cursor over a profile plot line will display the corresponding image x,y,z values in the **Info** window if it's enabled (see section [Info Window](#)).

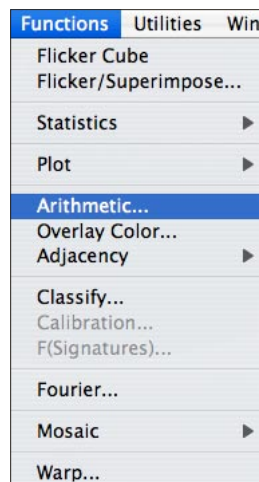


Resulting profile plots displaying the underlying 16 bit data and the 8 bit display data corresponding to the line drawn on the geo-referenced shaded relief image above. The starting, ending, length and scale of the profile line is displayed along with the geometric angle (ccw from the start point horizontal line) and the heading angle (cw from the start point vertical line).

Image Arithmetic

HyperCube allows you to enter general “C” type arithmetic expressions where the operands are images and constants. The result is an image whose size equals the minimum spatial size and maximum pixel depth of the individual image operands. Since the arithmetic generally produces values outside of [0,255] you can specify how to remap the resultant display image. All arithmetic is performed to floating point precision and remains attached to the resultant image. If an expression contains a color image the output will be a color image limited to the display depth of 8 bits per color without the attached floating point values (see examples).


Menu **Functions -> Arithmetic...** opens the dialog window shown below. The scrollable area contains a list of all of the current images that you can use to form an expression and is in order of creation. Each image name is preceded by the letter “s” and a number. These become the symbols used to form the expression instead of using the explicit image name. The expression itself will be used as the resultant image name and it too will be added to the image list. The expression is typed into the upper edit box and must conform to the “C” syntax. Listed below and on the next page are the operators and functions recognized by HyperCube where s1 and s2 may be images or constants. The result is always an image of the minimum spatial size of the image operands with possibly underlying floating point attached data.



(,)	left and right parenthesis	==	equality
+, -, *, /	arithmetic operators	!=	inequality
%	modulo	>=	greater than or equal
&,	bit-wise and, or	<=	less than or equal
^	bit-wise exclusive or	>	greater than
<<, >>	bit-wise shift left, right	<	less than
+, -	unary signs	pi	constant (3.14159...)

fabs(s1), sqrt(s1), pow(s1,s2), exp(s1), log(s1), log10(s1)
 sin(s1), cos(s1), tan(s1), asin(s1), acos(s1), atan(s1), atan2(s1,s2)

The Image Arithmetic dialog window displaying a list of images that you can use as operands in an expression. The example shown at the right generates an image mask of 0 and 255 representing the areas that have an x and y slope magnitude greater than 1.



Additionally, the following non-C functions are available where n is a constant. As is the case with the C type functions each result is an image (see `hist()` for additional outputs).

<code>shift_left(s1,n)</code> <code>shift_right(s1,n)</code> <code>shift_up(s1,n)</code> <code>shift_down(s1,n)</code>	Shift image s1 a distance of n pixels (n may be fractional) in the specified direction.
<code>slopes(s1,n)</code> <code>slopey(s1,n)</code> <code>slopexy(s1,n)</code>	Generate the slope of image s1 over a distance of n pixels, in the x, y or combined xy direction (n may be fractional).
<code>int(s1)</code>	Truncate the fractional value of image s1 to integer.
<code>min(s1)</code> <code>max(s1)</code> <code>mean(s1)</code> <code>mode(s1)</code> <code>median(s1)</code> <code>sd(s1)</code> <code>cov(s1,s2)</code>	These 6 functions create an image with constant values equal to the specified statistic of image s1. Standard deviation of s1. Generate the covariance between images s1 and s2.
<code>norm2(s1,s2)</code>	Generate the normalized distance, $(s1 - s2)/(s1)$, between images s1 and s2.
<code>red(s1)</code> <code>green(s1)</code> <code>blue(s1)</code>	Extract the red, green or blue plane of a true color image. If s1 is gray then the function output is s1.
<code>hue(s1)</code> <code>saturation(s1)</code> <code>intensity(s1)</code>	Extract the hue, saturation or intensity plane of a true color image. If s1 is gray then the output is s1.
<code>gray(s1)</code>	If s1 is true color then return a gray image weighted as $.33*red + .33*green + .33*blue$ (equal weights). If expression s1 is not true color then it is remapped to [0,255] using a 98% interval about the median.
<code>rgb(s1,s2,s3)</code>	Form a true color image from 3 arguments where s1, s2 and s3 represent red, green and blue respectively. Equivalent to Image -> Convert -> R,G,B to Color .
<code>hsi(s1,s2,s3)</code>	Form a true color image from 3 arguments where s1, s2 and s3 represent hue, saturation and intensity respectively. Equivalent to Image -> Convert -> H,S,I to Color .
<code>ternary(s1,s2,s3)</code>	Similar to the C ternary operator, $(s1)? s2:s3$. Use image s2 if s1 is not zero and image s3 if s1 is zero, pixel by pixel.

<code>conv(s1,s2)</code>	Perform the convolution of s1 and s2 (frequency domain).
<code>corr(s1,s2)</code>	Perform the cross correlation of s1 and s2.
<code>dft_power(s10)</code>	Generate the Fourier frequency power spectrum of s10.
<code>dft_phase(s10)</code>	Generate the Fourier frequency phase spectrum of s10.
<code>dft_nterms(s1,n1)</code>	Fourier transform s1, keep largest n1 % terms, invert result. Demonstrates how few Fourier coefficients are needed.
<code>class(s1,n)</code>	Extract, as a color mask image, the n'th class from class map image s1.
<code>band(s1,n)</code>	Retrieve the n'th band from cube image file s1, i.e., access the cube file itself.
<code>band_sum(s1,n1,n2)</code>	Form the simple summation of bands n1 through n2 from cube file s1. Ignore disabled bands.
<code>hist(s1)</code>	Produce, as a separate output, the histogram(s) of s1. Also, outputs s1.
<code>hist_masked(s1,s2)</code>	Similar to <code>hist()</code> but uses image s2 as a [0,!0] mask to define the areas of s1 to use in the histogram.
<code>save_file(s1)</code>	Save image s1 as a TIFF file. If s1 is a cube file (not a specific band) save its corresponding window.

The following 8 functions compute and use the discrete wavelet transform (dwt or its inverse) of argument s1. If the spatial size of s1 is not a power of 2 then it is automatically zero filled to the next power of 2 size (also making a square image) and the transform taken. Arg n1 is the number of Haar or Daubechies coefficients to use, either 2 (Haar), or 4,...,20 (Daubechies). Argument n2 is the multiresolution analysis level and may range from 1 to log2 of the image size (e.g., 9 for 512 by 512). A special case of n2 = 0 generates the maximum level. Because of the power of 2 requirement, functions `dwt()` and `dwt_inv()` should first be used by themselves and the resulting power of 2 size wavelet or inverse then used freely in subsequent expressions. The remaining 6 functions can be used directly within any expression since their image argument, s1, is zero filled as necessary, processed and extracted back to original size.

<code>dwt(s1,n1,n2)</code>	Compute the forward discrete wavelet transform (dwt) of s1.
<code>dwt_inv(s1,n1,n2)</code>	Compute the inverse discrete wavelet transform.
<code>dwt_low(s1,n1,n2,n3)</code>	Compute the forward dwt of image s1. Keep all transform values with magnitude < n3, zero the rest and perform the inverse transform to produce a filtered image.
<code>dwt_high(s1,n1,n2,n3)</code>	As above, keep values with magnitude >= n3, zero the rest.

`dwt_energy(s1,n1,n2,n3)` Compute the forward dwt of image `s1`. Rearrange the dwt in decreasing energy order. Normalize and form summation until energy sum $\geq n3$ where `arg n3` is a percentile [0,100] of the total dwt energy. Zero transform terms less than last energy addition and invert to form an image.

`dwt_nterms(s1,n1,n2,n3)` Compute the forward dwt of image `s1`. Rearrange the dwt in decreasing energy order. Keep the greatest `n3` terms where `arg n3` is a percentile [0,100] of the total number of dwt terms. Zero the remaining terms and invert to form an image. For example, if `s1` is 512x512 then the expression `dwt_nterms(s1,4,0,10)` transforms `s1` to the highest level (9) using Daub4 coefficients, keeping the greatest $26214 = (262144 \cdot .1)$ energy terms and zeroing the rest. The remainder is inverted to form an output image.

`dwt_avg(s1,n1,n2)` Compute the forward dwt of image `s1` to level `n2` using `n1` coefficients. Zero all detail (fluctuation) terms, i.e., just keep the level `n2` trend terms. Invert the result to produce an averaged image.

`dwt_detail(s1,n1,n2)` Compute the forward dwt of image `s1` to level `n2` using `n1` coefficients. Keep only the level `n2` fluctuation terms. Zero the trend and level 1 thru `n2-1` fluctuation terms. Invert the result to form a detail level `n2` image.

Note: For example, if image `s1` is 512x512 and `n1 = 4` then the following are true:

$$s1 = \text{dwt_avg}(s1,4,n2) + \text{dwt_detail}(s1,4,1) + \text{dwt_detail}(s1,4,2) + \text{dwt_detail}(s1,4,3) + \dots + \text{dwt_detail}(s1,4,n2).$$

i.e., an image is the sum of inverted dwt trend and fluctuation terms.

and:

$$s1 = \text{dwt_high}(s1,n1,n2,n3) + \text{dwt_low}(s1,n1,n2,n3)$$

$$s1 = \text{dwt_energy}(s1,n1,n2,100)$$

$$s1 = \text{dwt_nterms}(s1,n1,n2,100)$$

for any valid `n1, n2, n3` values.

The relational operators (e.g., <, >, ==) produce a mask image consisting of 1's where the pixel by pixel relation is true and 0's where it is false. For example: the expression `s1==s2` creates a mask with 1 for each pixel of `s1` and `s2` being equal and zero otherwise. The expression: `(s1==s2)*s1` shows those pixels of `s1` that match `s2`, whereas `(s1!=s2)*s2` shows those pixels of `s2` that differ from `s1`.


The expression: `((red(s3)>=green(s3))&(red(s3)>=blue(s3)))*s3` displays all areas of the true color image `s3` in which the red intensity is greater than or equal to both the green and blue intensities; black elsewhere. See more examples later.

In addition, a new constant type operand is available that consists of a triplet set of numbers represented as `{n1,n2,n3}`. When used in an expression such as `s1*{.95,1.3,.6}` the red component of `s1` is multiplied by .95, the green component by 1.3 and the blue component by .6. You can use the triplet constant anywhere that a single constant is valid. It will force the pixel depth to true color.

All image operands are converted to floating point before arithmetic is performed. The result of the expression must then be converted back to 8 bits/plane to display an image (the underlying data remains the results of the expression and is left attached to the image).

The **Image Arithmetic Options...** button allows you to determine how the display pixels generated by the expression are remapped back to the [0,255] range. If the **Remap** check box is unchecked then the expression is truncated (but not the underlying data) such that values above 255 are set to 255 and values less than 0 are set to 0. This is useful when the expression only contains bit-wise operators. Generally, expressions produce pixels well out-of-range. In this case select the **Remap** check box and choose the **min & max** radio button, then the expression range [min,max] will be linearly remapped to [0,255]. For expressions that produce a wide non uniform range of values choose the **median** radio button. Then, the value of the expression is remapped so that values in the range [median - %/2, median + %/2] are linearly transformed to [0,255]. Values outside of this range are set to 0 and 255 respectively. The **Do each color separately** checkbox specifies whether one set of statistics is applied globally for all three colors or whether three sets of statistics are separately used to remap each color independently of the others. **Set divide...** defines the result of any zero divisors that may occur (almost a certain consequence of using the divide operator). The 3 possibilities set the resultant pixel to zero, one or 255.

Expressions are evaluated from left to right unless parentheses force the order of execution. An expression can consist only of constants in which case the output image is a uniform color. All 8 bit pseudo color images are converted to their respective true color values during computation. The result remains true color, it is not reconverted to pseudo color. Any expression containing only gray results in a gray output image at full precision. Any expression containing color produces a color output image at 8 or 16 bits per color.




The Arithmetic Options dialog allows you to generate a sequence of outputs that are assembled into a “synthetic” cube where each plane is the result of one evaluation of the sequence. To use this option check **Generate a sequence**, and enter the number of outputs and range that will be linearly substituted for the variable ‘t’ in the expression. For example: the arithmetic expression `slope(s1,t)` with the options shown will generate a cube consisting of 10 planes where each plane is the slope in the x direction over an interval of ‘t’ pixels, i.e., 1,2,3,...,10. The variable ‘t’ is not limited to integer values and may be substituted anywhere in the expression that the substitution produces correct syntax. Another example: `((s4>=mean(s4)-t*sd(s4))&(s4<=mean(s4)+t*sd(s4)))*s4` sequentially displays the values of image s4 lying within the expanding interval $[\text{mean}-t*\text{sd}, \text{mean}+t*\text{sd}]$ with ‘t’ set to range from 0 to 3.

The **Save File Name** button is used to designate base name for the `save_file()` function. Each file name will be the ‘base name’ concatenated with a 3 digit sequence number.

The **List constant function result** checkbox will display the result of constant functions in a separate text window. These functions include: `mean()`, `mode()`, `median()`, `min()`, `max()`, `sd()` and `norm2()`. Also, the various discrete wavelet transform functions, `dwt_energy()`, `dwt_nterms()`, ..., display detail in the text window. Uncheck this option to disable the text window output.



Band 60 (s1) of HYDICE image cube (available at www.tec.army.mil/Hypercube/index.html).



An 8 class map (s2) generated via K-means using the class mean as the class color. See section [Spectral Classifications](#).



The result of combining the above left gray shade image, s1, and the above right color class map, s2, using the arithmetic expression: `hsi(hue(s2), saturation(s2), gray(s1))` with Remap disabled in Image Arithmetic Options.

Examples:

where: s1 is a data image; s2,s3,s4 are bands of a cube and s5 is the shaded relief of s1

`((s1>1000)&(s1<1200))*s1` Keep s1 values between 1000 and 1200, zero rest.

`log(s1)` Should also choose a remapping option.

`atan2((s2-s3),(s2-s4))` Also try various remappings.

`shift_left(s1,3)` Shift the s1 image 3 pixels to the left with wraparound.

`shift_up(s1,.5)` Shift the s1 image up .5 lines (fractional shifts are allowed).

`slopex(s1,2)` Generate the slope of image s1 over an interval of plus and minus 2 pixels in the x direction.

`gray(s5)` Use the display gray values of s5 instead of any attached data.

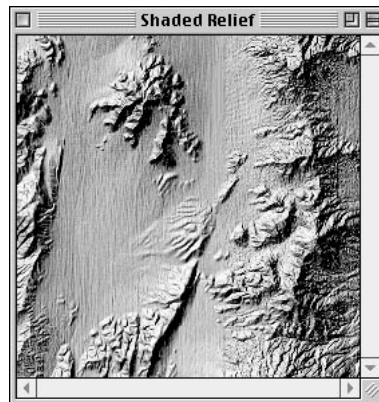
`(s1>1200)*{32,0,0}+gray(s5)` Generate a red tint image of s5 greater than 1200.

`(fabs(slopex(s1,1))>1&fabs(slopey(s1,1))>1)*255` A mask of areas with x and y slope magnitude greater than 1.


`hist_masked(s2,s2==s3)` Generate a hist of s2 wherever pixels s2 and s3 are equal.




A data source image.



Corresponding shaded relief generated via menu [Applications->Shaded Relief](#).



The mask image resulting from the last expression in the examples.



This image of the above source image is the result of applying the the following arithmetic expression:

`(slopex(s1,1) > 0)*{255,0,0} +
(slopex(s1,1) < 0)*{0,0,255} +
(slopex(s1,1) == 0)*{0,255,0}`

so that positive x slopes are colored red, negative x slopes blue and zero x slopes green. Expressions may be copied and pasted to and from a text window ([File->New](#)) and then saved.

More examples:

where `s1` is a gray or true color image or cube file band.

`dwt(s1,4,0)` Compute the discrete wavelet transform of image `s1` using Daubechies 4 coefficients to level 8 (arg `n2 = 0` implies the maximum level).

The resulting transform image is always a power of 2 in size.



Source 256 by 256 by 8 bit gray image.



The 8 level dwt using 4 Daubechies coefficients.

`dwt_energy(s1,20,0,99.8)`



Compute the 8 level dwt of image `s1` using 20 coefficients and keep 99.8% of the energy of the transform. Zero terms representing the remaining .2% and invert the transform to form a filtered image.




A profile line displayed over the original image along with its plot below.



The resulting filtered image with the same profile line and plot shown below.






© 1965 Kodak

Original 512x512 8 bits/color image.

Shown at the left and bottom are two images representing the source and result of [dwt](#) processing, respectively. The bottom image has been re-constructed from only 10% of the dwt transform terms and is visually indistinguishable from the source image. This represents a 10:1 compression. Also shown is the color by color result of the transformation. Argument level 0 uses the max level which in this case is 9 ($\log_2 512$).



© 1965 Kodak

Reconstructed image from the arithmetic expression:

[dwt_nterms](#)(s1,20,0,10)

where 20 is Daub20, 0 represents level 9 and 10 is %.

dwt_nterms:

red:

Daub20 level 0
% energy 99.9122
% terms 10.0002
terms 26215
abs threshold 13.7915

green:

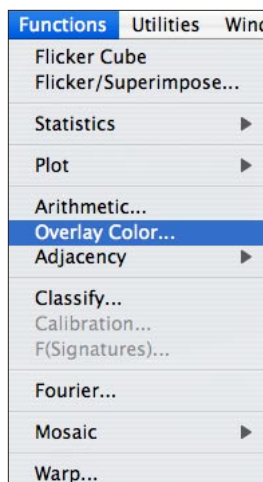
Daub20 level 0
% energy 99.913
% terms 10.0002
terms 26215
abs threshold 13.3667

blue:

Daub20 level 0
% energy 99.8488
% terms 10.0002
terms 26215
abs threshold 15.3345

Arithmetic detail output produced via [Arithmetic Options](#) List constant function result.

Overlaying Images



A special image function is provided to allow overlaying one image on top of another. This could be accomplished by splitting the images into red, green and blues planes followed by various logical operators previously described in Image Arithmetic and then re-combining the resulting planes. The Functions -> Overlay Color... menu makes this much easier: you select the two images and the program will construct a third true color image that is the result of overlaying image 1 upon image 2. The first image is the mask or top and the second image is the base or bottom. Every black pixel in the mask becomes clear when placed upon the base thus allowing the base image to show through at that pixel. Every non-black pixel in the mask overwrites the base image. The spatial sizes and pixel depths of the source images do not have to be the same, the output image will be the size of the overlap and true color. Pseudo color and gray scale are automatically accounted for. Note that the mask criteria is based upon color, in this case black, and not on a zero index if the mask happens to be a pseudo color image.

Choosing Overlay Color... from the Functions menu brings up the following dialog:






Image 1 (mask)




Image 2 (base)




Color Overlay

Adjacency

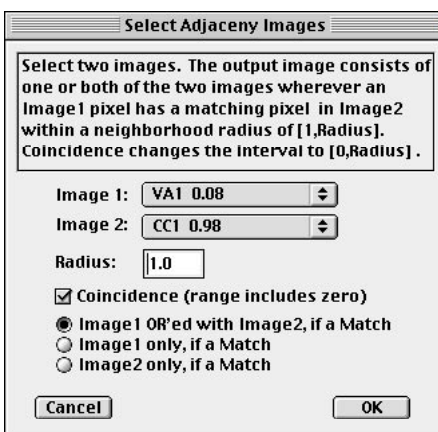

This image function can be used to delineate where two or more pixels (classes) have the same value within a selected spatial neighborhood. Four variations are available: the first, **Images**, cross compares the pixels in two different images, the second, **Colors**, uses a single image in the comparison and the third and fourth, **Classes** and **Singularities**, pertain to spectral classifications. The idea is the same for all: generate an output image consisting of a combination of input image pixels that match some criteria anywhere within a specified radial distance of each other. The neighborhood can range from zero to some arbitrary radius.




Menu Functions -> Adjacency -> Images is a generalization of equality. Given two images it sequentially compares a pixel of the first image with the neighborhood surrounding the similarly located pixel in the second image. For example, a radial neighborhood of 1 is interpreted as comparing the source pixel with any of the 4 adjacent pixels in the second image (a radius of 1 is not large enough to include the diagonal neighbors). If Coincidence is checked (see below) then the central pixel is also included. A radius of 2 increases the valid comparison region to 12 or 13 pixels depending upon Coincidence. Fractional radii are allowed, e.g. a radius of 1.42 would include 8 or 9 pixels. If the radius is set to zero then the function reduces to equality. The output pixel can be set to just the source pixel, just the matching neighborhood pixels or, both as selected by the three radio buttons. The image examples below demonstrate how two spectral classification output masks are combined to produce a new mask in which the algorithms agree with each other allowing a spatial uncertainty of 1 pixel, e.g. where **Vector Angle** and **Correlation Coefficient** matched exactly (Coincidence set) and within 1 pixel of each other.




The two source spectral classification outputs used as input to the Adjacency function producing the output shown at the lower right.





Dialog is set to produce an exact match as well a spatial mismatch of a one pixel radius.



Menu Functions -> Adjacency -> Colors operates on the front image window and is used to extract pixels with two specific colors that are co-located within a selected radius. This operation could be done by first separating the source image into two masks of the given colors (see section [Image to Mask](#)) and then using Adjacency -> Images as above but Functions -> Adjacency -> Colors greatly simplifies this. The image example below shows how a multiple spectral classification output can be transformed into an output showing the outline of everywhere one class is exactly adjacent to another class, e.g. grass and tree canopy demarcation. Again, the selected radius allows a considerable latitude in the equality range. The two colors to be compared are best chosen by first positioning the cursor over the source image and using the special command-option-C combination described in section Copy, Paste & Duplicate. The colors will then automatically be filled in when Functions -> Adjacency -> Colors is selected. Note that checking Coincidence has an effect only when the two selected colors are identical in which case the resultant image will consist only of that extracted color (see section [Image to Mask](#) to accomplish the same thing).

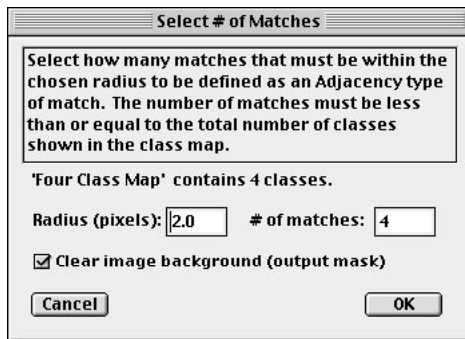


The dialog above was used on image Six Class Map to extract tree [43,96,43] regions that are adjacent within a 2 pixel radius to grass [195,255,195]. The output consists of only the grass border.




The third variation, **Functions -> Adjacency -> Classes**, is only available when the front most image is a classification overlay produced via the **Classify** function. This type of image displays the member class as the cursor is moved over it. The intent is to extract all of the areas within the overlay that have a common adjacency but instead of specifying colors you specify how many of the class members must be present within a certain radius. If you select less than the actual number of classes that the overlay possesses then this subset could and would change throughout the image. This permits generating a spectral classification using similar signatures from a library (again using **Classify**) and then entering how many of these signatures are needed (found) within a certain radius to reduce the ambiguity of false “hits”. The output may be an image or a mask depending upon the **Clear image background** checkbox. The image appears where the criteria fails and some combination of the class member colors where it succeeds. If a mask then black appears where the criteria fails. The mask can be overlaid on the original classification source using menu **Functions -> Overlay Color**.

The following user dialog and examples illustrate how you can increase the accuracy of a spectral library classification.




Class map image Four Class Map consists of 4 classes. All four classes are used and searched within a 2 pixel radius to generate the output mask image shown below right.



This classification overlay image was generated by using 4 signatures from a spectral library, each depicting some characteristic of a vehicle.

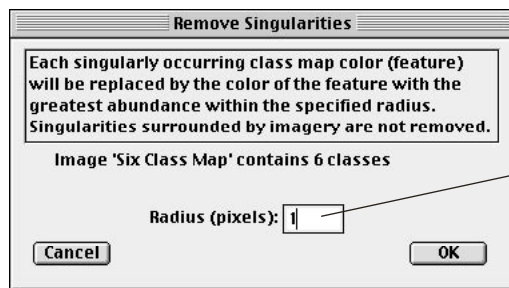
Vehicles




After applying Adjacency -> Classes to the left hand image.

The fourth variation, Functions -> Adjacency -> Singularities, is also only available when the front image is a classification overlay produced via the **Classify** function. This type of image displays the member class as the cursor is moved over it. The idea here is to scan the image for all occurrences of a single class map color (feature) that is surrounded by a different feature or features within a specified pixel radius. Wherever this happens the singularity is replaced by the surrounding feature with the greatest abundance within the radius. If the single class feature is surrounded by imagery (gray) then no replacement is performed. The final output of this function is an almost identical class map with single pixel “noise” replaced by surrounding features. Since this modifies what the **Classify** function thought was correct you must carefully inspect the new class map to determine the validity of this operation. The best use of this function is to de-clutter class map images before converting them from raster to vector form (See section [Class Map Editor](#)).


Below are the user dialog and an example of applying this function.



As the radius increases, the probability of finding a singularity greatly decreases. With a radius of 1, only the top, bottom, left and right pixels need to be different.



*A three times magnification of a portion of a six class map as output from the **Classify** function. There are 169 singularities within this subset.*




The same subset after removing each occurrence of a single class surrounded by different classes.

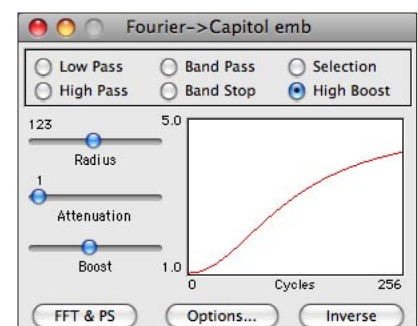
Fourier Operations

Frequency domain filtering is initiated by selecting menu **Functions -> Fourier**. The dialog presented allows you to interactively change various filter parameters and apply them to an image. The filter's characteristic curve is shown within the plot window as sliders are moved. The radii sliders represent the half magnitude point. The attenuation slider determines how sharp the filter is. All filtering uses the same scenario: select part or all of an image, click the **FFT & PS** button, select the filter type, move the sliders for the desired effect, and click the **Inverse** button to generate the filtered image. If **Dynamically** is enabled in **Fourier Options** then the filtering is performed as the sliders are moved. **Fourier** is an operator like gray mapping, it is unique to the active image window when it was selected. There can be several **Fourier** operators, each connected to its own source image.

An important visual key to interactive frequency domain filtering is the power spectrum (PS). It displays the relative amplitude of the frequencies comprising the image. The center of the power spectrum represents the lowest frequency with frequency increasing radially outward. As you move the sliders the radii change indicating which frequencies will be affected for the type of filtering chosen.




The images above show the effect of applying the High Boost filtering operation. At the left is the Power Spectrum corresponding to the upper left source image. See functions `dft_phase()` and `dft_nterms()` in section [Image Arithmetic](#) for additional Fourier operations.



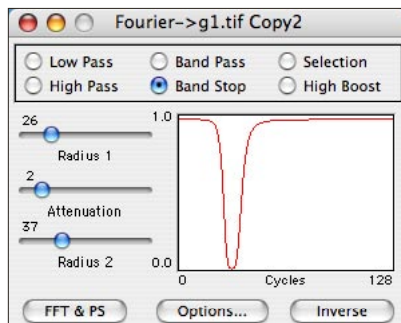
Fourier contains five slider controlled filters and one user selection filter. These are defined as follows:

Low Pass	Attenuate frequencies above Radius 1 . Use this slider and the Attenuation slider for desired response.
High Pass	Attenuate frequencies below Radius 1 . Also use Attenuation slider to shape response curve. Because the average brightness is zeroed remapping (Options) must be chosen.
Band Pass	Pass only those frequencies between Radii 1 & 2. Radii sliders select pass band, Attenuation slider determines sharpness. Also requires remapping option.
Band Stop	Attenuate frequencies between Radii 1 & 2. The inverse of a pass band.
High Boost	Similar to High Pass . The boost varies from 1.1 to 9.9 and is controlled by the second radii slider (label changes to Boost). Remapping should be used for data type images.
Selection	Attenuate frequencies within a user selection rectangle. Manually draw a selection rectangle on the power spectrum image.

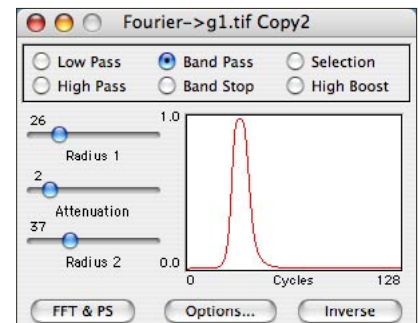
The image area to be filtered must be a power of two sized square. HyperCube accomplishes this by automatically embedding the source image or selection rectangle in the smallest power of two square image that will enclose it. The border is set to the mean of the image (selection) to minimize ringing in the inverse transform. All operations then proceed using the embedded image. The result will be extracted as a separate image and may, optionally, be pasted back into the source image. If the source is an 8 bit pseudo color image you will be directed to convert it to a true color copy using menu **Image -> Convert Image -> Indexed to Color**.



The Fourier Options dialog shown above presents additional controls. The top check box determines whether the inverse image also overwrites the original image patch. The second check box changes the filter characteristic from a tapered one to a vertical cutoff. The **Remap** check box works exactly as that in the [Image Arithmetic](#) section and is a must for high pass and band pass filters which zero out the fundamental frequency (image average). This is also especially true for any images that represent data. The PS Compression popup is used to present a better display of the power spectrum. If **Dynamically** is check then a new invervse image is generated for every movement of the sliders.



Typical characteristic curves for Band Stop, left, and Band Pass, right, Fourier filtering.



Mosaic Points

Two or more images can be fitted together to form a single image. They are joined by selecting conjugate image points on each of the images. The geometry used to force the joining ranges from a simple orthogonal transformation to a plane to plane projection. However, the transform must be the same for all of the images. **Mosaic** is more general than [Warp](#) (described in the next section) in that the end result is to stitch images to a common geometry for a greater view whereas warp's objective is to match the geometry of one image to another.

By choosing menu **Functions -> Mosaic -> Points**, a dialog will appear with several scrollable list areas. Only the **Available Images** list initially contains anything, namely all open images that are candidates for mosaicking. Click on each **Available Images** list name to highlight it and click the **Add ->** button to move the image name to the **Mosaic Images** list. Double clicking an **Available Images** name will also move it directly. The order of selecting images is not important (see Load/Save points). To remove an image from the **Mosaic Images** list click its name and click the **Remove** button or double click its name, it will be returned to the **Available Images** list. Removing an image after starting the control point selections will remove all control points. You are warned of this and are allowed to cancel the removal.

Once the mosaic list has been determined the program is ready to begin the control point selections. Position the cursor over any image and shift-click the mouse, a colored numbered reticle will be left there. Move to the conjugate position on another image and shift-click there. This can be repeated until each desired image contains a matching control point. Subsequent attempts are disabled until the **New Image Cntrl Pt** button is clicked. This permits selections for the next control point number. Each control point must appear on at least 2 images before going to the next control point. **New Image Cntrl Pt** is automatically disabled until at least 2 images contain a given control point. A control point does not have to appear on every image, only on at least two. As a control point is chosen it appears in a scrollable list. A point may be deleted by clicking on it and choosing the **Delete Image Cntrl Pt** button. It will be deleted from each image it appears upon and the remaining points will be renumbered. The number of control points determines the degree of transformation (see section [Warp](#)) that will be allowed. Therefore if two images have only 2 control points in common then only the **Orthogonal** transformation will be available regardless of how many common points may appear on other images.

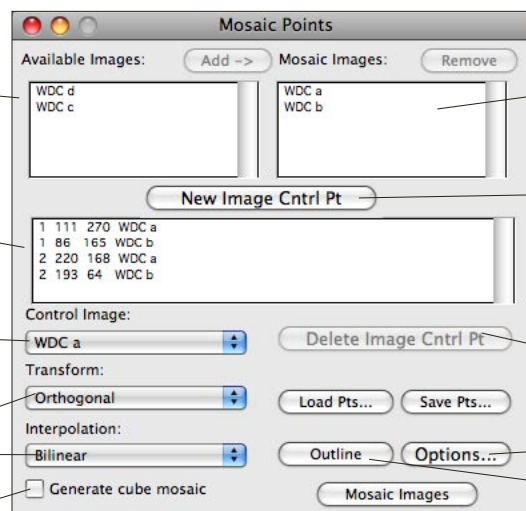
Unselected available images. Data type images may also be mosaicked together.

Two control points and their coordinates that have been measured.

The image chosen to be the geometric control

Type of transform and interpolation.

Mosaic each cube band.



List of images to form the mosaic.

Activate another set of selection points.

Delete the selected control point.


Described on subsequent pages.

After all of the control points have been selected you can mosaic the images by clicking the **Mosaic Images** button or generate an outline showing how the images will be placed in the mosaic by using the **Outline** button. This generates a colored line drawing image that can be overlaid on top of the final mosaic using [Overlay Color...](#) in the **Functions** menu. The **Control Image** popup is used to select which one of the mosaic list images defines the geometry that the remaining images are conformed to. The control image shape stays the same and the other images are warped to its geometry using the selected **Transformation** and type of **Interpolation**.

Any of the mosaic images may be chosen as the **Control** and it will appear “stationary” in the output. The biggest problem with mosaicking to a control image is the cascading of errors: e.g., Image 1 is the control and has points in common with image 2 which has points in common with image 3 but image 1 and 3 have no (or not enough) points in common. Then any errors between 1 & 2 and between 2 & 3 cause very obvious misfits between overlap boundaries because the global least squares fit is to the control image’s geometry. A way to minimize the effect of poor control points is to choose a good image as the **Control** image and then only mosaic it and one other image together, then save the resultant mosaic image and restart a new mosaic operation using the added image as the new control image (the current mosaic will generally have black areas that would overwrite if it were chosen as the control image). This has the effect of breaking the geometric fit into regions (more degrees of freedom) and results in a more aesthetically pleasing final mosaic.



The result of mosaicking the two images in the previous dialog is shown on the left along with the outline generated for the same pair to the right.



Below is a description of each of the **Mosaic** window controls along with the possible options and their effect.

Add -> : Move the selected image from the Available Images list to the Mosaic Images list. Double-clicking does the same.

Remove: Move a selected Mosaic list image back to the Available list. Also may double-click.

New Image Cntrl Pt: Used to step to the next control point, is disabled until the current point has been selected on at least two images.


- Control Image:** Selects which of the mosaic images will be the control geometry. Unless averaging (see below) the control image will always be the top most image in the mosaic, i.e., it will overlay the others in common overlaps.
- Transform:** The type of geometric transformation. See section [Warp](#).
- Interpolation:** What type of interpolation to use in the resampling. See [Warp](#).
- Outline:** Generates an output image the same size as the mosaic images but only shows the outlines of the individual mosaic images. Becomes enabled when a mosaic is possible. Also see [Options](#).
- Load...:** Load a previously saved set of control points.
- Save...:** Save the current set of control points as an ASCII file.

Note: Saved and reloaded control points must exactly match the images and order of the original mosaic. However, once reloaded they can be deleted, added to, and resaved.

- Options...:** Bring up the **Mosaic Points Options** dialog. See below.

Mosaic Images: Execute the mosaic operation.

Selecting the **Options...** button displays the following user dialog.

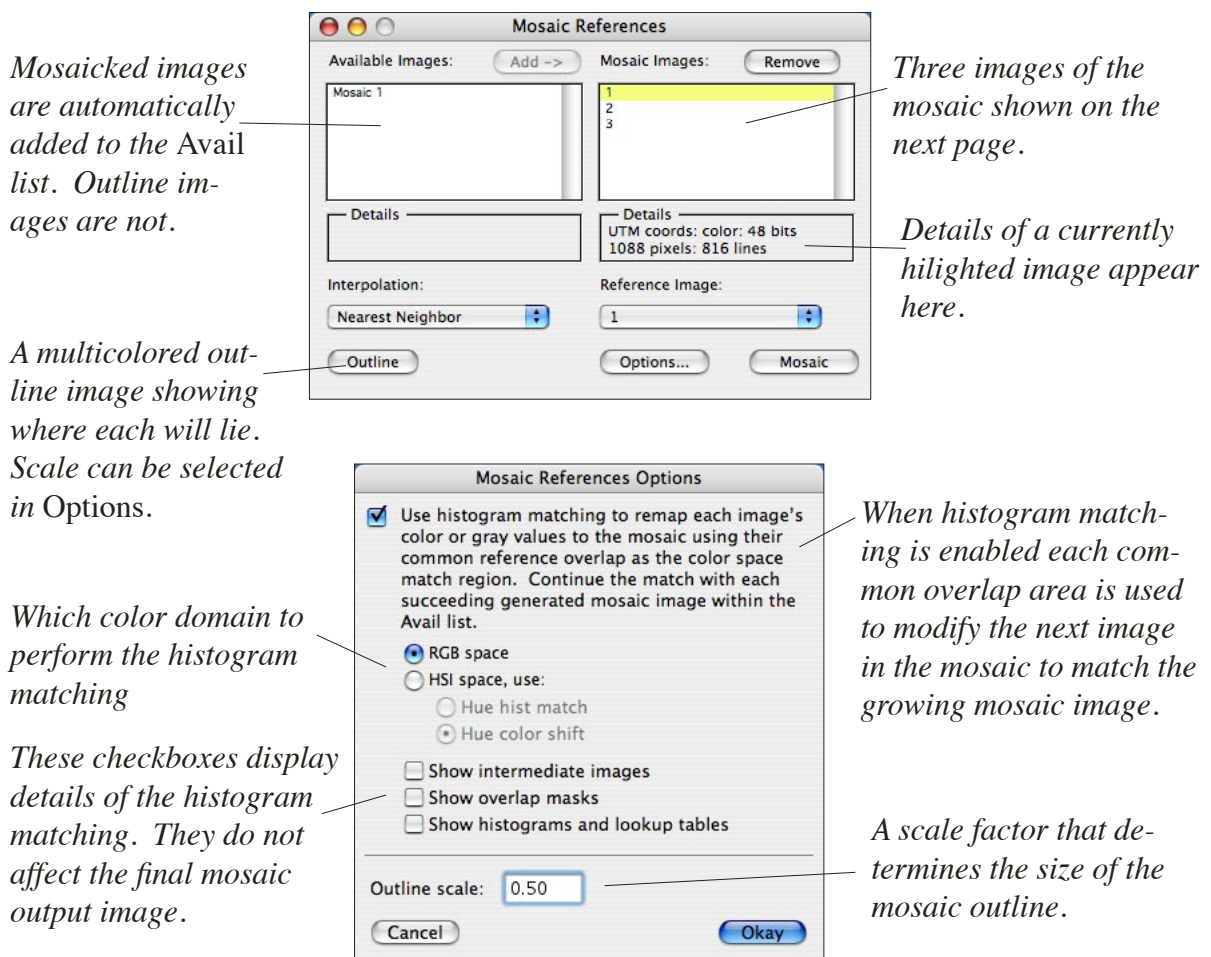


- Average...:** Where the mosaic forms overlaps the control image's pixel takes precedence followed by the order of the **Mosaic images** list. When this option is checked all of the images within an overlap will be averaged together to form the output pixel.
- Display mini...:** When checked, the **Outline** button will generate a miniature instead of a full scale line drawing. The mini cannot be correctly overlaid with the Mosaic output.
- Void areas fill:** The gray value that is used to fill new undefined areas.

Mosaic References


This function also mosaics two or more images together except that no image points are used. Instead, each of the images must have a geometric reference in a common coordinate system (see section [Image to Reference Coordinates](#)). Menu Functions -> Mosaic References function computes the outer boundary of the images in the Mosaic Images list and then places each image in the proper location based upon its coordinates. The resulting mosaic is created in the order of the Mosaic Images list with each image overlaying the mosaic under construction. Note: If the overlaying pixel is zero then it does not overwrite the lower layer. There are several options that determine how the images are combined in the overlay region. One of the Mosaic Images list determines the scale of the final mosaic. That is, even though all images must be in the same coordinate system they don't have to have the same scale. For example: one image could represent a meter per pixel while another could be .5 meters per pixel. The scale of the selected Reference image determines the final scale of the mosaic. Pixels are warped as necessary to make everything merge correctly. The interpolation used is selected by the Interpolation popup menu and can be nearest neighbor, bilinear or cubic resampling. Nearest neighbor works best for images that are essentially the same scale.

The Mosaic References dialog is similar to the [Mosaic Points](#) dialog. They both show a list of available images from which specific images can be selected for mosaicking (this also determines the mosaic order). Double-click an image in the Avail list to add it to the Mosaic list. Double-click a Mosaic list image to move it back.





Three source images sharing a common coordinate system, UTM, to be mosaicked together. In this case, they all appear to have a similar scale. However, that is not necessary. The Reference image popup in the Mosaic References dialog determines the final scale of the mosaic and popup Interpolation determines the resampling.



The result of selecting the Outline button in the Mosaic References dialog. The Reference image outline is always shown in red.



The final mosaic image with image 1 selected as the Reference image. Therefore, the remaining images have been properly scaled using Nearest Neighbor as the interpolation.

Saving an outline or mosaic image as [NITE](#) will also embed the Rational Polynomial Coefficients ([RPC](#)) in the image header.

Warp

Menu **Functions -> Warp** allows you to force the geometry of one image to fit the geometry of another image. You can force two completely different types of imagery, such as optical and radar, into registration, combine them and thus produce a new product. The registration points, the pixel resampling, the final output scale and the type of transformation used are all selectable.

After choosing the **Functions -> Warp** menu, you select two images from the dialog popup list of images. All single images and cubes are permissible. Using the mouse control points are selected alternately between the images as follows: position the cursor on control point 1 of image 1 and click the mouse while holding down the shift key, then position the cursor on the corresponding point 1 of image 2 and shift-click, repeat until all control points have been selected. Each pair of points leaves a cross hair mark behind, usually red on image 1 and green on image 2. The coordinate pairs are listed in the **Warp** dialog window where they can be edited or deleted. Either image may be [zoomed](#) to help position the cursor. The number of points needed depends upon the transformation, higher order requires more evenly distributed points.

The result of warping is a new image whose output size is determined by the value in **Scale factor**. A value of 1 produces an image the size of the destination image. Depending on the transformation there may be pixels in the output image that have no correspondence in the source image, these will be set to **Off image value**. Transformations are performed in a least squares fit manner to distribute the error across the control points. The **Residual** button shows the mean error of the fit without producing an output image.

Shown here is the **Warp** dialog for a projective transformation.

Warp

Image 1: WDC512 Image 2: WDC512 Warp: ☒ 1 -> 2 ☐ 2 -> 1 ☐ Cube

Instructions:
Shift-click a control point in image: 'WDC512'.
To remove a point, select it in the list and press the Delete button.

Control Points:

1	19	25	22	33
2	21	237	21	223
3	276	237	268	225
4	248	24	251	33
5	134	32	134	48
6	134	241	132	230
7	130	123	148	125
8	29	135	34	143
9	263	107	252	110
10	277	183	274	186
11	88	28	88	22

Transform: Projective Interpolation: Cubic

Edit Del

134 32 134 48

Off image value: 0 Scale factor: 1

Load Pts... Save Pts... Residuals Warp

Warp residuals

RMS residual: 8.83, Projective
Point by point:
X res, Y res, dist

1	-3.026022	-1.180439	+3.248114
2	+3.278908	+12.236832	+12.668515
3	-1.523285	+2.054245	+2.557405
4	-0.218521	+2.974651	+2.982667
5	+4.432907	-7.050091	+8.327931
6	+3.757156	+4.866318	+6.147949
7	-14.937515	+3.133507	+15.262642
8	-2.717484	-3.557953	+4.477024
9	-9.733688	+3.446083	+10.325705

Warp coeffs

Coefficients:
Projective
X coef Y coef

1	+0.913406	-0.027090
2	-0.022181	+0.968090
3	+2.258384	-5.583946
4	-0.000244451	-0.000244451
5	-0.000279393	-0.000279393

Residuals and coefficients for the above example when the Residual button is clicked.


The transformation selected has a great affect on the shape of the output image. The non-linear transforms are very sensitive to the number of control points and their positions. The transforms along with their arithmetic are defined below. Each definition is actually a pair, one for the pixel coordinate and one for the line coordinate. The Projective transformation has the same denominator for both pixel and line. The value within parentheses is the minimum number of points needed for the specified transformation.

- Orthogonal (2): $ax + by + c$ (with constraints on a,b for both pair of equations)
- Affine (4): $ax + by + c$ (no constraints on coefficients)
- Quadratic (4): $ax + by + cxy + d$
- Quadratic (6) $ax + by + cxy + dx^2 + ey^2 + f$
- Quadratic (8): $ax + by + cxy + dx^2 + ey^2 + fx^2y + gy^2x + h$
- Cubic (8): $ax + by + cxy + dx^2 + ey^2 + fx^3 + gy^3 + h$
- Cubic (10): $ax + by + cxy + dx^2 + ey^2 + fx^2y + gy^2x + hx^3 + iy^3 + j$
- Projective (4): $(ax + by + c)/(dx + ey + 1)$ for pixel
 $(fx + gy + h)/(dx + ey + 1)$ for line

The orthogonal transform preserves shape, allowing only rotations and a single scale change. Two points are required to define it. The second, affine, is a linear transform that preserves parallel lines and is a generalization of orthogonal. It will transform a square into a rotated parallelogram and is the most useful transform for registering two similar images. The nonlinear transforms include cross products, second and third degree terms and produce curved lines.

Three methods of pixel interpolation are possible: nearest neighbor, bilinear and cubic resampling. Nearest neighbor is the fastest, it simply takes the closest source pixel and uses it as the output. That is important if multiband imagery is warped and then used in a classification as no new intermediate “fringe” classes are created. Bilinear interpolation uses 4 surrounding pixels to form a positive weighted average. Cubic also uses weights but some weights are negative. This produces a slight sharpening in the output image and appears superior to bilinear for natural images but produces ringing in artificial images that contain discontinuous gray values.


The figures below demonstrate the transformation of an image warped to a duplicate of itself. The control points were purposely chosen to be slightly different. The higher the degree of warp (e.g., cubic) the greater the number of control points needed. And they must be well distributed across the image.



Orthogonal.



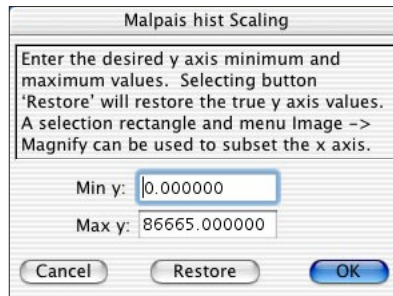
Projective.




10 Term cubic.

Plot Scaling/Selection


When a plot is first generated the y axis limits are determined by the minimum and maximum of the functional values and the x axis values represent the range of the function. They remain so until changed, even if the plot window is resized. The vertical axis limits of certain plots can be explicitly changed by choosing menu **Functions -> Plot -> Scaling** and the x axis range can be changed by drawing a selection rectangle around the area of interest and choosing menu **Image -> Magnify Plot** (the magnification value is set to 1). Each type of plot has a corresponding scaling dialog (e.g., [Plotting Signatures](#)). Shown below is an example of these operations applied to a histogram plot.




The dialog corresponding to the histogram shown on the left, below. The rescaled histogram is shown on the right after changing Max y to 100,000.



Original histogram with the y axis limit set to the maximum functional value. A selection rectangle has been drawn to select an x axis subset.



The resulting plot after changing Max y to 100,000.



The new plot displaying only values within the selection rectangle after choosing menu Image -> Magnify Plot.

Multiband Functions

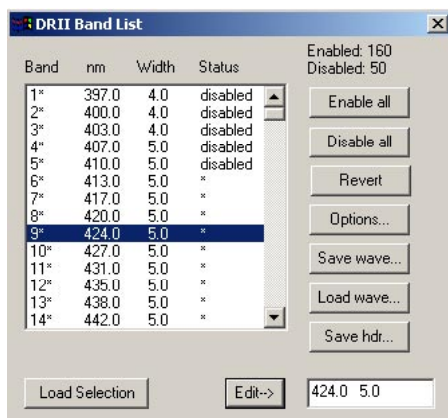
Band List

Every multiband image has an associated band list that specifies the wavelength and status of each band. This is necessary in order to correctly plot spectra and perform library searches and comparisons. You can view/modify a band list by selecting menu **Windows -> Show Band -> List**. A resizable window with a scrollable list of 4 columns and several buttons will appear. The columns display the band number, wavelength, band width and status for each band. Except for band number, the items are editable by: clicking them, typing a new value in the edit field and clicking the **Edit** button. Double clicking the band number in the Mac version will automatically load that band as a single image not connected with its cube parent. You must use **Load Selection** in the Windows version in order to load a specific band after highlighting it in the list.


The band list is the window display of the [*.wvl](#) file described in the [Wavelength File](#) section. If a [*.wvl](#) (or older [.wav](#)) file exists when the multiband image was loaded then HyperCube converts and loads its contents into the band list. The [*.wvl](#) file may also be loaded via the **Load wav...** button. By choosing the **Options...** button, the band list can be modified as desired and then saved as a new [*.wvl](#) file via the **Save wav...** button. Also, it can be copied to the clipboard via menu **File -> Save** or **Edit -> Copy**.

Band list window controls have the following definitions:

- List column:
- 1 Sequential band number, same as image cube band numbering. If the band wasn't loaded, i.e. only a subset of the full cube, an asterisk appears after that band number. You can load any band as a stand-alone image by double clicking its band number.
 - 2 Band center wavelength in (usually nanometers or microns).
 - 3 Band width in same units of measure as wavelength.
 - 4 Status, if "disabled" (no check mark) then the band is ignored in spectral classifications, library searches and histograms. It will be highlighted in red on spectral plots.



The Windows version has an explicit Load Selection button, instead of double clicking.



Mac version band list showing the same selection area.

Enable all	Set the status of all bands to enabled (check mark). Toggle individually by clicking on Status (4) column entry.
Disable all	Set the status of all bands to “disabled”. Also toggle individually as above.
Revert	Restore the band list to the last loaded wave length file.
Options...	Display the Band List Options dialog. Sets up various combinations of enabled and disabled bands. See details below.
Save wav...	Save band list as *.wvl file.
Load wav...	Display a standard open file dialog and load the *.wvl file.
Save hdr...	Save the current cube definitions in a *.hdr file.

In combination with the **Band List Revert** button, **Band List Options** allows you to define a variety of enabled bands for experimentation. There are 7 mutually exclusive options: the first 2 change the wavelengths, the third rescales (a convenient way to convert between nanometers and microns) and the last 4 enable or disable groups of bands. However, the first 3 of the Enable/Disable options will not enable any band that is currently disabled in the **Band List**.

Start & end is used to set a uniform sequence of wavelengths. Delta is computed.

Start & delta is similar to start & end with a user defined delta.

Multiply each wavelength and width by the scale entered below. Band width and times delta change to Scale factor and times wave respectively when selected.

Only those wavelengths within the interval values defined in Image -> Cube Color Composite -> Definitions are enabled.

The complement of the above option, i.e., the Color Comp Defs are disabled.

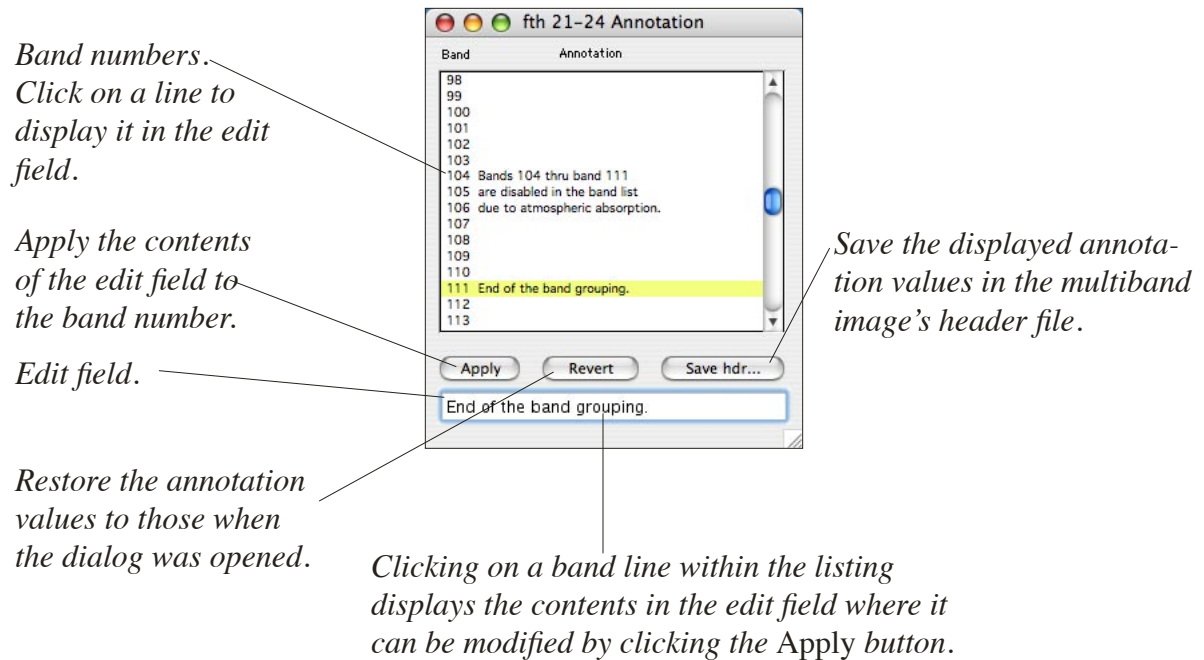
Only enable every n'th band that was already enabled, e.g., a value of 2 enables every other band currently enabled.

The last of the 4 Enable/Disable options: Specify the Start and End band numbers to be enabled or disabled.

Band List Options dialog.

Annotation List

You can associate a short (up to 64 characters) free text message with each band of a multiband image. The messages are displayed in the band scrolling area of a multiband window (section [Band Scrolling](#)). Menu Windows -> Show Band -> Annotation displays a dialog containing a listing of each band. Clicking on a band line displays the corresponding message in the edit field that can be modified and updated with the **Apply** button. Choosing **Save hdr** within this dialog or similarly in the [Band List](#) dialog will include the annotation in the header ([.hdr](#)) file and automatically display it whenever the multiband image is



Band Scrolling

Every multiband image is loaded into HyperCube so that you can easily scroll through the bands in a movie like manner. The band scroll bar (buttons in the Windows version) is across the top of the image cube window and works like any other Mac scroll bar. The current band number and corresponding wavelength and any annotation is shown at the top of the image. If a specific band's status is disabled then its band number is shown in red else it is black. A specific band can be given its own display window by scrolling the cube to that band and then choosing **Duplicate** from the **Edit** menu. It becomes completely separate from the cube. Any band, whether it was loaded with the cube or not, may be loaded by double-clicking a band number in the **Band List** (see sections [Band List](#) and [Band Annotation](#)).

Band scroll bar, number, wavelength and annotation area.



Windows version.




Mac version.

All controls except for Band Scrolling operate in the same manner in both versions.

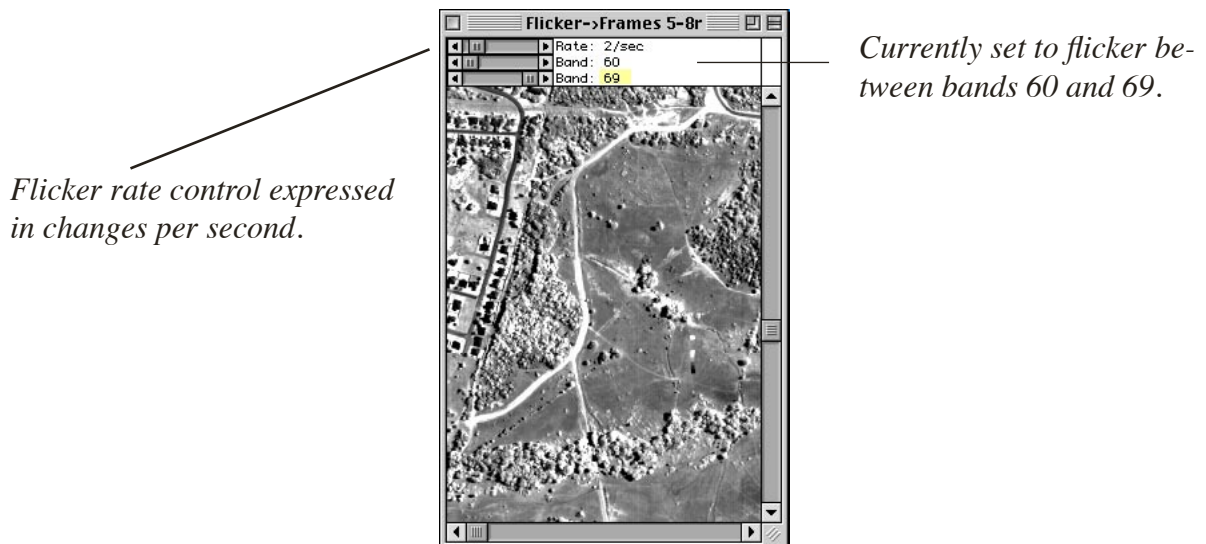
Flicker/Superimpose

While band scrolling permits traversing an image cube in a contiguous fashion sometimes it's desirable to jump from one band to another to compare differences. HyperCube contains two functions that facilitate this. The first is the menu **Functions -> Flicker Cube** and is enabled whenever a multiband cube window is front most and the second method is menu **Functions -> Flicker/Superimpose** which is enabled whenever at least one image window is open. **Flicker Cube** differs from **Flicker/Superimpose** in that it is specifically for the bands of an image cube and is adjustable, whereas **Flicker/Superimpose** is fixed to the two specific images (of any type) which drive it (see **Flicker Image** dialog below).



Flicker Cube creates a new window similar to the cube window but with three scroll bars at the top which control the rate and the two bands to alternate. **Flicker Cube** is an operator like graymap and if the cube window from which it is derived is closed then flicker disappears. You can change the spatial offset between the two flickering images by using the keyboard arrow keys while the flicker is the active window. Each keystroke moves image 1 a single pixel in the arrow direction. Enable menu **Windows -> Show Info** for a dynamic listing of the current x,y offsets and the alternation between the pixel value under the cursor (see section [Info Window](#)). The flicker image can be printed and also copied.

The figure below is the flicker window generated from the previous band scrolling example. The flicker rate is nominal and some PC's may not support the higher rates. Alternation continues even if the flicker window is not the active window.



Flicker rate control expressed in changes per second.

Currently set to flicker between bands 60 and 69.

The second flicker type function, **Flicker/Superimpose**, is more general in that it permits any two image windows (and selection rectangles within) to be alternated between or superimposed one on another as if viewing through a porthole. The images are selected from two popup menus that appear when menu **Functions -> Flicker/Superimpose...** is selected. The images do not have to be the same size, the viewing window is sized to their minimum. The image depth, i.e. 8, 16 or 32 bits, also doesn't matter. If one of the source images is modified such as by gray mapping or filtering then the change will be automatically reflected in the view window. Similar to **Flicker Cube**, if either source image is closed or rotated 90 degrees the view window will disappear. The window title alternates between the source images when **Flicker** is chosen. Also, the keyboard arrow keys can be used to spatially offset image 1 relative to image 2. Enable menu **Windows -> Show Info** to view the current off-sets (see section [Info Window](#)).

Below is the dialog in which you select the two images and either **Flicker** or **Superimpose**. If **Flicker** then the rate of alternation for flicker can be set. If the resultant flicker is too fast or too slow close the flicker window and generate another one. **Superimpose** places a rectangular window (or inset) of the first image upon a base consisting of the second image. This rectangle can be repositioned by clicking the mouse within it and dragging it anywhere within the window dynamically superimposing image 1 upon image 2. To reverse the ordering, reselect **Flicker/Superimpose** and swap **Image 1** and **Image 2** names in the popup menus. **Superimpose** is most useful when comparing an image band with a classification output image (see Products).

Flicker alternates between the two images whereas Superimpose displays a rectangular viewing port of image 1 upon image 2.



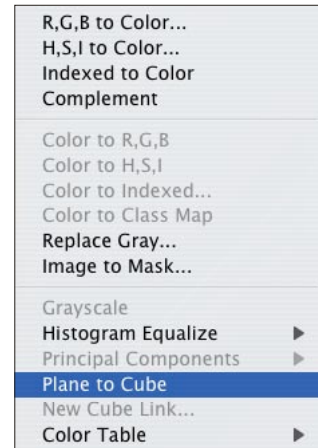


Image 1, Six_classes..., has been selected to superimpose itself upon image 2, r49_g55_b18... The resultant window is shown at left in which the selection rectangle can be dynamically moved about the window.

Constructing/Modifying a Cube

HyperCube can generate a new image cube from existing images of the same spatial size having various pixel types (e.g., byte, short, long, float for gray and byte or short (3) for color). Any image is a candidate including bands of an existing image cube. However, the new cube must be all of the same kind, either gray or color. The bands for the new cube may be already resident in memory or reside in a folder (see following page). The in memory technique for construction involves two steps: selecting the first band which determines the spatial size and adding subsequent bands to the cube. As new bands are added the pixel type of the cube automatically changes to accommodate the maximum number of bits/pixel required. HyperCube can also directly convert a true color image into a 3 band cube. After completing the construction save it via menu File -> Save As..., Multiband .

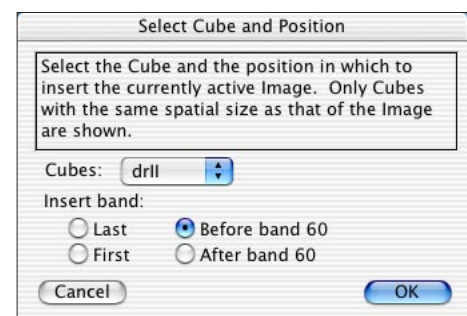




1. Click on the first image component. This will become the first band.
2. Choose menu Image -> Convert Image -> Plane to Cube.
3. Click on the second component.
4. Choose menu Image -> Add to Cube... A dialog appears asking you where to insert the new band (see bottom right).
5. Repeat steps 3. & 4. to complete the cube. Menu Image -> Delete Band can be used to delete the current band face (i.e., first scroll to band to be deleted).
6. Select menu Windows -> Show Band List and manually set up any wavelengths for each band (the default will be band numbers).
7. Choose menu File -> Save As... and select Multiband to save the cube.

If converting a color image, an additional dialog is presented asking whether to directly generate a 3 band cube (red, green, blue) or the first band of a synthetic color cube as shown in the dialog at the right.

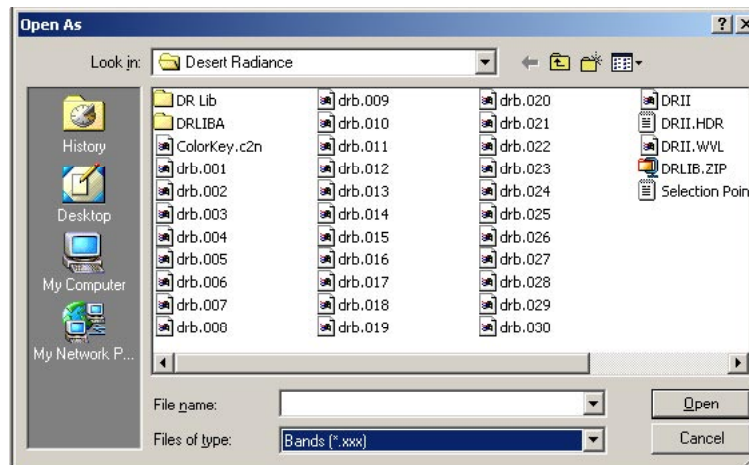
New bands may also be added to an existing image cube file by loading the entire cube. A warning dialog similar to that shown at the lower left will appear when the first band is added stating that the cube will be disconnected from the existing file, becoming an all new cube that must be saved.

After adding or deleting bands, choose menu Window -> [Show Band](#) List to edit the list and then File -> Save As..., Multiband (use a different name than the original cube file).

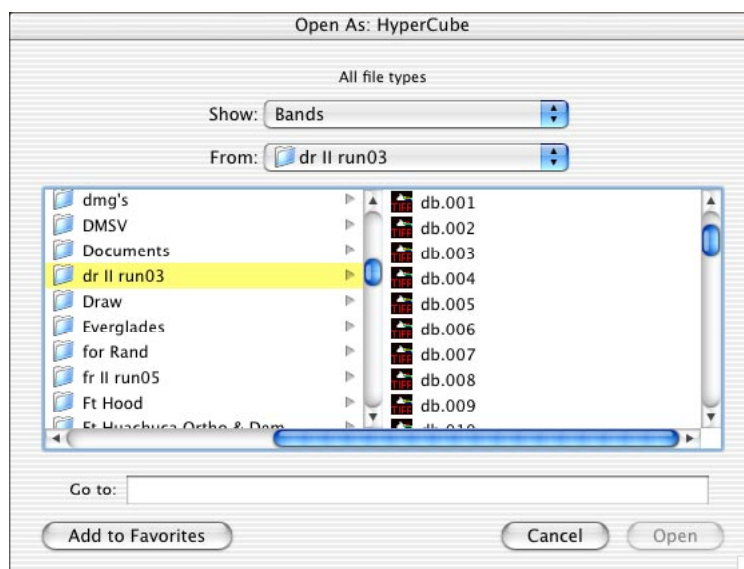


The second method to construct a cube is to have an existing set of TIFF image files within a folder. These may have been created by a program such as MATLAB®. Each file name must have the format: *.xxx, where * is a fixed base name and xxx is a zero padded value in the range [001,999], e.g., dr.001, dr.002, dr.003,... . The numbers don't have to be consecutive, just increasingly ordered and may reside with other files. The files are loaded and automatically constructed into a cube by choosing menu File -> Open As..., Bands and then selecting any one of the ordered files as shown by the Windows and Mac dialogs below. The first band will be the image file with the lowest number value in the folder. Note: reading a list of files (bands) like this is the inverse of selecting a cube image window and then menu File -> Save As..., Bands (see section [Saving Images](#)).

You can scroll the cube bands and use menus Image -> Delete Band, Edit -> Duplicate and Image -> Add to Cube to modify it. After constructing, choose menu Windows -> Show Band List and edit the band values as necessary before saving (File -> Save As..., Multiband).



Windows Open As..., Bands.




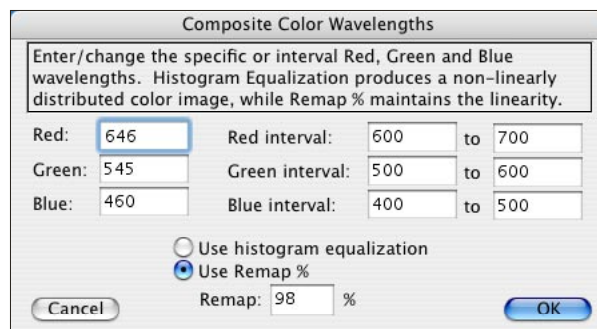
Mac OS X Open As..., Bands.

Color Composite

HyperCube provides a direct function to construct a true or false color image from a cube. Menu Image -> Cube Color Composite allows the user to automatically generate a color image for the currently active (front) image cube window and to link it to the cube for point selection and classification (see sections [Spectral Plots](#) and [Classify Functions](#)) breaking any previously linked color image. Two methods are provided via a submenu: **Specific Wavelengths** and **Interval Wavelengths**. The former finds the three closest bands to the wavelengths set in the **Definitions...** menu (see below) and the latter integrates over three spectral intervals to generate the color components. In each case the data is extracted from the image file. Disabled bands (see section [Band List](#)) are ignored so it is possible to have a null interval and HyperCube will warn you if this occurs. The resultant color image is correctly mapped so as to maintain the original brightness relationships, i.e. the red, green and blue are normalized as a group.

Choosing menu Image -> Cube Color Composite -> Definitions displays the dialog shown at the right. If the image cube header file contains “color...” keywords (see section [Keyword Format](#)) then they appear in the dialog edit fields else the initial wavelength values for the **Specific** wavelengths red, green and blue component bands are 646, 545 and 460 nm (or μm) respectively. These correspond to a set of Kodak® color separation filters and may be changed by editing the dialog fields. **Interval** wavelengths values are similarly entered specifying the spectral range of integration to be used to form each color component. Note, identical or overlapping values are allowed. The **Remap %** determines the area about the median of the composite histogram that will be remapped to the [0,255] for each color. A good starting point is 98%. **Use histogram equalization** performs a nonlinear stretch of the components and is useful when the image cube has an asymmetrical histogram. Changes may be incorporated in the image header file upon closing the dialog and selecting menu Window -> Show Band List, Save hdr... (see section [Band List](#)).

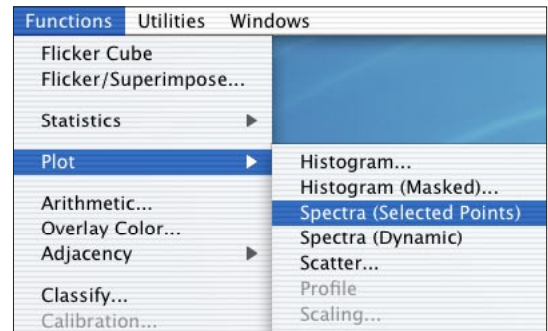
If menu Edit -> Options -> Detailed Messages is checked then the composite histogram of the three bands is generated. This is similar to the histogram described in section **Statistics** except that it represents only the three bands or band intervals. A typical composite histogram is shown below.



Bands 49, 35 and 17 were automatically selected from the active image cube file to form a composite color image. The histogram highlight represents 98% of the data in these 3 bands. File values between 14 and 223 will be remapped to the display range of [0,255] for each color.

Spectral Plots

Two methods of plotting spectral vectors from an image cube are provided: statically via **Functions -> Plots -> Spectra (Selected Points)** or **Functions -> Plots -> Spectra (Dynamic)** as a function of the cursor position. The former, described here, allows up to 128 points to be plotted along with library signatures superimposed upon the same plot (see section **Spectral Classifications**). The latter allows the simultaneous display of the closest matching library signatures to the current cube vector (see **Dynamic Displays and Plots in Spectral Classifications**).




Points are selected by positioning the cursor and clicking the mouse while holding down the shift key. This draws a numbered colored cross hair on the cube face. The Plot window need not be open when points are selected and if no plot points exist when a Plot window is opened it will contain a message indicating what to do. If the wavelength file is incorrect no plotting is possible and the Plot window will contain a warning to that effect. Having the Info window open while you select points facilitates the process by dynamically showing the spatial coordinates. As points are selected they are graphed in the Plot window along with a color patch used for toggling them on and off. The number of color patches that can be toggled is determined by the size of the Plot window, resizing the window (lower right grow icon) displays more patches. All of the points can be erased at one time by choosing **Clear** from the **Edit** menu or selectively in the **Classify** dialog. When the cursor is positioned over a static Plot window the spectral and reflectance coordinates will be listed in an opened Info window.

The static Plot window contains several controls for customizing the plots: three check boxes, a standard button and dynamic colored push-buttons that toggle on and off their matching color plots. Cube bands that have been disabled in the wavelength file have a red highlight superimposed on them. Controls have the following definitions:



- Scale:** Rescale the vertical plot axis between the minimum and maximum of all of the vectors. Otherwise, the minimum and maximum will be set to the arithmetic range of the pixels.
- [0,1]:** Change the vertical axis labelling to the range [0,1] and map the plotted values to it. Also see **Scaling...** example below.
- Hide Bad:** If checked then suppress the plot wherever a band has been disabled. See section [Band List](#).
- Clear Lib:** Remove any superimposed library signatures.



When a plot window is active Scaling may be selected from the Plots submenu to produce a dialog window which will result in [0,1] changing to [0,1.5].



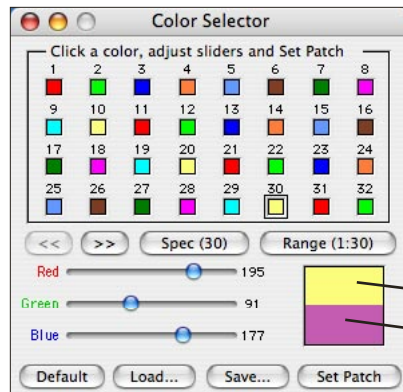
Four spectral plots from the 70 band image shown at the right. Each number corresponds to the numbered cross hair in the image. Several bands at 750 nm have been disabled in the wavelength file due to a sensor error. These points may also be selected from a linked red, green, blue composite image constructed from cube bands (see section [Components to Color](#)).



A different image cube spectral plot after resizing to show all of the corresponding color patches. Numbers 2 and 11 have been toggled off. Because Hide Bad is checked, plots within the two major water absorption regions are not drawn (disabled in the Band List). The correspondence between a point number and its color may be changed using the Color Selector item in the Options menu.

Selection Point Colors

When HypeCube is launched the correspondence between a selection point number and its color is set from an internal default table that will be used when plotting spectral signatures and classifying (see **Products in Spectral Classifications** section). The color sequence is arbitrarily repeated in groups of 10 with each group slightly different in value. Therefore, although points 1, 11, 21, 31, 41, 51 and 61... appear as the same red, they are not. The point # to color correspondence can be customized at any time by choosing menu **Edit -> Options -> Color Selector**. Up to 1024 colors may be uniquely redefined. When a change is made all windows using selection point colors will be updated to reflect the change. Classify output products that have already been generated will not be changed (Overlays and Masks will use the new color on a subsequent classify). However, products can be reasigned colors and labels after the fact by using menu **Utilities -> Color to Name Editor** (see **Dynamic Displays and Plots**). A color that has been copied to the clipboard can be pasted as the selected color (see section [Copy, Paste & Duplicate](#)). Also, see **Classify -> Options -> More Options -> Use source spectra color** in section **More Options** as an alternate method of specifying class map colors.



Color patch #30 has been selected and the sliders adjusted to a new color. The components of each color are shown to the right of the sliders and range from 0 to 255. You can not select a gray value (i.e. all three colors with the same value). The Selector will automatically modify the blue value to be different.

Current palette color.

New palette color.

To change a single color click on its patch to highlight it (a small frame is drawn around it) and adjust the three sliders for the desired color. The top half of the palette shows the current patch color while the bottom half of the palette changes to show the new color.

- | | |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| << >> | These buttons are used to navigate through the 1024 colors moving forwards and backwards 32 colors at a time. |
| Spec (N) | Generate a <u>rainbow</u> spectrum in the first N colors ranging from red through yellow, green, cyan and blue. N is the currently toggled color patch. If no patch is selected then N will equal 1024. |
| Range (1:N) | Generate a smooth range of colors between the colors in patches 1 and N. If no patch is toggled then 1024 colors will be generated. |
| Set Patch | Change the patch color to the new palette color and update all windows. |
| Default | Change the selected color patch to its original HyperCube launch color and update all windows. If no patch is currently selected, a dialog appears asking you if all patches are to be set to their defaults. |
| Save/Load... | The current color key correspondence is saved/retrieved to/from an ASCII file. The output is the same format used in the Color to Name Editor , i.e., red, green and blue values followed by a name which in this case is the index number. The default file extension is "c2n". |

Spectral Libraries

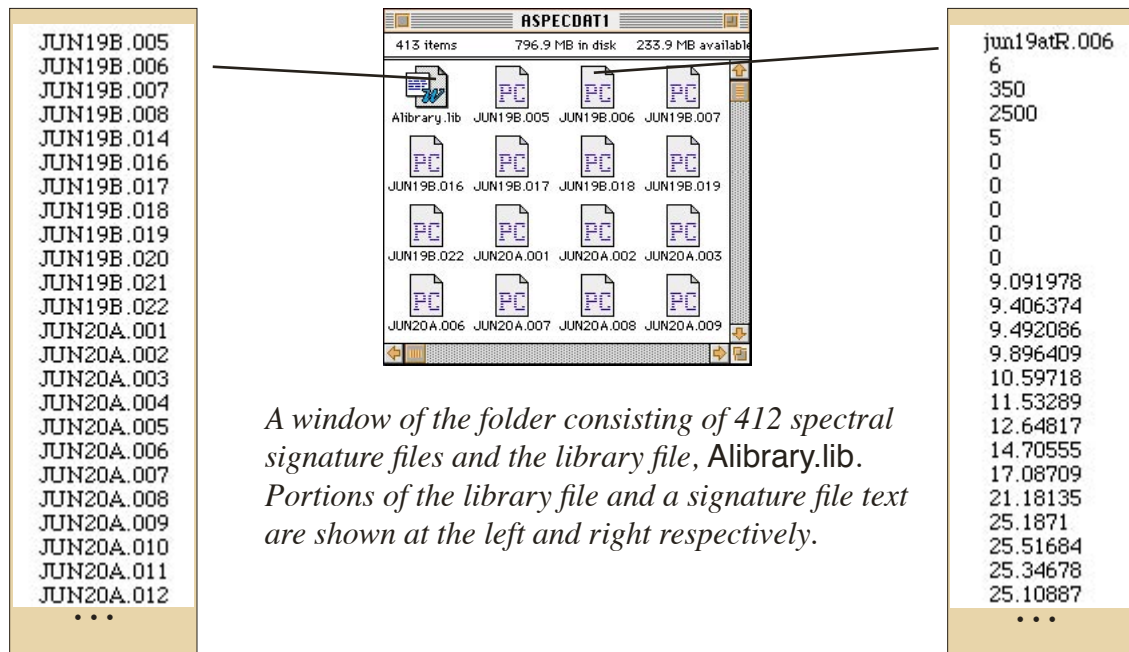
Format

The HyperCube program currently recognizes 3 types of spectral signature files: a signature format developed at USATEC, the NGA NEFDB (Nonconventional Exploitation Factors Data Base) and a simple list of wavelength and value pairs. One signature file represents the spectra of a single material. A set of signature files constitutes a library of materials. The various signatures belonging to a library do not have to be related in any manner, they are often grouped together to correspond to a particular sensor run.

The format of a USATEC signature file is defined as: a small header followed by an ordered set of linearly spaced reflectances. All of the elements are in ASCII text and can be easily edited in a word processor or spreadsheet application. Almost any ASCII text file can be read, edited and resaved in HyperCube, including signature files (see section [Loading/ Saving Text Files](#)).

A library is a text file whose elements are signature file names, i.e. a library contains no data except file names. The library file and all of the signature files must be in the same folder.

The relationship between the library elements in the USATEC library is depicted below: Each spectral signature, regardless of size, uses the following format:



Alibrary.lib file with an entry for each signature file name.

This USATEC spectral signature file contains a total of $431 = (2500-350)/5+1$ reflectance values.

<u>Line:</u>	<u>Parameter:</u>	<u>Example:</u>	
1	File name	jun19atR.018	
2	Spectrum number	18	
3	Starting wavelength (nm)	350	
4	Ending wavelength (nm)	2500	<i>Lines 3,4 and 5 implicitly determine the number of reflectances.</i>
5	Wavelength increment (nm)	5	
6	Scan index number (hr)	0	
7	True solar time (hr)	0	
8	Solar azimuth (degrees)	0	
9	Solar azimuth (degrees)	0	
10	-- not used --	0	
11	First wavelength reflectance (%)	11.15	
12	Second wavelength reflectance (%)	11.05	
.	.		
N	N-10 th wavelength reflectance	##.##	
.	.		
438	428 th wavelength reflectance	25.27	
439	429 th wavelength reflectance	26.61	
440	430 th wavelength reflectance	0	<i>These last two measurements are filler in this example.</i>
441	431 th wavelength reflectance	0	

For USATEC format spectra HyperCube only needs the information in lines 3,4 and 5 plus the reflectance values.


An alternate format of a spectral signature is simply a list of wavelength and value pairs. The wavelengths and values can be in any units as long as the wavelengths are increasing. Each pair should be space or tab separated followed by a carriage return. HyperCube determines the format of a spectral signature by first trying to parse it as if it were in the USATEC or NEFDB format. If that fails it assumes it is wavelength, value pairs. An example of the later is shown below.

<i>Wavelengths</i>	399.000000	0.028000	<i>Values</i>
	402.000000	0.027000	
	405.000000	0.018000	
	408.000000	0.023000	
	412.000000	0.021000	
	415.000000	0.024000	
	419.000000	0.028000	
	422.000000	0.028000	
	426.000000	0.027000	
	429.000000	0.029000	
	433.000000	0.033000	
	437.000000	0.035000	
	440.000000	0.034000	
	444.000000	0.035000	
	448.000000	0.036000	
	452.000000	0.039000	


Plotting Signatures

Viewing a spectral signature requires first opening a library by choosing **Open As..., (Library)** from the **File** menu. This brings up the standard Mac file dialog showing only text type files. Select the library file containing the signature names (**Alibrary.lib** in the examples so far) and click **Open**. A new sizeable window containing a scrollable list of the signatures within the library will appear. Click on any signature name, select **Plot Selection** and a subsequent window, also resizable, will open showing a list of reflectances with the corresponding graphical plot. The average of two or more signatures can be plotted by double-clicking the desired signatures, setting **Average selections** and then the **Plot Selections** button. Double-clicking an item toggles its checkmark on and off. The plot range displayed for any particular plot can be modified by choosing menu **Functions -> Plot -> Scaling**.

Double-clicking an item toggles it on so that an average of all items with a checkmark can be plotted.




Scroll to each signature of interest, click on it to highlight it, and then click the Plot button.



The corresponding signature to the library list after clicking in the graph at the 1265 nm position.

Spectra plots have two check boxes that affect the appearance of the graph. The **Scale** option toggles the vertical axis (reflection) of the plot range. If unchecked the range is fixed to [0,1] else it is set to [min,max] of the reflectances of the viewable graph. The second check box, **10X**, expands the horizontal axis 10 times centered upon the current spectral marker, if any. Clicking within the graph area sets the marker and forces the list to scroll to the corresponding wavelength. Conversely, clicking the list sets the marker within the graph. Spectral plots may be printed by choosing **Print...** from the **File** menu and a snapshot of the window can also be copied to the clipboard. The list contents can be written to a file in the form of a library (see section **Spectra Libraries**) by selecting **Save** from the **File** menu or copied to the clipboard via **Edit -> Copy**.



Line 1 of the signature file.

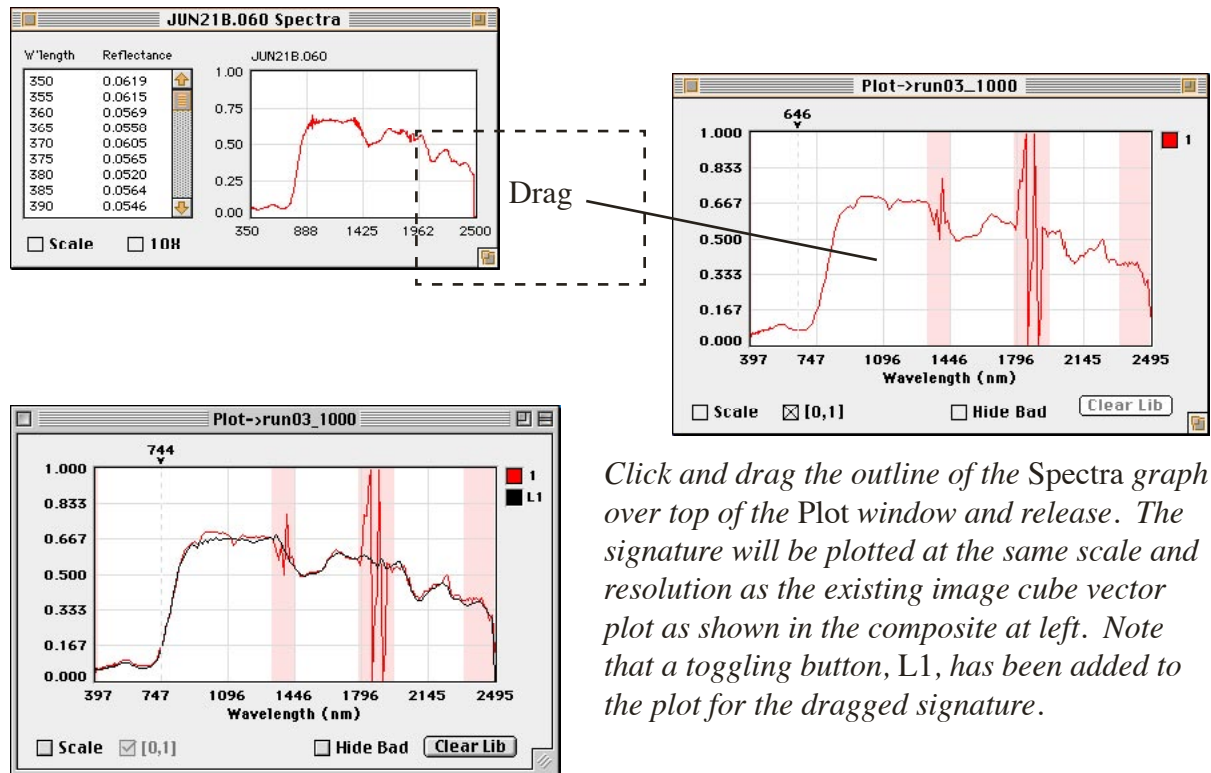
Click within the graph to set the wavelength marker; the matching list element will be highlighted.

Click and drag icon to resize.

The previous spectra plot example repeated with both check boxes set.

Superimposing Signatures

If HyperCube is running in system 7.5 or higher then the Mac's "drag and drop" feature can be used to overlay signatures on a common spectral plot. First open a multiband image and choose Functions -> [Plot -> Spectra](#) (Selected Points). Then click the mouse in a signature plot and while holding down the mouse drag the rectangular outline, which will appear, over the spectral plot graph and release the mouse. The signature will then be plotted in black superimposed with any selected cube points and will have its own toggle button. Up to 8 signatures can be "dropped" on the spectral plot window. The signature will be re-sampled to match the range of the multiband image.



Click and drag the outline of the Spectra graph over top of the Plot window and release. The signature will be plotted at the same scale and resolution as the existing image cube vector plot as shown in the composite at left. Note that a toggling button, L1, has been added to the plot for the dragged signature.

Spectral Classifications

General Algorithms

Spectral classification is the process of segmenting a multi or hyperspectral image into regions of similar spectral statistical characteristics or of finding the best match between pixels given a specific metric. HyperCube is designed to perform both. For example, Mahalanobis classify uses previously extracted statistical classes and decides which class, if any, a pixel belongs to, whereas Vector angle directly compares image pixel vectors with selected pixel vectors. The criteria of a match is the closeness of the comparison. This is a function of the algorithm used to define the closeness and HyperCube provides several from which to choose. No one algorithm is always better than the others. Some are very sensitive to noise, e.g., Mahalanobis distance, derivative difference, and others, e.g., Euclidean distance, are not. The best technique is to experiment.

Listed below are the algorithms from menu **Functions -> Classify** along with their arithmetic definitions. Summation extends over all enabled bands (see [Band List](#)).

Abs diff	$\sum \text{abs}(\text{reference} - \text{signature}).$
Corr coef	$\text{Covariance}(\text{reference}, \text{signature}) / (\text{sd}(\text{reference}) * \text{sd}(\text{signature})).$
Deriv diff	$\sum \text{abs}(\text{delta}(\text{reference}) - \text{delta}(\text{signature})).$
Euclid distance	$\text{Sqrt}(\sum (\text{reference} - \text{signature})^2).$
Fourier	Compare corresponding harmonic terms. See Options .
Fourier (no DC)	Same as Fourier excluding the zero (average) harmonic.
Fourier phase	Compare only the phase of corresponding harmonics.
Fourier phase delta	Compare the phase differences from harmonic to harmonic.
Kmeans	Unsupervised classifier that segments an image into a user specified number of classes.
Mahalanobis classify	<u>Statistical</u> difference between each image vector and the mean vectors of a set of covariance matrices.
Mahalanobis distance	Statistical distance between reference points and image points.
Mahalanobis list	A listing of the statistical difference between reference points and the mean of a set of covariance matrices.
Matched filter abundance	Generate a gray shade output image consisting of projections of the source data against a constructed matched filter of a signature(s).
Matched filter mask	Generate a mask over the image cube wherever the abundance exceeds a set tolerance (see description on next page).

Template	Max(abs(reference - signature)). Useful for finding very similar signatures.
Vector angle	N-dimensional angle between reference and signature.

Closeness is defined differently for each of the algorithms. **Euclidean distance** uses minimum as the criteria whereas the **correlation coefficient** approaches 1 for the best match. This is automatically taken into account by HyperCube but it is not possible to equate the numeric values of one method with corresponding ones in another method because the distribution of each method is totally different.

It is important to remember that a classification extends only upon bands whose status is **Enabled** in the **Windows -> Show Band List** and is spatially confined to any selection rectangle. The **Band List** makes it easy to test the contribution of a given band by running a classification, enabling/disabling the band, and rerunning the classification to see the effect. Any statistics required by the algorithm will automatically be recomputed.

Specific Algorithms

Matched filter abundance does not use a tolerance, instead it directly generates a data image where each pixel has a value approximately in the range $\sim[-2,2]$. It is possible to have numbers greater than 1.0. If the pixel used as the matched filter signature is also part of the classification region then it will have a value of exactly 1 at its position within the output image. The matched filter and abundance are defined as:

matched filter = $(\text{cov})^{-1} (\text{signature} - m) / ((\text{signature} - m)^T (\text{cov})^{-1} (\text{signature} - m))$
abundance = inner product of each data vector minus m and the matched filter.
where:

cov is the NxN covariance matrix of N bands of imagery

m is the Nx1 mean data vector

signature is the Nx1 pixel vector, N being the number of enabled bands.

Note that the denominator is the Mahalanobis distance between the signature vector and the mean vector.

When **Matched filter** is chosen as the classify **Method** the program must first compute the band to band covariance of the enabled file bands of the image cube or the corresponding subimage selection rectangle, if any. The determination of the statistics region is via **Classify -> Options -> Stats Region** (see Mahalanobis description, next). The computation is done once and the covariance matrix remains the same even if the current selection rectangle is moved or a different selection rectangle is drawn. This allows the statistics for the covariance to be generated for an image region other than the area chosen for the matched filter classification region. To force a recomputation of the covariance choose **Classify -> Options -> Redo stats**. Also, changing the state of enabled/disabled bands in menu **Windows -> Show Band List** will cause a recomputation of the statistics on the next classify.

Mahalanobis distance is a generalization of Euclidean distance that uses the statistics of the imagery. Specifically, the Mahalanobis distance squared between 2 vectors, X and Y is defined as:

$$D = (X - Y)^T (\text{cov})^{-1} (X - Y)$$

where:

D is the resulting scalar (i.e., 1xN times NxN times Nx1 matrices)

X and Y are each pixel vectors of length N

$(\text{cov})^{-1}$ is the inverse covariance matrix of N bands of imagery
and N is the number of enabled bands.

In HyperCube, X and Y are defined by the various Mahalanobis classify methods as:

Mahalanobis classify - X is each image pixel vector.

Y is the mean pixel vector of the region used to generate a covariance matrix. There can be one or more regions, each producing a mean vector and covariance matrix. Classification involves computing the Mahalanobis value of an image vector for each region and choosing the region (class) with the smallest value within a user supplied tolerance. Selection points are not used. Regions are defined by user selection of a mask image in **Classify -> Options -> Stats Region**. The statistics (covariance) are automatically generated by HyperCube based upon the mask definition.

Mahalanobis dist -

X is each image pixel vector.

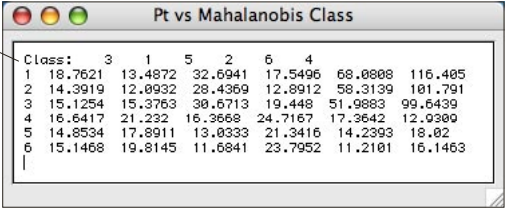
Y is one or more selection points.

The image statistics are the same as defined for **Mahalanobis Classify** except that the mean vector is not used and only one statistical region is used. This produces the statistical distance between one or more selection points and an image, not the distance between the statistical class mean and image.

Mahalanobis list -

Same mathematics as **Mahalanobis Classify** except selection points are used and the output is a numerical listing showing each selection point's distance from the mean vector for each of one or more classes.

Class names are those names that appear as the cursor is moved across the class map that was used to generate the statistics classes. Also, note their order. In this case the points themselves were used to create the class map so the minimum class matches the point number.



Class:	3	1	5	2	6	4
1	18.7621	13.4872	32.6941	17.5496	68.0808	116.405
2	14.3919	12.0932	28.4369	12.8912	58.3139	101.791
3	15.1254	15.3763	30.6713	19.448	51.9883	99.6439
4	16.6417	21.232	16.3668	24.7167	17.3642	12.9300
5	14.8534	17.8911	13.0333	21.3416	14.2393	18.02
6	15.1468	19.8145	11.6841	23.7952	11.2101	16.1463

Six selection points and their Mahalanobis value for each of six statistical regions.

The image area used for computing the covariance matrix depends upon **Classify -> Options -> Stats Region**. This displays the dialog shown below. There are 4 ways to define the regions. The first type computes the covariance statistics using all of the pixels within the image cube or a portion if a selection rectangle is present. The remaining three methods use an image mask. The pixels used for generating the covariance statistics are determined by an intermediate image mask that spatially matches the image cube. A simple way to generate a **Class map** image mask for the fourth method is by using the **Classify** function itself (e.g., select image points followed by method **Vector angle**). **Classify** can also be used to generate a detailed zero, non-zero mask by selecting image points, toggling neighborhood points and choosing **Form Neighborhood Mask Image** in the **Classify** dialog.

Statistics are automatically computed when the **Classify** button is executed (section [Classify Function](#)) if there has been any change in the **Classify -> Options -> Stats Region -> Select Image Mask** selection. They remain in effect so that a different area on the image cube can be used for classification based upon the same statistics. Also, any enable/disable change in a band list (see section [Band List](#)) will cause a re-computation of the statistics. Button **Classify -> Options -> Redo stats** forces a re-computation and is useful if the selected image mask was generated as a class map sequence (see section [Classify Options](#)). Then the current class map is taken as the sequence cube face and can be used as follows: Select a new face, **Redo stats**, and then **Classify** to see the effect of a mask change within the sequence (saves having to repeatedly go to **Classify -> Options -> Stats Region -> Select Image Mask** each time the sequence face is changed).

Some additional considerations: the covariance matrix is very sensitive to noise. Noise, being defined here, is anything that is not representative of that class. Even a single outlier, such as striping in the image, will skew the statistics and produce a non representative covariance matrix. If the imagery is noisy then try using fewer widely spaced bands (see section [Band List](#)). If a two state mask image is selected or in use and **Mahalanobis Classify** or **List** is chosen as the **Classify Method** then the outputs reduce to one statistical class.

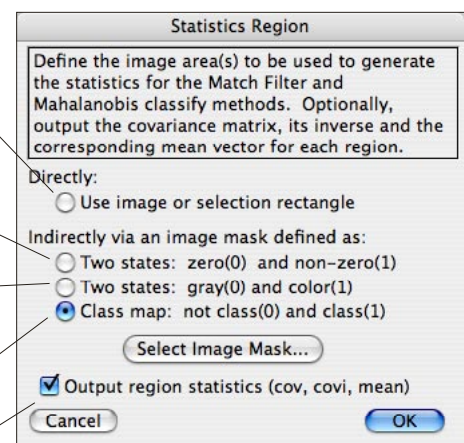
The covariance, its inverse and the mean vector for the selected region(s) can be displayed by checking **Output region...** in the **Statistics Region** dialog. If already computed then check **Redo stats** in the [Classify -> Options](#) dialog. If fewer than 16 bands are enabled then the matrix results will be text windows else they will be image windows. The matrices and plots can be saved or copied to another program.

Use every pixel within the image or selection rectangle.

Skip image pixels wherever the mask has a zero, use image pixels wherever mask is not zero. Produces a single covariance statistics class.

Skip image pixels wherever the mask is gray (including zero) and use image pixels wherever mask is color. Also produces a single covariance statistics class.

Mask must be a class map (section [Image Products](#)). Each color produces a covariance statistical class.



Output covariance, its inverse and mean vector each time new statistics are generated.

K-means is an unsupervised classifier that does not use any selection points and requires minimal input. The algorithm segments the image into n classes by initially guessing the mean vectors of the n classes. It then scans the image and determines (using one of the four classification metrics listed below) which class each image vector belongs to and a corresponding set of new mean vectors. This procedure is repeated until the number of image vectors that changed classes is below a certain minimum or a maximum number of iterations has occurred. The number of classes and both termination criteria as well as other parameters are user selectable. Choosing **Options** in the **Classify** dialog while the classify **Method** is **Kmeans** displays the dialog shown below (note: the **Options** dialog associated with **Kmeans** is different from the options shown for the other classify algorithms).

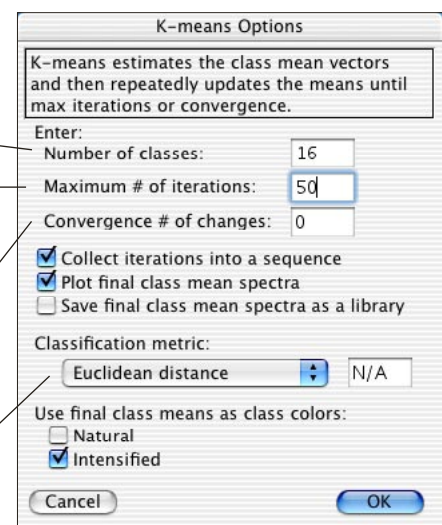

See section [Image Arithmetic](#) for an example of combining a **K-means** classification and the gray values of a specific band using the `hsi()` function.

The number of desired classes is limited to 500 or less. The actual number of classes generated may be less than the number requested. Use menu Functions -> Plot -> Histogram on the classify output to list the classes.

A Progress window dynamically displays the iteration count and the number of class changes. When the number of iterations reaches this value Kmeans classify stops and outputs the results up to this point.

When the number of class changes is less than or equal to this value classify terminates.

Various metrics popup menu.



Remaining options:

Collect iterations: Each iteration (classify) is assembled into a class cube that can be scrolled to view the convergence sequence.

Plot final class... Each mean vector forms a spectral plot window.

Save final class... The final class mean vectors are saved as a spectral library.

Classification metric The algorithm that is used to determine which class an image vector belongs to by comparing it the mean of each class.

Use final class means as class colors:


Natural - Determine mean class colors directly from mean vectors.

Intensified - As above, but intensity the color.

Red, green, blue wavelengths are determined by menu
Image -> Cube Color Comp -> Definitions.

Classify Function

The HyperCube vehicle for spectral classifications is initiated by choosing the **Functions -> Classify** menu and is available whenever a cube or a linked color band composite window is front. Each image cube can have at most one **Classify** window. A classification scenario has three basic steps: decide on the signature to be matched (what to look for), choose the search domain (where to look), and pick the method (how to look). The output product is determined by the domains and numerous options. The simplest classification is to shift-click on a point in the image cube window, double-click on it in the **Classify** list, choose one or more neighborhood points and generate an overlay image of the current **Search Domain** cube face band using **Vector Angle** as the **Method**. This is the default setting of the **Classify** window and does not involve a library. A **Classify** window is logically connected to its image cube and closes when the cube window closes.

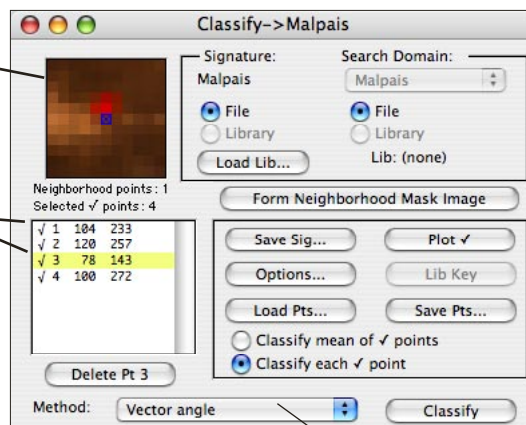


As you select cube points their coordinates are added to the **Classify** list. These points are then used to reference one or more surrounding neighborhoods from which you can toggle one or more signatures to be averaged together. The list coordinate is the central pixel in the neighborhood and has the same color as the cube selection point reticle but with a transparent interior. It is also the same point that is graphed if a plot window is open (see **Spectral Plots**). By clicking the mouse within the neighborhood, points may be added or removed from the signature mean. Shift-clicking adds or removes a rectangular region between the shift-click and last non shift-click point. Any combination of neighborhood points is allowed. If more than one pixel is toggled-on then their mean is formed. Multiple neighborhoods are selected by double-clicking each list point number. Clicking the **Plot Mean** button produces a graph of the resulting signature and this graph may be superimposed upon an open spectral **Plot** window (see **Superimposing Signatures**).


The neighborhood image is generated from the band currently shown on the **Signature** cube window and therefore will change as the cube is band scrolled. However, if a linked color image has been constructed from cube bands (see sections [Components to Color](#) and [Color Composite](#)) then the neighborhood is extracted from the linked color composite.

An 11 by 11 neighborhood surrounding the selected image cube coordinate.

All four list points have been selected as the neighborhoods in which additional points may be toggled-on to form a mean signature for each neighborhood. Clicking once on point number 2 highlights it and displays its neighborhood.



The Search Domain does not have to be in the same cube as the Signature which is fixed to the selection point cube.



The following is a description of each **Classify** window control.

Signature: The signature source is always taken from the point selection cube and its name is shown beneath **Signature**. The only exception is when **Library** is chosen (see below).

File Use a selected signature from the image cube.

Library Use a selected signature from the current library as the signature.

Search Domain: The popup menu beneath this title is used to select the cube in which to search. The **Search Domain** does not have to have the same spatial or spectral dimensions as the **Signature** source. A different cube is only permitted when **Signature** and **Search Domain** are both set to **File**. Also see **Library** below for another exception.

File Match the selected signature against image cube signatures.

Library Match the selected signature against all signatures in the library.

Method: Popup menu of available search algorithms.

Library... Open a standard file dialog to select the library to use. The name of the library will replace **Lib:(none)** in the **Classify** window.

Lib Key Active after a library load. Generates a window showing the correspondence between a toggled-on library item and the color that will be used in the classification. See the later section on **Products** and earlier section **Changing Selection Point Colors**.

Form Neighborhood Mask Image:

A separate gray image is generated spatially matching the current selection area and having 255 gray values at each toggled selection point and neighborhood with zero elsewhere. See **ROC Curves**.

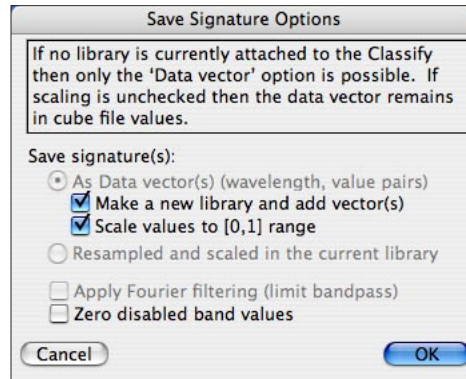
Plot \checkmark Generate a spectral signature plot of toggled-on neighborhood points. If **Classify each \checkmark point** is checked then the plot will contain each neighborhood and their mean. If **Classify mean of \checkmark points** is checked then one mean of all the selected list points and their neighborhood points is formed and plotted. If **Method** is **Fourier** then the plot is pre-filtered.

Options... Open the **Classify Options** window.

Delete Pt # Delete the highlighted selection point from the list and move lower points, if any, up. Also affects any current plots.

Load/Save Pts... Save the current list of selection points including their neighborhood definitions. Upon loading all points are automatically toggled on.

Save Sig... If **Signature** is set to **File**, then the currently toggled on selections can be saved as a spectral signature of raw values, as a new library or added to the currently opened library. The following dialog is presented.



Classify mean of $\sqrt{}$ points
Classify each $\sqrt{}$ point

These form a pair of mutually exclusive radio buttons that determine how the toggled-on ($\sqrt{}$) points are used in the classification process. The default is **Classify each**. Then each $\sqrt{}$ point generates a classification region in competition with the others. That is, cube pixels are compared with each point and the closest point is defined as the proper classification. The program requires that the classification criteria be set to **Threshold** in the **Classify Options** dialog. If no point matches within the **Threshold** then that pixel remains unclassified (see the next sections, **Options & Products**). If **Classify mean** is chosen then all of the $\sqrt{}$ points are averaged together to form a single mean signature that will produce a single classification region. Any classification criteria may be chosen in **Classify Options**.

Note that if there is only one $\sqrt{}$ point then either of these choices will produce the same result if **Threshold** is the criteria. Also, any $\sqrt{}$ point that has multiple neighborhood points selected will still form an average of that neighborhood to represent the signature of that $\sqrt{}$ point regardless of the choice here.

The color that is used to delineate classification regions is the same as each selection point color for **Separate** and the color of the first selection for **Mean**. The colors can be changed using the **Color Selector** described in section **Changing Selection Point Colors**.

Classify/ROC Curve:

Perform either a classify or ROC curve. This button toggles between these two as a function of the **ROC curve** checkbox in **Classify Options** dialog.

Classify Options

Pressing the Options... button in the Classify window opens the following window.

Each classify Method has a theoretical mathematical range and is shown along with any parameter that must be set for that Method.

There are 4 ways to compare a signature within the search domain. If Classify each is set then the criteria will be forced to Threshold.

The output format is determined by these radio buttons and the checkboxes.

The number of signatures to list when Search Domain is a library.

If enabled then each spectral signature is filtered based upon its energy distribution before being used in a classification or plot.

The screenshot shows the 'Classify Options' dialog box. It has a title bar 'Classify Options'. The first section is 'Selected Method Parameters' with a dropdown menu showing 'Mahalanobis (no param)' and a text field with 'N/A'. There is a 'Redo stats' button and a 'Stats Region...' button. The second section is 'Classification Match Criteria' with four radio buttons: 'Threshold, Enter tolerance, [best,worst]:' (selected), 'Percentage, Enter tolerance, [0,100%]:', 'Equal distribution, Enter classes, [2,64]:', and 'Countable, Enter N closest, [1,32767]:'. The third section is 'Classification Output(s)' with three radio buttons: 'Color mask (black where no match)', 'Color image (cube face shows where no match)' (selected), and 'Generate a sequence, Enter # of outputs:'. There are also checkboxes for 'ROC curve' and 'ROC Mask...'. The fourth section is 'Library Parameters' with a dropdown menu showing 'N/A' and a text field with 'Closest library matches'. There are also checkboxes for 'Lib:File match lists', 'Ignore the Disabled Bands List in Lib:Lib searches', and 'Enable dynamic display in File:Lib searches'. The fifth section is 'Pre-filter' with a checkbox for 'Enable spectral wavelet denoising' and a 'Parameters...' button. At the bottom are 'Cancel', 'More Options...', and 'Okay' buttons.

Selected Method Parameters:

This section shows the currently selected classification Method and allows you to change a specific parameter associated with that method. Presently, only Deriv diff and the various Fourier methods have such a parameter and are described below. The Theoretical best and worst tolerances refer to the mathematical closest and furthest difference between the signature and search domain vectors. These are used to set thresholds and sequence limits and vary depending upon the algorithm (Method). For instance, Vector Angle and Fourier Phase range from 0 to π radians whereas Corr Coef ranges from 1 (best) to -1 (worst). The various metric distances such as Abs Diff are normalized to the range [0,1] regardless of the arithmetic range of the search image cube. Euclidean distance is normalized to the sqrt of the number of enabled bands. The Methods which allow parameters are described below.

Deriv diff delta: The sample distance from signature point to point used to compute the derivative. Generally, signatures are quite noisy and using a delta greater than 1 provides a better indication of the true slope. However, if there are few or widely spaced bands some important signature variation may be skipped.

Fourier terms: The low order number of Fourier transform harmonics to be used in comparing signatures in the frequency domain. The greater the number of terms the finer the discrimination and sensitivity to noise. The smaller the number the greater the filtering and immunity to noise. A good starting point is 8 and this value must not exceed the number of enabled bands divided by 2.

Redo stats: Available with classify methods Mahalanobis and Matched Filter.
Stats Region: Used to define the region for generating needed covariance statistics. See description in section Algorithms.

Classification Match Criteria:

The four mutually exclusive radio buttons in this area specify how the mathematical difference between the source signature spectra (vector) and the search domain spectra is used. The first one, **Threshold**, is the most direct. It allows you to enter a tolerance in the range given by the **Selected Method Parameters** section just described. Then, each computation is compared with this tolerance and if it is between it and **best** then the image cube pixel is defined to be classified as the source spectra, otherwise not. If more than one point is being classified, i.e. **Classify each**, then the closest point to the best within the tolerance, if any, determines the classification.

Threshold: A value in the range: [best,worst] that serves as a breakpoint. Computations within [best,value] are classified, outside this range are not. Example: if the **Method** is **Vector Angle** then the tolerance range will be listed as 0 (best) and π (worst). Setting the tolerance to .08 means that all vector angle differences between 0 and .08 will be considered to be “classified”. The color given to a classified region is the same as the selection point color and may be changed via the **Color Selector**. **Threshold** is the only criteria possible if **Classify each** is selected.

Percentage: Generate a color image using the selection point color to indicate the closest percentage of the entire distribution of the match. For example, if the percentage is 10% then the upper (best) 10% matching pixels are defined as classified.

Equal Distribution: Generate an N color output image where N is chosen by you. The entire match distribution relative to the source spectra is divided into N equal areas. Each area then becomes a classification region. Unlike the other criteria, the closest region has the color of patch 1 in the **Color Selector**, the second closest region has the patch 2 color and so on.

Countable: Generate an output image with the N closest matches delineated. For example: if N equals 50 then the 50 closest matches will define the classification region. Good for showing very small targets.

Classification Output(s):

This section defines the form and number of output products. There are 2 types and both produce an image. Each generates the appropriate color classification regions for the **Match Criteria** listed above. Their difference is how an unclassified region appears. The **Color Mask** leaves these areas black while the **Color Image** shows the current cube face band where there is no classification. Both selections produce a true color image. Note that the colored regions correspond exactly to the color currently defined for points in the **Color Selector**. If the colors don't seem appropriate for the given classification change them in the **Color Selector** and redo the classification. Also, see the color note above regarding **Equal Distribution**. **Flicker/Superimpose** is very effective in dynamically comparing a classification output with a source image (i.e. the linked **Color Composite** or an image cube band).

- Color Mask:** The output image has a color where a classification is define as successful and black otherwise. The only output possible if **Equal Distribution** is selected.
- Color Image:** A copy of the current cube face band will be used as a base image and the classified regions will be overlaid.
- Generate Sequence:** If checked then the 2 tolerance values entered below the checkbox define a range of **Match Criteria** in which a sequence of classification images will be generated. The number of outputs is user selectable. The outputs are arranged into a synthetic cube for display (see **Products**). The example below will generate a synthetic cube of 10 images where image 1 represents a classification tolerance of 0 and image 10 a tolerance of .5 (in this case radians).
- ROC curve:** This option is available if **Classify mean of $\sqrt{}$ points** and **Generate a sequence** are both selected. No image output is

- generated, instead a plot of the number of classify hits versus tolerance is generated. See **ROC Curves** section for details.
- ROC Mask...** Displays a popup menu dialog allowing you to select the “ground truth” mask image to be used in generating ROC curves.

Library Parameters:

This last section is applicable only when the **Search Domain** is **Library**. In this case the output is not an image but ordered sets of matching library signatures.

- Closest lib matches:** You can specify how many library signatures to list (see the **Products** section).
- Lib:File match lists:** Generates a list of the library elements found in the image and a list of the source library elements that didn’t match.
- Ignore Disabled bands:** Available when **Signature** and **Search Domain** are both set to **Library**. In this case you may want to perform matching at every spectral band value by including any wavelengths that are disabled in the **Band List**.
- Enable Dynamic Display:** Available when **Signature** is **File** and **Search** is **Lib**. As the cursor is passed over the Cube or Link image a real time display of the closest library matches is dynamically displayed. If a dynamic spectral plot is enabled, then the closest 2 library spectra will be plotted.

Pre-Filter:


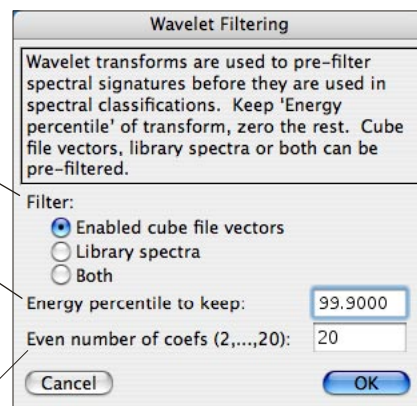
Enable wavelet spectral denoising: If checked then each enabled cube file spectral signature is filtered using the discrete wavelet transform.
The parameters are determined by the dialog window shown below.

Parameters... Displays the following dialog.

*Filtering can be applied selectively:
only file spectra, only library spectra or both.*

The percentile of the total energy of a spectra in the wavelet transform domain to keep. Generally, this number should be very close to 100%. By displaying a cube selection point and enabling filtering you can plot (Classify:Plot) and drag the filtered signature onto the selection point plot window as a comparison.

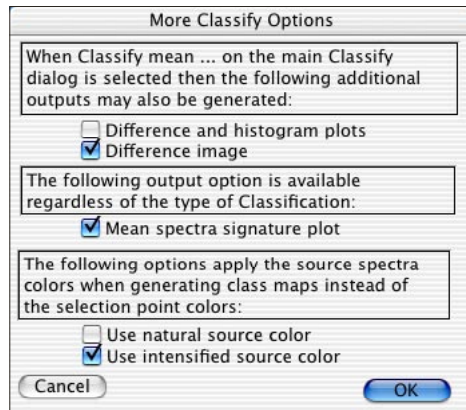
Specifies the number of Haar (2) or Daubechies (4,...,20) coefficients to use in the wavelet transform. The greater the number, the smoother the filtering.





This plot, [Functions -> Plot -> Spectra \(Selected Points\)](#), displays a cube selection point (red) versus the same signature (black) used for [Classification](#) with Enable wavelet... checked (see section [Superimposing Signatures](#)). The wavelet filtering values applied are those listed in the dialog window above. Disabled bands, highlighted in red, are not filtered and therefore remain the same in both signatures.

More Options



When performing a classification you may request several output products that can be generated in addition to the class map. By selecting **More Options...** in the **Classify Options** dialog the following dialog appears.



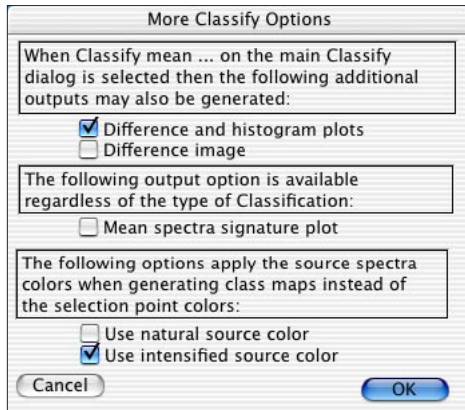
Three options have been selected: produce a difference image between the selected classification point and each spectral signature within the image cube, compute the average of all spectral signatures within the image, and use the source color (determined by menu Image -> Cube Color Comp -> Definitions) of the selection point as the class map overlay color. The mean spectra does not depend upon the classification points whereas the difference image does.




A single selection point within the delineated area shown on the left is used to produce the overlay type classmap appearing to the right.




*The plot of the mean of all 47,872 spectral signatures in the classification area. Bad bands are zeroed. A plot that includes bad band values can be generated by choosing **Enable All** in the **Band List** and then reclassifying. The difference image (shown right) is coded so that the brightest values represent the closest match to the chosen selection point(s) for the given classification method. However, the data attached to the difference image gives the actual metric difference (e.g. radians for Vector angle) as the cursor is moved over the image and the Info Window (menu Windows->Show Info) is active.*



The selected option will produce two plots: The actual value produced by the Method used in the classification for each pixel and a histogram that plots a relative frequency count of how often a specific difference occurred. Any section of each plot can be magnified by dragging a selection rectangle about the area of interest and then choosing the Image -> [Magnify](#) -> 1x menu. The magnification factor is ignored, instead, the selection rectangle is expanded to full size.



Each of the 49,970 spectral signatures is plotted against the Vector Angle difference between it and the chosen classification point shown on the previous page. The maximum difference is 1.22 radians. By dragging a selection rectangle anywhere within the plot and choosing Image->Magnify->2x a new plot consisting of just that region will be constructed. This may be repeated on the new plot to get even finer detail.



This Histogram is essentially the same plot as the Difference plot reordered so that the horizontal axis represents Vector Angle differences (in radians) verses how frequently this difference occurred. This plot may also be magnified repeatedly.

Image Products


Each **Classify** produces an image(s), an ordered set of spectral file names or an ROC plot. The image can be a color representation of the goodness-of-fit or a composite consisting of the color representation overlaid on an image background. The background is automatically chosen to be the **Search Domain** band face so for example if band number 35 is visible, a duplicate is made and used as the overlay base. The colors of the mask/overlay are the same colors as $\sqrt{}$ points in the classify list. If a selection rectangle is present on the **Search** cube face then the classification and resultant image is limited to the rectangle size.

An output image has one or more classification regions depending upon the various options described in the previous section. The default is one region per $\sqrt{}$ point or $\sqrt{}$ library spectra. However, if the **Classify mean** radio button is chosen then the average of all $\sqrt{}$ points is formed for the source signature for comparison and therefore only one classify region is generated. Shown below is a source image with four selected points, and 2 classification outputs using the default **Vector Angle Method** with tolerances .08 and .8 radians respectively. Each region color corresponds to the selection point color and can be changed via the **Color Selector** (see section [Selection Point Colors](#)) or Use source spectra color in **More Options**.

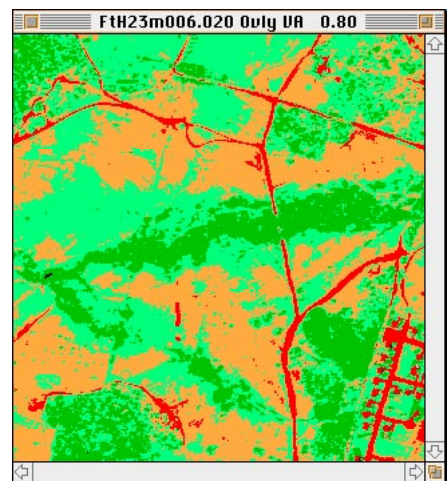
The utility of the classification output can be greatly expanded by using **Flicker/Superimpose** in the **Functions** menu or compositing it with another image via **Overlay Color** or using image **Arithmetic**. Furthermore, classifications can have labels that dynamically change as the cursor passes over them (see section [Dynamic Displays and Plots](#)).



Four selected points on a linked color composite.



The top right output is the result of using the Vector Angle Method with a tolerance of .08 radians. The base image cube shows through where the angle exceeded the tolerance. The right image is the same classification re-run with a tolerance of .8 radians which just about forces each pixel into one of the classification regions thus obliterating the underlying image.



One of the more interesting products of **Classify** is the generation of a synthetic cube consisting of a linear sequence of classification tolerances. For example, the **Classify Options** dialog shown below will generate 11 outputs ranging in tolerance from 0 to .1 radians that will be automatically collated into an 11 band synthetic cube also shown below. Besides band scrolling, most of the HyperCube functions that can be applied to a spectral image cube can also be applied to the synthetic cube, such as: [Gray & Color](#) mapping, [Flips, Rotates](#), [Flicker/Superimpose](#), [Magnify](#) and [Filtering](#). The exceptions are Shifts, Histograms and Statistical Measures.

The **Classify Options** dialog box is shown with the following settings:

- Selected Method Parameters:** Vector Angle (no param) is set to N/A. Theoretical best Tolerance is 0.00, and Theoretical worst Tolerance is 3.14.
- Classification Match Criteria:** Threshold is selected with a tolerance of N/A. Other options like Percentage, Equal distribution, and Countable are also listed with N/A values.
- Classification Output(s):** Color image (cube face shows where no match) is selected. The option "Generate a sequence, Enter # of outputs:" is checked and set to 11. Starting tolerance is 0 and Ending tolerance is .1.
- Library Parameters:** N/A Closest library matches is selected. Other options like "Ignore the Disabled Bands List" and "Enable dynamic display" are unchecked.
- Pre-filter:** "Enable spectral wavelet denoising" is unchecked.

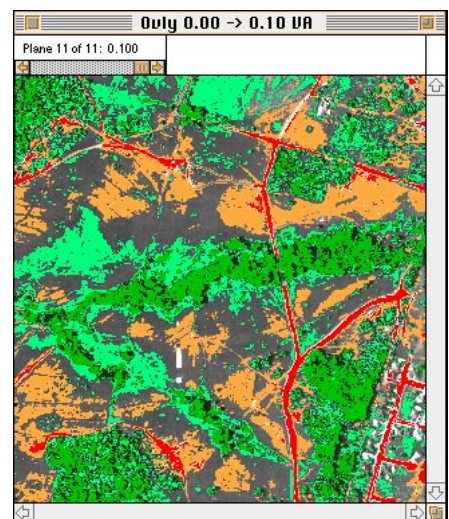

Vector Angle is the chosen Method.

The number of outputs in the synthetic cube. In this instance the output sequence will be: 0, .01, .02, .03, .04, .05, .06, .07, .08, .09 and .1 radians tolerance.

Starting tolerance of the first output.


If checked then a plot of "hits" versus "false alarms" will be generated.

Ending tolerance for the last output.



Shown left and right respectively are two views of the resulting synthetic cube. The left one has been scrolled to the fifth output band in which the overlay tolerance is .04 radians while the right view has been scrolled to the last band showing a tolerance of .1 radians.

The classification regions of a classify image can be tabulated by selecting menu Functions -> Plots -> [Histogram](#) when the image is active. The result is list of each class name and the frequency with which it appears in the image. Zero frequency classes are omitted. The list can be selected (normal drag selection and command-C key) and pasted into another program, e.g., Excel®. Below is a portion of a very detailed Lib:File classification and the histogram of the selection rectangle area within the image.



#	Count	Name
5	4	BG-1, asphalt, light
6	1	BG-2, asphalt, light
13	16	BG-1, roadway, packed
14	1	BG-1, roadway, gravelly
15	1	BG-2, roadway, gravelly
16	186	BG-3, roadway, gravelly
17	165	BG-4, roadway, gravelly
19	36	BG-6, roadway, gravelly
22	11	BG-1, PW3b, sand
23	144	BG-2, PW3b, sand
24	15	BG-3, PW3b, sand
25	39	BG-4, PW3b, sand
26	29	BG-5, PW3b, sand
27	17	BG-6, PW3b, sand
28	112	BG-1, F1c, gravel
29	302	BG-2, F1c, gravel
30	3	BG-3, F1c, gravel
31	343	BG-1, T1c, gravel
32	97	BG-2, T1c, gravel
33	50	BG-3, T1c, gravel

At the left is a portion of a classmap generated via a Lib:File classification. The library contained several hundred signatures. At the right is part of the histogram frequency count listing for this classmap. Each library signature has multiple measurements. The first column contains the index of the library signature in the color key. See sections: [Library Searches](#) and [Class Map Editor](#) for a description of the class name and color association.

F(Signatures)

Whenever the **Classify** dialog is active you can construct the eigenvectors and/or orthonormal basis of the selected signatures via menu item **Functions -> F(Signatures)**. Either image or library signatures must have been toggled on and **Classify each...** must be selected. If from a library they are resampled to the enabled bands of the attached image cube. **Eigenvectors** requires 2 or more signatures while **Orthonormal** 1 or more. The covariance or the correlation of the signatures may be chosen for the eigenvectors computation. Since the number of signatures is usually much less than the number of bands the resulting eigenvalues and hence eigenvectors have little significance after the N-1 value where N is the number of signatures used. Compare this with the HyperCube menu item **Image -> Convert Image -> Principal Components** that uses all of the pixels within a cube (or selection rectangle) to compute the covariance. **Orthonormal** uses the Gram-Schmidt method and the number of signatures can not exceed the number of bands. **Orthonormal** also produces the corresponding projection matrix defined as: $A(A^T A)^{-1} A^T$ where A is the column matrix of orthonormal vectors. Both methods generate matrices that are saved to a file and can be used in **Utilities -> Apply Transformation**.

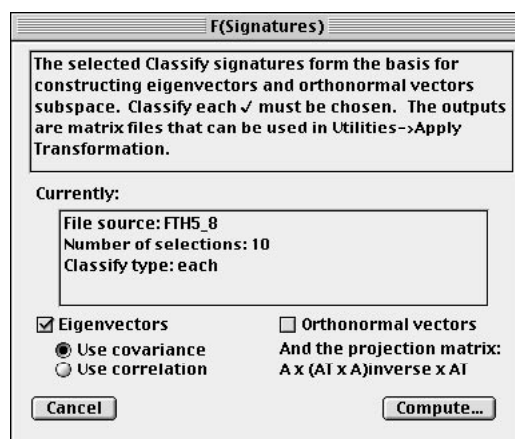
Mathematically, the following matrices are produced for the particular instance of an image cube with 160 of 210 bands enabled and 10 signatures.

Eigenvector matrix (Matrix.ev): 160 rows x 160 columns.

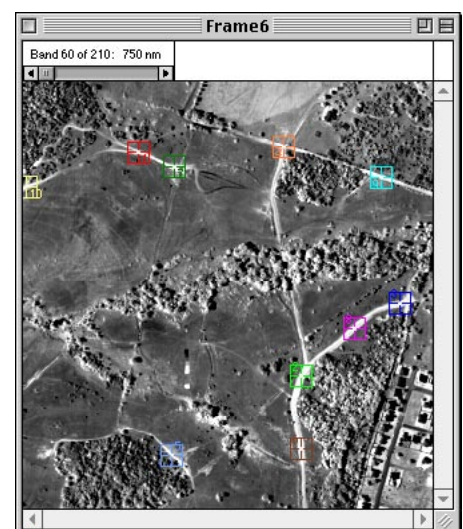
Orthonormal matrix (Matrix.gs): 160 rows x 10 columns.

Projection matrix (Matrix.pm); 160 rows x 160 columns.


Below and on the following page are examples of the **Eigenvectors** method and the result of applying the matrix to the original image cube using menu **Windows->Apply Transformation**.



F(Signatures) dialog set to compute the eigenvectors corresponding to 10 Classify signatures. The result is a matrix file.




Classify selection points.



First principal component.

The left and right images are the first two bands of the cube resulting from applying the eigenvector matrix to the original data using [Utility Apply Transformation](#).



Second principal component.

ROC Curves

One type of classification output is a ROC (Receiver Operating Characteristic) curve. This produces a plot of the probability of correctly classifying a data set (statistically not committing a Type I Error) verses the probability of erroneously classifying a data set (statistically of committing a Type II Error). This relationship is commonly referred to as the “hit rate” verses the “false alarm” rate as a function of a varying decision criteria. In HyperCube the criteria is the tolerance for a given classification algorithm. The validity of whether or not a signature is classified correctly is determined by a previously constructed “ground truth” image mask matching the classification area. A mask pixel of zero means there shouldn’t be a match and a nonzero pixel means there should be. During the ROC computation as each image point is classified (matched against a given spectral signature) a tabulation is maintained of whether an image point was classified (within tolerance for the given algorithm) and if the classification was correct according to the “ground truth” image mask. Each tolerance value produces one tabulation. A sequence of linearly spaced tolerances (~250-500 values) produces tables that are then normalized to form the ROC plot.

The ROC option is available when **Classify mean of $\sqrt{}$ points** is selected in the main [Classify](#) window along with **Search Domain** set to **File** and the **Generate a sequence in Classify Options** checked. It is similar to generating the classification image sequence previously described. After checking the **ROC curve** box the **ROC Mask...** button is clicked to select the ground truth image mask from a popup list of images. This image must spatially match the classification area used in generating the ROC sequence of classifications. Note that “**Classify mean of $\sqrt{}$ points**” always results in a single classification region so the summation is over a single class.

Listed below is a representative scenario for generating a ground truth mask from an image cube and then forming the various ROC plots shown on the following pages.



1. Choose several image points representing the same classification feature.
2. Select one or more neighborhoods around each point (to form a good average).
If these points exactly delineate the ground truth then click the **Form Neighborhood Mask Image** button and **Classify mean of $\sqrt{}$ points** and continue at step 7.
3. Choose an algorithm such as **Vector Angle** and **Classify mean of $\sqrt{}$ points**.
4. Choose **Classify Options** and set up a sequence of ~15 outputs and **Color Image**.
5. Click **Classify**. Choose the output within the sequence of images that best represents ground truth for this feature. Select **Convert Image -> Image to Mask** menu to extract the mask. This mask can then be further edited (e.g. PhotoShop®) to give it more validity.
6. Clear all of the selection points that were used to generate the mask and select a new point from the cube that represents the feature in the mask. Or, open a library and select a signature from the library list (see section [Library Searches](#)) and select **Library** as the **Signature** in the **Classify** window.
7. Choose an algorithm from **Methods** in the **Classify** window. Select **Classify Options**, check **Generate sequence** and **ROC curve** and then **ROC Mask...** .
Select the mask image generated above from the popup (this only needs to be done once). Return to **Classify Options** and enter the number of outputs (~250) and a broad range of tolerances for the chosen algorithm. Return to the **Classify** window.
8. The **Classify** button will change to **ROC Curve**. Click it to generate the plot.
9. Repeat steps 7 and 8 for each algorithm. Step 6 may also be repeated to see its effect.



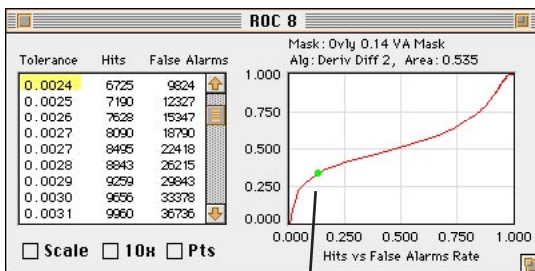
The mean of six points representing a typical tree canopy was used for this image mask that was extracted from a sequence of 15 classifications using Vector Angle.



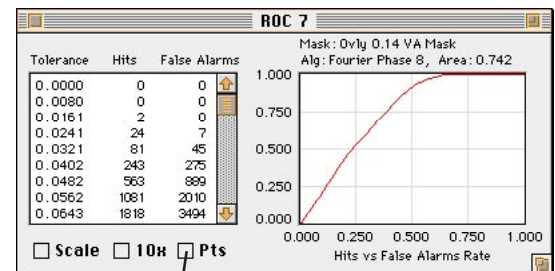
The selected image point on this linked color composite was the classification signature to be matched when generating the ROC plots shown below.





From above clockwise are the ROC plots corresponding to Vector Angle, Euclidean Distance, Fourier Phase and Derivative Difference *respectively*. The vertical axis is the hit rate and the horizontal axis is the false alarm rate. The greater the area under each curve the more discriminating the algorithm. Ideally, the Hit Rate would approach 1 while the False Alarm Rate stays near zero producing an area close to 1.





Clicking within the graph displays a green dot and highlights the corresponding tolerance value. Conversely, clicking within the scrolled list shows the corresponding plot position.





The Pts checkbox switches the plot from continuous to a discrete scatter plot which the ROC actually represents.



Fifteen selection points along with various neighborhood points were used to generate the resulting selection mask shown to the right by choosing Form Neighborhood Mask Image in the Classify dialog above (see section [Classify Function](#)). Each target point represents the same physical material. The library spectral signature of this material was used in the classification to produce the four ROC curves shown below. Note that the classification area and the mask selected for the ROC curves must spatially coincide which in this case was automatic.



These plots correspond to the same four methods shown on the previous page that were used in discriminating the tree canopy. Here they give a quantitative measure of target versus background discrimination.



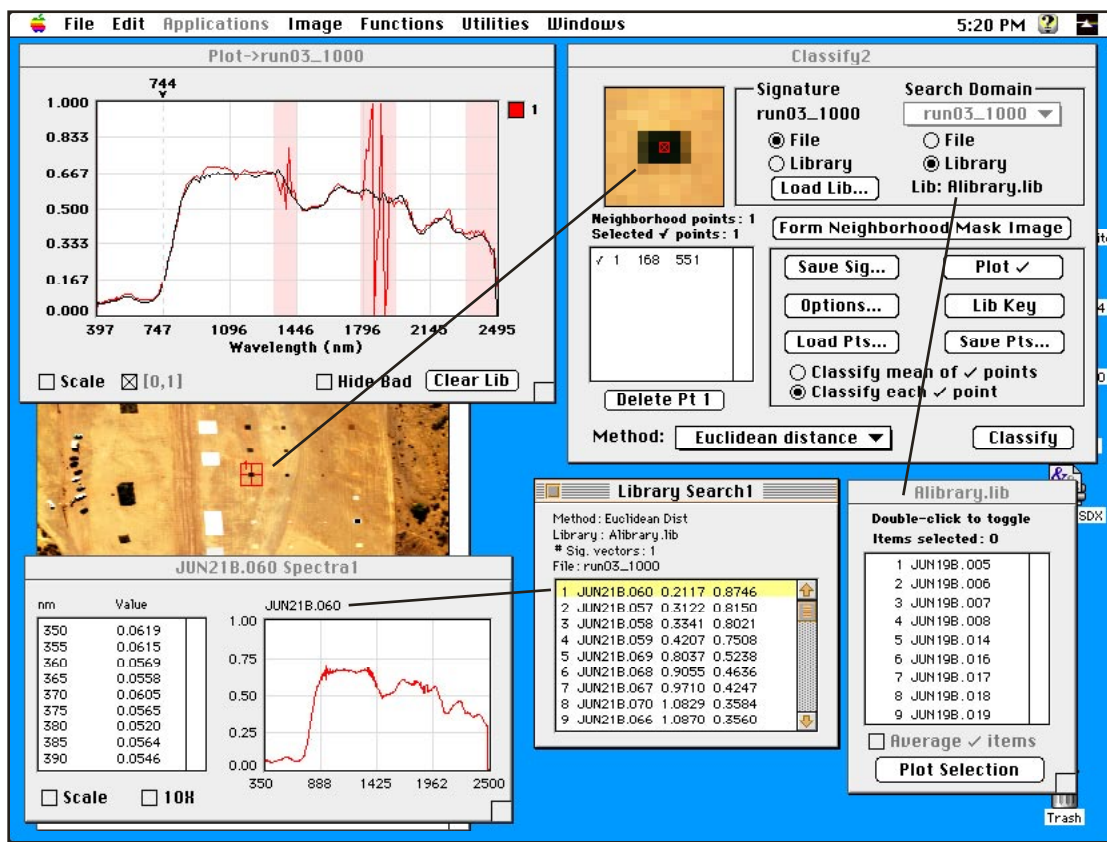
Library Searches

Choosing **Library** as the **Search Domain** produces a list of library signatures that best match the criteria selected by the user. The list is ordered in decreasing closeness with a numerical value derived from the classification **Method**. The numbers are only meaningful within the same method and are most useful in detecting a cluster of matching signatures, i.e. items that are physically similar. The spectrum of any listed signature can be plotted in a spectral signature window by double clicking it and can be added to the spectral plot window by dragging its graph (see the [Superimposing Signatures](#) section).

A library is loaded by choosing the **Library...** button in the **Classify** window. Once loaded the name will replace **Lib:(none)** and the **Library** option radio buttons under **Signature** and **Search Domain** will then become available. Loading a library connects it to the classification process whereas using menu **File -> Open As...**, and selecting **Format: Library** does not and that library remains independent. However, signatures from the later may also be dragged onto the spectral plot window.

The steps to perform a typical library classification are as follows:

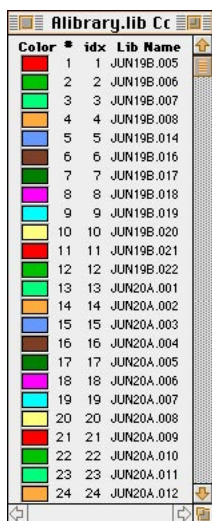
1. Load the image cube file and select one or more (if averaging) image points.
2. Choose [Classify...](#) from the **Functions** menu.
3. Click the **Library...** button to load a spectral library.
4. Select (✓) one or more cube points from the **Classify** coordinate list.
5. Click the **Signature: File** and the **Search Domain: Library** options.
5. Choose a classification from the **Method** popup menu.
6. Click the **Classify** button. A list(s) of matching library signatures will appear.



The Library Search window lists the library signatures matching the selected image cube file point. The top signature in the list has been superimposed upon the Plot window.

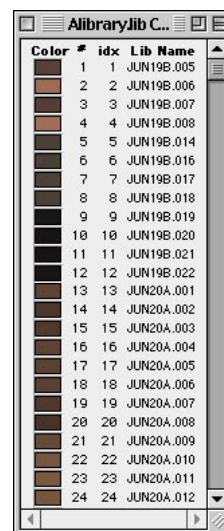
The opposite of the previous example is also possible: i.e., choosing one or more library signatures and generating a classification output (Lib:File search). The region colors are then taken from the patches in the **Color Selector** as follows: the lowest $\sqrt{\text{library signature index}}$ corresponds to the first color patch, the next higher $\sqrt{\text{library signature index}}$ to the second color patch and so on (see right hand example below for an alternate method of defining the color correspondence). In the same manner that the selection point color correspondence can be chosen, so can the library signature color correspondence by choosing menu **Edit -> Options -> Color Selector**. The color to name correspondence can be displayed by clicking the **Lib Key** button in the **Classify** window. This generates a palette like image of the correspondences. This can be saved like any other image. Shown here is a typical output. To select (toggle on) a complete library, click on the library window and then choose **Select All** from the **Edit** menu. Up to 1024 signatures can be selected. To deselect (toggle off) all choose **Edit->Clear** menu or double click individual selected spectra. Section [Dynamic Displays and Spectral Plots](#) describes how to change both the color and the associated name after a Lib:File classification has been made.

If Lib:File match lists is checked (default) in **Classify -> Options** then two lists are generated in addition to the classmap. One contains a list of the library selections that are in the classmap and the other list contains the library selections that are not in the classmap. Toggling a list item will display its corresponding spectra. Empty lists are not generated.



Color	idx	Lib Name
1	1	JUN19B.005
2	2	JUN19B.006
3	3	JUN19B.007
4	4	JUN19B.008
5	5	JUN19B.014
6	6	JUN19B.016
7	7	JUN19B.017
8	8	JUN19B.018
9	9	JUN19B.019
10	10	JUN19B.020
11	11	JUN19B.021
12	12	JUN19B.022
13	13	JUN20A.001
14	14	JUN20A.002
15	15	JUN20A.003
16	16	JUN20A.004
17	17	JUN20A.005
18	18	JUN20A.006
19	19	JUN20A.007
20	20	JUN20A.008
21	21	JUN20A.009
22	22	JUN20A.010
23	23	JUN20A.011
24	24	JUN20A.012

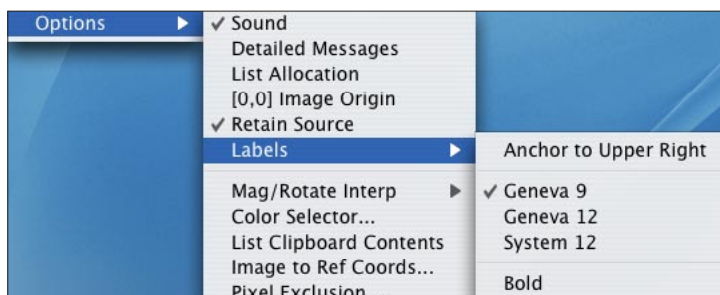
At left is a portion of 412 color to library signature correspondences after Select All was chosen in the Edit menu to toggle-on all of the library entries showing the default colors. At right is the same library color coding after Use source spectra color has been selected in Classify: Options: More Options.



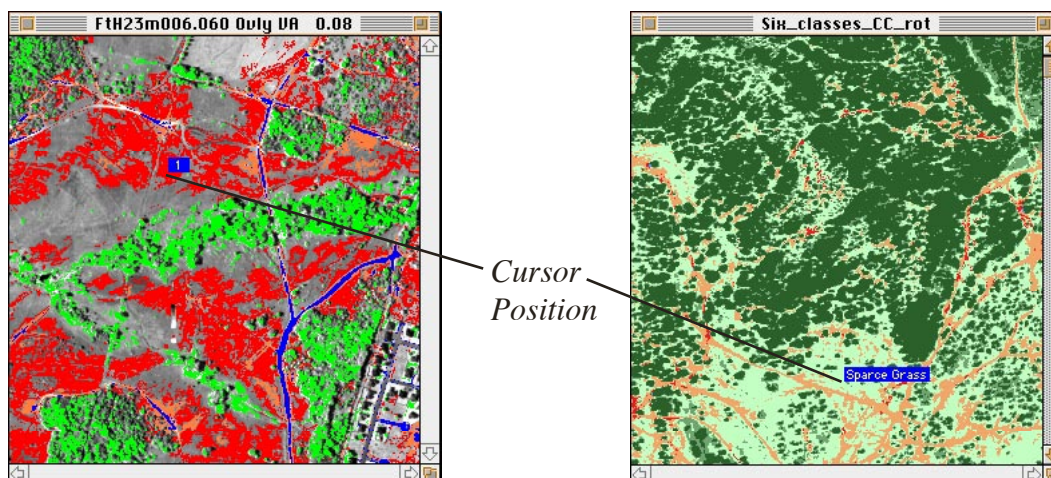
Color	idx	Lib Name
1	1	JUN19B.005
2	2	JUN19B.006
3	3	JUN19B.007
4	4	JUN19B.008
5	5	JUN19B.014
6	6	JUN19B.016
7	7	JUN19B.017
8	8	JUN19B.018
9	9	JUN19B.019
10	10	JUN19B.020
11	11	JUN19B.021
12	12	JUN19B.022
13	13	JUN20A.001
14	14	JUN20A.002
15	15	JUN20A.003
16	16	JUN20A.004
17	17	JUN20A.005
18	18	JUN20A.006
19	19	JUN20A.007
20	20	JUN20A.008
21	21	JUN20A.009
22	22	JUN20A.010
23	23	JUN20A.011
24	24	JUN20A.012

Dynamic Displays and Spectral Plots

Each time a classification output image (class map) is created a list is generated that associates the colored regions with the selection point number or library signature name corresponding to that region. Then, as you move the cursor over the classified image the name appears as a small label next to the cursor. The labelling font and its position can be changed by choosing menu **Edit -> Options -> Labels**. Saving a classification also saves its correspondence in a text file with a “.c2n” extension. The correspondence between the color of a region and the label name that appears is completely editable by using the **Utilities -> [Class Map Editor](#)** described later.




Below are two examples of classification outputs showing the region labels for a specific cursor position.



On the left is the output image from a 4 point selection File:File classification showing point #1's region label. Vector Angle was used with a tolerance of .08 radians. To the right is the classification output of a different image cube resulting from a Lib:File search. The label and color of either image can be changed via menu Utilities->Class Map Editor.

Attach Class Map



An image can be made into a “virtual” class map by linking it with an actual class map using menu **Image -> Attach -> Class Map**. The actual class map image must be front most. This displays the dialog at right listing “virtual” class map image candidates. Once an image is selected then moving the cursor over it displays the label corresponding to the same position in the actual class map. Closing the class map breaks the attachment.



A dynamic display similar to a classification output image can be generated when performing a File:Lib search. Normally, this consists of selecting an image cube point and classifying it against a library of signatures list of best matches. However, if the Enable Dynamic Display check box in the Classify Options dialog is then the match vector





is chosen from the cursor's position on the image cube. A label containing the best matches will be computed and displayed as the user moves the cursor. The dynamic search remains in effect until any mouse click in the Classify window. It must be restarted by clicking the Classify button. If a color composite Link image of the cube image has been created then the dynamic display will also appear when the cursor is moved over the linked image.




The dynamic display of 4 library spectra that most closely match the image cube spectrum at the cursor position. The graph is the corresponding dynamic spectral plot of the top 2 library spectra (red and green) along with the cube plot in black. In this case the Method chosen was Vector Angle.

A dynamic display of the current image cube vector position can be shown by selecting Spectra (Dynamic) from the Plot submenu. As the user moves the cursor over the image cube or linked color composite the plot is automatically updated. The data is read from the image disk file and works best when the cube file is in BIL or BSQ format. If the user has enabled Dynamic Display in the Classify Options dialog, described previously, then the two closest library spectra will also be plotted on the same graph. Similar to the static Selected Points plots, you can 'Scale' and 'Hide Bad'.

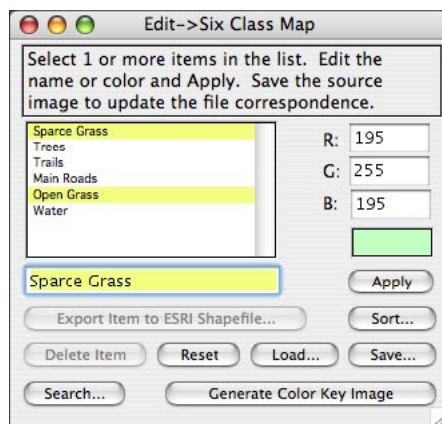



Dynamic spectral plots of the same cube vector point with and without the optional check boxes selected.



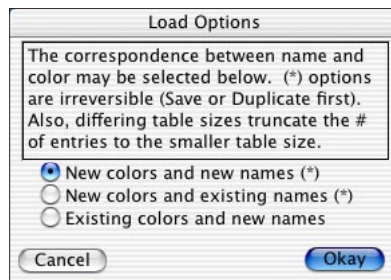
Class Map Editor

The correspondence between the color and name of an image classification can be modified by using **Utilities -> Class Map Editor**. You can do this after the image is generated or retrieved from a saved file (the saved file has an associated *.c2n file defining the correspondence). A classification image is a true color image where each pixel is either a gray value (including black if it's a mask) or a specific color. Associated with each color is an ASCII string comprising the name that will be displayed whenever the cursor is over that color or no display when over a gray (image) pixel. The Editor allows changing that association so that a color or group of colors may be “renamed” or the color itself changed with a corresponding change in the image. For example: a classification region titled “Water” and represented by a light cyan color can be changed to “Shallow Pond” with a medium blue color. A single colored region can be changed to any other color (not gray) as long that color is not already used for a different region. You can choose multiple items using the shift key for contiguous and the command key for non-contiguous selections (Windows version uses the Ctrl key for non-contiguous). As each item is chosen the corresponding class on the image will flicker in white. So, as multiple selections are made they produce a larger and larger composite region. The **Apply** button will consolidate all of the selected items into the displayed name and color and change the image appropriately. The original class map and its associated color correspondence can be restored at any time by the **Reset** button.



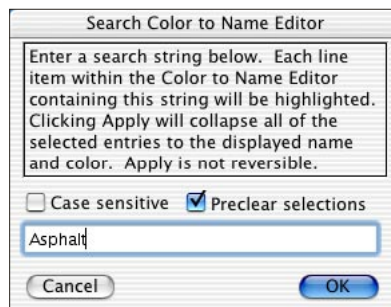
The Editor window on the left shows 2 highlighted items that have been selected and are about to be consolidated into a single class with the name and color shown. The classmap image on the right shows where these 4 items are located by flickering them in white. When Apply is clicked both items will have the color: {195,255,195} and the item list will be updated. The Search button displays a dialog allowing you to enter a text string and every list item containing that string will be highlighted as if manually selected. The Export... button saves a single item in the ESRI Shape File format (see section [ESRI Shape Files](#)).

- Scrollable List** For each classification region the associated name and color. Click to select, edit name and color followed by **Apply** or return. Multiple lines may be selected via shift click and command click.
- R:G:B:** The color components as each is selected. Edit as desired. The color patch shows the composite.
- Apply** Change the selected list name to that of the edit field and color and update the image. Hitting the return key is same as Apply.
- Delete** Remove the selected list item and replace its color with black.
- Color Key** Generate a legend type image (see section [Library Searches](#)).
- Load/Save** The color to name correspondence may be saved without having to resave the image. It must be saved as “name”.c2n. A previous (e.g. master) may be loaded to overwrite an original correspondence. This is an ASCII file and is also compatible with Edit->Options->Color Selector. See section [Selection Point Colors](#).



Various combinations of the current and new colors and names may be selected.

- Search** This displays the dialog shown below and is used to select items based upon their common names. Each item containing the search string will be highlighted. No modifications are made until **Apply** is clicked.



If Preclear is checked then any existing highlighting is cleared else items matching the search string are added to currently highlighted items.

Sort



All of the elements in the list, i.e., the class names will be re-ordered according to this dialog option. This makes it easier to group a large class map by color such as classifying against a library.

Export Item... Convert a single item into a shape file as specified in the ESRI Shapefile Technical Description - July 1998. Brings up the user dialog shown below presenting a number of options.





Save files: Solicits a file name using the item as the base name and produces the three files listed. The .shp extension is a must, the other two are automatic.

Plot shape: Generates a vector plot of the result at 2 times size. This plot will exactly match a 2 times magnified image of the class map. See sections [Magnify](#) and [Flicker/Superimpose](#).

Flip output: Invert the y coordinate so that the origin is at the upper left instead of the image's lower left.

Transform: Convert shape file coordinates to the transformation specified by Image to Reference Coordinates.



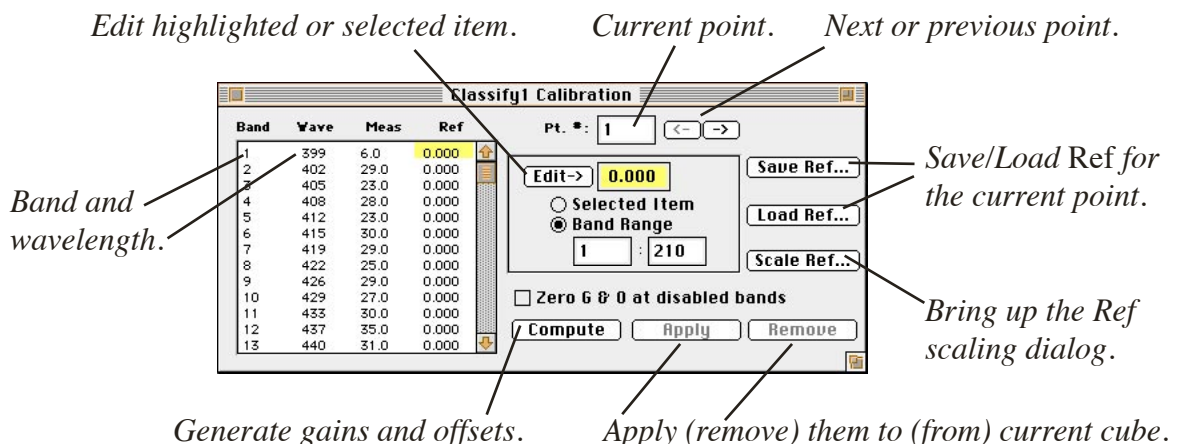
The left image is a 2x magnified view of a six class map. The main dark green regions are the selected item in the Class Map Editor that was Exported to a shape file. The right image is the vector plot output showing the class outline in green and any interior voids in the complement color, magenta.

Spectral Calibration

HyperCube has the ability to empirically calibrate a cube against a set of two or more reference spectral vectors. The reference may be from an external file, another cube or a library. Calibration uses the same point selection and neighborhood toggling as is done for a Classification. It can be invoked whenever a Classification window for a given cube is active by choosing **Functions -> Calibration**. Calibration produces a gains and offsets file that may be saved or immediately applied so that you can see the effect on a classification or plot by toggling it on and off.


The idea behind calibration is simple, try to compute a separate linear transformation, gain (or scale) and offset, for each band that will convert that band's pixel values to the same values in a reference set. The assumption is that the difference between the measured and reference spectral values is closely represented by a linear change. In practice this isn't exact so more than two spectral vectors (selection points) are used in a least square solution.

A typical calibration scenario is as follows: Suppose there are four panels in the source cube that have known spectral reflectance signatures (somewhere else). Then choose four selection points over these panels and toggle them on as if doing a classification. You may then choose additional neighborhood points around these if an average spectral signature for each panel is desired. Bring up the Calibration dialog as stated above. The **Meas** column will be the measured spectral vector for the given point and the **Ref** column will be the desired reference spectra and initially will be zero. The next step is to put data into the **Ref** column that you want the **Meas** value to be transformed into. This could be done by manually entering the data as described below (very tedious for hyperspectral data!). However, the easiest method is to drag an existing reference spectral signature plot over top of the calibration window list. It will be automatically interpolated from its wavelengths to the cube's wavelengths (see section [Plotting Signatures](#)). The signature plot can come from a library list (see [Spectral Libraries](#)) or from another cube. The later is extracted by loading a second (reference) cube, selecting a point, opening a classification window, toggling that point on and then clicking the **Plot** ✓ button. The resulting signature plot can then be dragged onto the calibration window. The final method of entering **Ref** data is to read it from a previously generated file. Its format is ASCII tab delimited text pairs of wavelength and value. Note that it is not defined by a band number. The wavelengths are used to properly interpolate the values into the cube's wavelengths. If the cube's spectral range extends beyond the **Ref** values then the corresponding **Ref** data is zero filled.

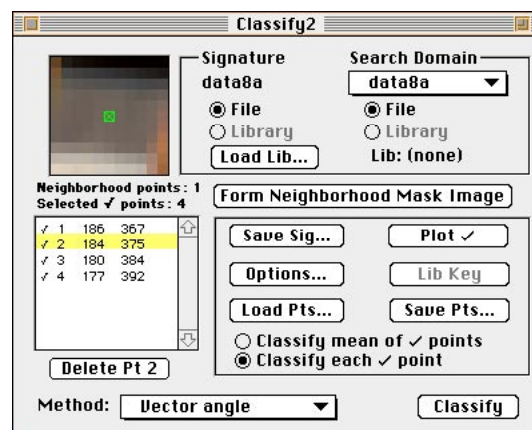


The various **Calibration Dialog** controls are defined below:

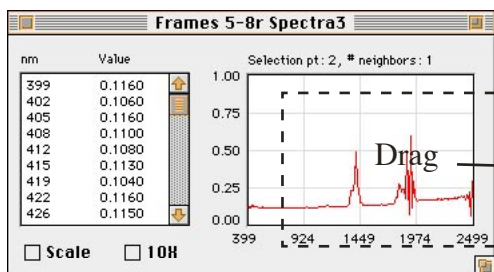
Scrollable List:	Contains four columns listing the Band , Wavelength , Meas value and Reference value for the given Pt. # . Clicking on a Meas or Ref column item allows that item to be edited, see Edit button.
Pt. #:	The current selection point number shown in the list. May be manually edited or incrementally changed with the adjacent arrow buttons.
<- ->:	Step to the previous or next selection point.
Edit->:	Enabled when a Meas or Ref item is highlighted (clicked on). The highlighted item will be changed to the number in the edit field. Works in conjunction with the two radio buttons described next.
Selected Item:	If on then only the selected (high lighted) Meas or Ref item will be changed. Mutually exclusive with Band Range .
Band Range:	If on then all of the items within the Band Range specified will be changed to the same given value. Mutually exclusive with Selected Item .
Save Ref...	Displays a standard Mac save dialog and saves the current point Ref data as ASCII tab delimited wavelength, numeric value pairs.
Load Ref...	The opposite of Save , loads a previously defined Ref file and interpolates as necessary to overwrite the current point Ref data.
Scale Ref...	Brings up a dialog that permits you to enter a numeric value that will immediately be multiplied times a specific Ref column or all Ref columns. A convenience to change the range of the Ref data.
Zero G & O...	When checked the gain and offset computed for each <u>disabled</u> band is set identically to zero. When these gains and offsets are then applied to a cube the disabled band's spectral values will then be zero. See section Band List for a description of disabled bands.
Compute	This button becomes active when at least two Classify points are selected. When clicked the gains and offsets are computed and shown in a separate data list window that can be saved via Save in the File menu. Also generated is a data list showing the corresponding residuals from the least squares solution.
Apply	Becomes active after a Compute and immediately applies the gains and offsets to the current cube meaning that any plots are automatically updated and subsequent classifies will accommodate them.
Remove	Remove any gains and offsets that may be in effect for the given cube.



Four selection points from cube Frames 5-8r have been chosen as reference values. Point #2's spectra is shown below and is being dragged over the corresponding point #2 listing of data8a's calibration window



These are the four points of data8a that have been selected as the matching spectral vectors to the four reference vectors of Frames 5-8r along with the data8a calibration window below.



Spectral plot of point #2 of Frames 5-8r.

Band	Wave	Meas	Ref
1	434	0.126	0.000
2	440	0.127	0.000
3	445	0.130	0.000
4	451	0.125	0.000
5	457	0.129	0.000
6	462	0.129	0.000
7	468	0.130	0.000
8	473	0.130	0.000
9	479	0.131	0.000
10	484	0.130	0.000
11	490	0.129	0.000
12	495	0.130	0.000
13	501	0.132	0.000

Point #2 of data8a's calibration window before the drag operation.

Band	Wave	Meas	Ref
1	434	0.126	0.123
2	440	0.127	0.121
3	445	0.130	0.121
4	451	0.125	0.118
5	457	0.129	0.117
6	462	0.129	0.116
7	468	0.130	0.114
8	473	0.130	0.114
9	479	0.131	0.115
10	484	0.130	0.114
11	490	0.129	0.112
12	495	0.130	0.112
13	501	0.132	0.111

Point #2 after the drag operation. The spectral plot values have been interpolated and placed into the Ref column. This is repeated for each of the remaining points followed by the Compute button which then produces the Gains and Offsets values along with the Residuals for each band. These may be applied immediately (Apply) to data8a or saved (File->Save).

Band	Residuals
1	+0.0198
2	+0.0194
3	+0.0206
4	+0.0198
5	+0.0180
6	+0.0188
7	+0.0205
8	+0.0187
9	+0.0184

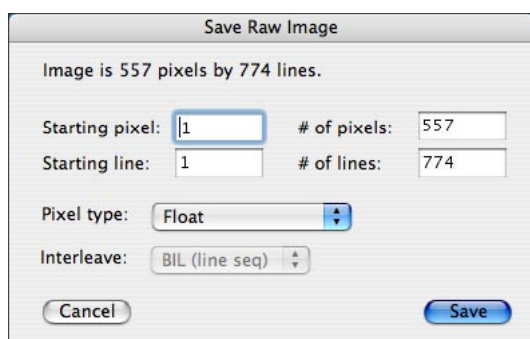
Band	Gain	Offset
1	+1.3375	+0.1014
2	+1.3738	+0.0976
3	+1.3700	+0.0974
4	+1.3684	+0.0945
5	+1.3883	+0.0887
6	+1.3708	+0.0893
7	+1.3728	+0.0881
8	+1.3567	+0.0886
9	+1.3549	+0.0867

Saving Images

HyperCube can save images in various formats depending upon the File menu item chosen and on the internal format of the image. If the image has not as yet been saved a standard save file dialog with the suggested name will appear asking for confirmation if the file already exists to overwrite it. Subsequent saves of the same image will perform the save directly automatically replacing the old image with the same name. However, if the image has been modified since the last save the confirmation dialog will again be presented. To save a specific band, scroll the cube to the band, choose **Duplicate** from the **Edit** menu and then **Save**. To save all of the bands within a cube image as sequentially numbered TIFF files, choose menu **File -> Save As...**, **Bands**. **File -> Open As...**, **Bands** will reload the bands into a cube (see section [Constructing/Modifying a Cube](#)).

File	Edit	Applications
New...		⌘N
Open...		⌘O
Open As...		
Load Selection...		⌘L
Close		⌘W
Close All		
Save		⌘S
Save As...		
Page Setup...		
Print...		⌘P

File->Save As... is used to save images in **Bands**, ERDAS Imagine®, HDF, JPEG, LASF, Multiband, NITF, Raw, SRF, Text or TIFF format. If ERDAS Imagine®, HDF, NITF, SRF or TIFF is selected from the **Format** popup then no other dialog is required. **Raw** is used to save a single image and allows specifying a rectangular subset and the pixel type. **Save As... LASF** will display an alert if the image does not have a geo reference but will continue. Selecting [Raw](#) displays:

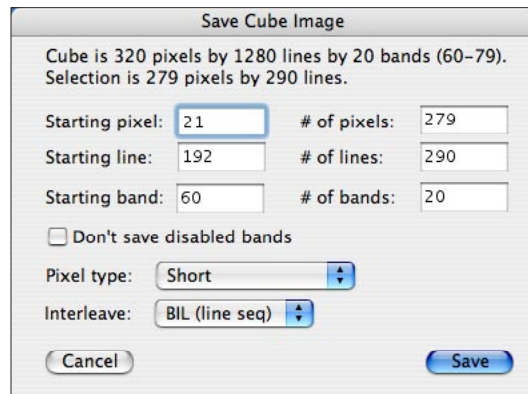


The Save As... Raw dialog defining the file format for a 32 bit float pixel type. If the image had a selection rectangle displayed upon it the coordinates of the selection area would have been automatically inserted into the various edit fields as defaults.

The dialog edit fields are defined as:

- Starting pixel:** First pixel defining the image rectangle to save. These first 4 edit fields are preset to the current dimensions of the image.
- Starting line:** First line to save.
- # of pixels:** How many pixels per line to save.
- # of lines:** Corresponding number of lines to save.
- Pixel type:** Either byte, short, float or long.
- Interleave:** If the image is true color then this popup will be enabled to allow the image to be saved as a 3 band cube with the specified interleave.

The **Save As...** function is also used to save an image cube as a multiband image. The dialog is similar to the previous **Raw** dialog with the addition of two more edit fields and a check box to selectively dump only bands whose status is **Enabled** in the [Bands List](#). Unlike the **Raw** dialog, the **Interleave** popup is always enabled, not just for true color. **Save As...** also generates a new [*.hdr](#), and for multiband a new [*.wvl](#), file.



This dialog will save all 20 bands of a multiband image within the area of a selection rectangle whose coordinates are shown in the edit field. The interleave format is BIL (Band Interleave by Line) and the pixel depth is 16 bits.

The additional dialog items are defined as:

Starting band: The first band of the cube to start saving. The default is the first band.

of bands: How many bands to dump beginning at the **Starting band**. The default is the number of bands in the image cube.

Don't save dis... If this check box is set then disabled bands specified in the **Bands List** will be skipped. This option is unavailable if all bands are enabled.

Note: To change the format of the multiband [file](#) use menu **Utilities -> Reformat Cube File**. It is much more general and operates directly on the file pixels images that have been loaded into memory.

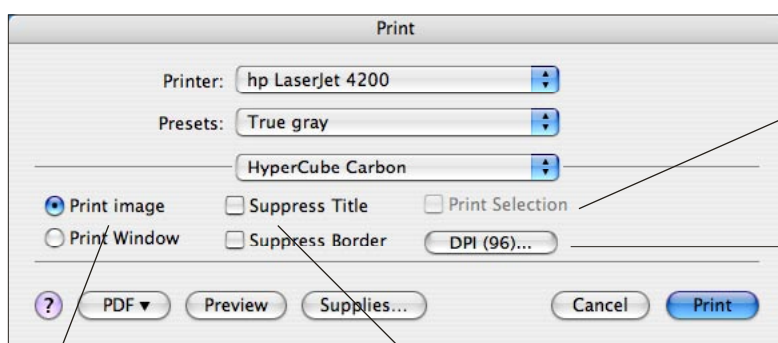
Save As..., **Bands** can be used when an image cube window is active. Its purpose is to save each band of the cube as a **TIFF** file forming a set of sequenced images that can be made into a QuickTime™ movie using a program such as Adobe Premiere™. The user chosen name will be concatenated with the a number (e.g. **name.001**) as each band is extracted and automatically saved. These band images can be automatically reloaded and formed into a cube using menu **File -> Open As..., Bands** by selecting any one of the numbered images.

Saving a [Class Map](#) image also saves a text file with the same base name and a **.c2n** file extension defining the correspondence between the image colors and their class names.

Printing

Any HyperCube window or underlying image can be printed to the currently chosen printer using the functions **File -> Page Setup...** and **File -> Print**. These bring up the standard Mac or Windows dialogs with options usually specific to the type of printer. The default is to print the contents of the front most window with its name below centered on the page. **Print Image** prints the entire image regardless of what is shown within the window whereas **Print Window** limits itself to just the way the window appears including any operations being performed on the window such as gray mapping or filtering. **Print Image** will be replaced by **Print Text** if the window contains any ASCII lists. Whether window or image, the output will always be resized to fit within the print area.

Macintosh Print Dialog details

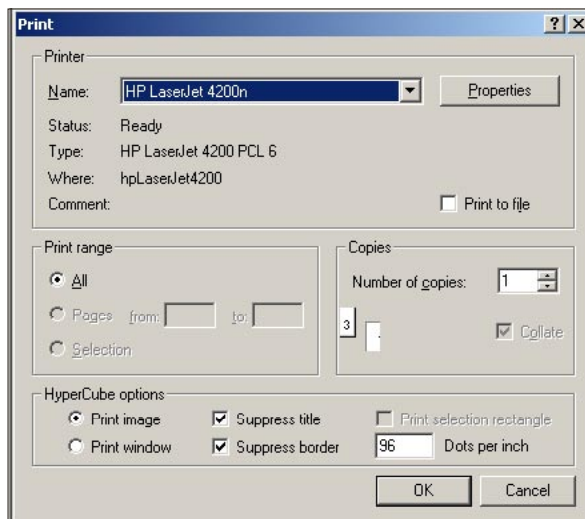


If a selection rectangle has been drawn then you have the option of just printing its contents.

The desired output resolution. However, the image will always be scaled, if necessary, to fit the page.

Selects either print the entire image or just that portion shown in the current window.

You have the option of omitting a title or border.

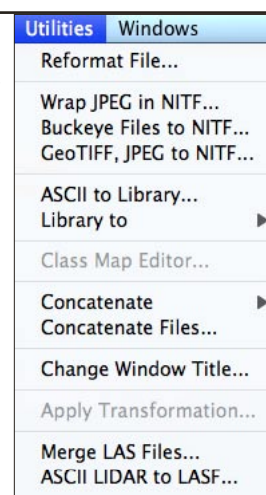


The corresponding Print dialog in the Windows version. The HyperCube options appear at the bottom.

Utilities

Reformatting a File

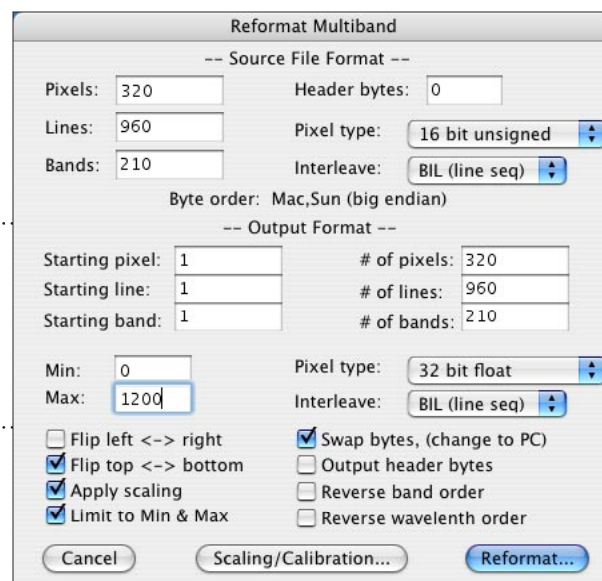
Menu Utilities -> Reformat File... operates on Multiband, Raw, NITF sequential multiband and NITF mono image files. If the file to be reformatted has a *.hdr file these values will be extracted and inserted as defaults into the appropriate edit fields of the **Reformat** dialog window. A new *.[hdr](#) file will be generated along with an updated *.wvl file if one existed for the source image cube. NITF multiband and mono files don't have a separate hdr or wvl file, all sizing information is contained in the source file. Reformatting a NITF file produces a cube or mono file, not another NITF file. Selecting a single output band from a multiband image essentially extracts the band and writes a single raw type image. The **Reformat** dialog appears below.



*The upper elements define the spatial format of the data. Normally, this is retrieved from the *.hdr file or directly from a NITF file.*

This dialog area specifies the format of the converted output file. Its initial values are set equal to the source file size and are edited as desired.

Check boxes used to change the order of pixels within a line, the order of the lines within a band and whether any scaling, limiting or byte swapping is to be applied.



The example shown will generate a full size cube scaling 16 bit signed pixels to floating point pixels while limiting the range to [0,1200] and reversing the order of the lines. Swap bytes will change the byte order to PC (little endian).

- # Pixels: The number of pixels per line in the source file.
- # Lines: The number of lines in the source file.
- # Bands: The number of bands in the source file.
- Header bytes: The number of bytes in the source file preceding the image data.
- Pixel type: Size and type of the source pixels.
- Interleave: Structure of the source file: band, line or pixel.

# Pixels, Lines, Bands:	The subset of the source file to be reformatted.
Start pixel, line, band:	Starting coordinates within the source file for subsetting.
Output Pixel type:	Output pixel size. If Limit... is checked then output pixels are kept within the Min & Max range, else they are automatically limited to the range of the pixel type.
Output Interleave:	The format structure of the output cube.
Min, Max:	The lower and upper ranges for limiting the output pixels.
Flip left <-> right:	Change the order of pixels within a line.
Flip top <-> bottom:	Reverse the order of the lines, i.e. first line last.
Apply scaling:	Modify each output band according to the type of scaling selected in Scaling/Calibration. If checked then each band is modified by: $\text{scale} * \text{band} + \text{offset}$.
Swap bytes...	This check box allows swapping the order of bytes within 16 bit and 32 bit pixels. Note that this applies to the output. The source byte order is taken from the header and if it is unknown you will be prompted for a best guess.
Output header bytes:	Directly copy the source file header bytes to the output file.
Reverse band order:	Output the bands in reverse order.
Reverse wavelength...	Also reverse the order of the output wavelength file.
Scaling/Calibration:	This brings up the following user dialog which is used to determine the type of scaling to be applied to the output. The Apply scaling option checkbox becomes active after returning from this dialog.

Apply this scale and offset to every band.

-- or --

Select a file containing a separate scale & offset pair for each band. The number of pairs must equal the number of bands in the source file.



A typical gains and offsets file content: a pair of values for each band...

```
0.00716025 -0.0044436
0.0054349 -1.367508
0.00553983 -1.2478907
0.00583894 -1.2276135
0.00571494 -1.1864042
0.00562544 -1.1921199
0.00553107 -1.5069996
```

...

...

Scale and offset pairs are separated by a space or tab.

Wrap JPEG in NITF

A JPEG (DCT) or JPEG (2000) image file can be embedded within a NITF file using this utility. Additionally, an IKONOS formatted [RPC](#) text file may be included. This is the complementary operation to the JPEG extraction described in section [NITF Files](#). Choosing menu **Utilities -> Wrap JPEG in NITF** displays the normal open file dialog. Select the JPEG file (either the older DCT type or a new 2000). A second open file dialog will appear allowing you to include a corresponding [RPC](#) (Rational Polynomial Coefficients) file.

Shown below is a portion of a NITF image header resulting from embedding a JPEG 2000 file. By setting the menu [Edit -> Options -> Detailed Messages](#) flag, subsequent image loads will display header information.

```
...
ENCRYP: 0
ISORCE:
NROWS: 00029696
NCOLS: 00035840
PVTYP: INT
IREP: MONO
ICAT: VIS
ABPP: 11
PJUST: R
ICORDS: G
IGE0L0: 013202S1445153E013208S1450248E014247S1450259E014242S1445143E
NICOM: 1
ICOM1: Saved via HyperCube, www.tec.army.mil

IC: C8
COMRAT: V039

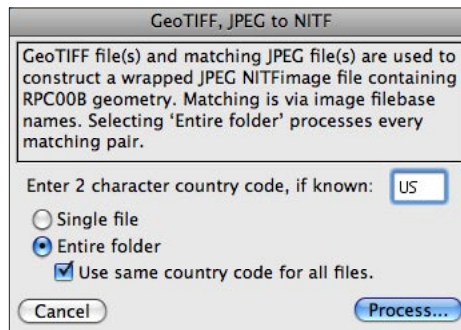
NBANDS: 1
  IREPBAND1: M
  ISUBCAT1:
  IFC1: N
  IMFLT1:
  NLUTS1: 0

ISYNC: 0
IMODE: B
NBPR: 0035
NBPC: 0029
NPPBH: 1024
NPPBV: 1024
NBPP: 11
IDLVL: 001
IALVL: 000
ILOC: 0000000000
IMAG: 1.0
UDIDL: 00000
IXSHDL: 01055
...
```

Part of the NITF image header after a large JPEG 2000 file has been embedded and re-opened. The value of C8 in the IC (Image Compression) field indicates that the image data is actually a complete JPEG 2000 file.

GeoTIFF, JPEG to NITF

The geometry of an existing [GeoTIFF](#) image in UTM coordinates along with a matching JPEG image can be converted to an encapsulated JPEG [NITF](#) file in geographic coordinates with geographic [RPC](#)'s. Either a single file or an entire folder of files can be converted. Selecting menu **Utilities -> GeoTIFF, JPEG to NITF** brings up the following dialog.



As shown in this example, country code 'US' will be embedded in the final NITF and will be used for each converted file within a folder. Selecting **Process** will then display a **Select Folder** dialog . If option **Single file** had been selected then the standard **Open File** dialog will appear and only that file will be converted. The program proceeds on its own and when finished displays a summary text window. It is important that for each [GeoTIFF](#) file (name.tif) there is a corresponding JPEG file (name.ext) where extension 'ext' can be 'jpg' (JPEG DCT) or any of: 'j2c, jp2, j2k, jpc, jpf or jpx' (all JPEG 2000).

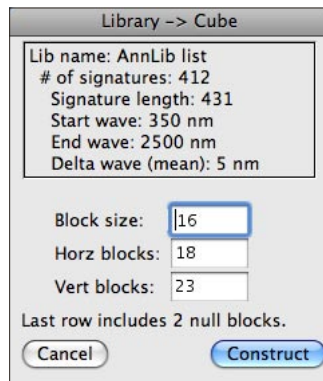
Buckeye Files to NITF

This function is only applicable to special Buckeye imagery and supporting meta data.

Converting a Library to a Cube File

HyperCube can create an image cube whose pixels are the spectral signatures from a library. This is accomplished via **Utilities -> Library to Image Cube File...** which is enabled whenever a **Library List** window is front most. Each signature is replicated into a square block (for clarity). Blocks are placed side by side and row by row in the same order that the signature names appear in the **Library List**. The total number of blocks and their rectangular arrangement (per row and the number of rows) is initially chosen by the program and can be edited by you. Null signatures are added to fill out the last row. If you enter values that produce a block total that is less than the number of signatures in the library then only that number of signatures will be converted.

After selecting **Library to Image Cube File...** the following dialog appears:



The structure of the library is determined by reading the first spectral signature in the Library List (see section [Spectral Libraries](#)).

The 431 band image cube will consist of 288 floating point pixels/line by 368 lines with the last 2 blocks containing zero filled signatures.

The upper portion shows the structure of the library while the lower portion defines the size of the resulting image cube. Note in this case that the total number of blocks (18×23) is 2 greater than the number of signatures, 412, in order to make the cube somewhat square and that the resultant image cube file size will be almost 183 megabytes, ($16 \times 16 \times 18 \times 23 \times 431 \times \text{sizeof(float)}$). These values can be edited and a **Block size** of 8 (~46 MB) still yields a usable image cube. The image cube is stored in band interleave by pixel (bip) format with floating point values along with newly created header and wavelength files. The **Reformat Cube File** utility can be used to reconstruct a different format image cube.

If the **Library List** window already has attached spectral data, as is the case when opened from a **Classify** window, then you are immediately presented with the save file dialog requesting a name for the image cube. If the library list does not have attached spectral data, i.e. opened using **Open As...**, **Library**, then there will be a delay of several seconds while the spectral signatures are being loaded before the save file dialog appears.

ASCII File to Library

A text file can be read and converted to a HyperCube [Library](#) file using menu **Utilities -> ASCII to Library**. This is basically the inverse of menu **Utilities -> Library to ASCII File** using the “Columns” option. Therefore, the text file must consist of white space delimited columns with the first column containing the wavelength values and subsequent columns representing the individual spectral values that will be paired with the wavelengths. A portion of a typical source ASCII file is shown below.


```
350 0.059 0.056 0.055 0.055 0.056 0.058 0.060 ...
355 0.060 0.057 0.056 0.056 0.057 0.058 0.061
360 0.060 0.058 0.057 0.057 0.057 0.059 0.062
365 0.061 0.059 0.058 0.058 0.059 0.060 0.063
370 0.062 0.060 0.059 0.059 0.060 0.062 0.065
375 0.064 0.061 0.061 0.061 0.061 0.063 0.067
380 0.066 0.063 0.063 0.063 0.063 0.066 0.069
385 0.068 0.065 0.065 0.065 0.066 0.068 0.071
390 0.070 0.067 0.067 0.067 0.068 0.070 0.074
395 0.072 0.069 0.069 0.069 0.070 0.072 0.076
400 0.073 0.071 0.071 0.071 0.072 0.074 0.078
405 0.075 0.073 0.072 0.072 0.074 0.076 0.080
410 0.076 0.074 0.074 0.074 0.075 0.078 0.082
415 0.078 0.076 0.076 0.076 0.077 0.080 0.084
420 0.080 0.078 0.078 0.078 0.079 0.082 0.086
425 0.082 0.081 0.080 0.081 0.082 0.085 0.089
430 0.085 0.083 0.083 0.084 0.085 0.088 0.092

.
2500 0.28 0.19 0.28 0.26 0.26 0.24 0.27...
```

*A text file subset in which the first column defines the wavelengths and subsequent columns contain the spectral values for a series of signatures. The result of menu **Utilities -> ASCII to Library** will be a file containing the signature file names and a separate file for each signature (see [Spectral Libraries](#)).*

Converting a Library to an ASCII File

The currently active (front most) Library List window can be saved as a text file by choosing Utilities -> Library to ASCII File. You are presented with the following dialog describing the size of the library along with several options defining the format to be used when generating the ASCII file.



The structure of the library is determined by reading the first spectral signature in the Library List (see section [Spectral Libraries](#)).

If Rows is selected then each row will be a complete signature (431 points in this case). If Columns is selected then a row will consist of all of the spectral values for each signature for a given wavelength (412 values per row shown here).

Include Wavelengths adds an extra row or column at the beginning containing the wavelengths.

If the Library List window is already loaded such as part of a Classify window then the standard save file dialog is presented else there will be a delay while the library data is extracted followed by the save dialog requesting the name of the ASCII file.

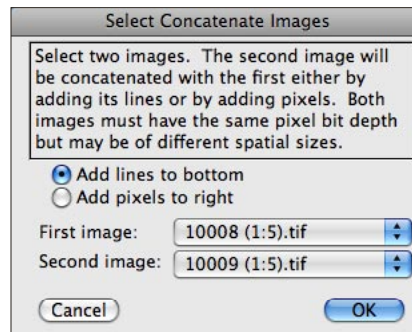
LibRow.text						
	A	B	C	D	E	F
1	350	355	360	365	370	375
2	0.0224	0.0228	0.0205	0.0208	0.0193	0.0197
3	0.025	0.024	0.022	0.0235	0.0213	0.0208
4	0.026	0.0228	0.024	0.021	0.0216	0.0217
5	0.0239	0.0247	0.0209	0.0238	0.0211	0.0217
6	0.0243	0.0222	0.023	0.022	0.0207	0.0223
7	0.0302	0.0272	0.0263	0.0276	0.0259	0.0285
8	0.0201	0.0168	0.0172	0.016	0.0158	0.0173
9	0.0202	0.0176	0.017	0.0150	0.016	0.0166
10	0.0178	0.0192	0.016	0.0197	0.0158	0.0179
11	0.0308	0.03	0.0262	0.0294	0.0253	0.0266
12	0.0315	0.0291	0.0279	0.0263	0.0266	0.0255
13	0.033	0.0294	0.0283	0.0279	0.028	0.0256
14	0.0207	0.0171	0.0177	0.0167	0.0158	0.0176
15	0.02	0.0206	0.0196	0.0192	0.0183	0.0189
16	0.0301	0.0263	0.028	0.0251	0.0256	0.0244

LibColumn.text						
	A	B	C	D	E	F
1	350	0.0224	0.025	0.026	0.0239	0.0243
2	355	0.0228	0.024	0.0228	0.0247	0.0222
3	360	0.0205	0.022	0.024	0.0209	0.023
4	365	0.0208	0.0235	0.021	0.0238	0.022
5	370	0.0193	0.0213	0.0216	0.0211	0.0207
6	375	0.0197	0.0208	0.0217	0.0217	0.0223
7	380	0.0188	0.0211	0.0203	0.0221	0.0197
8	385	0.0173	0.0198	0.0209	0.0212	0.021
9	390	0.0160	0.021	0.0192	0.0232	0.0203
10	395	0.0171	0.0194	0.02	0.0221	0.02
11	400	0.018	0.0206	0.0204	0.0238	0.0208
12	405	0.0186	0.0217	0.0222	0.0252	0.0214
13	410	0.0202	0.0234	0.0236	0.0273	0.0235
14	415	0.022	0.0249	0.025	0.0295	0.0244
15	420	0.0229	0.0267	0.0271	0.0311	0.026
16	425	0.0251	0.0298	0.0288	0.0341	0.0284

Small portions of a library saved in row and column form respectively and opened in Excel™. In each case the wavelength option was included.

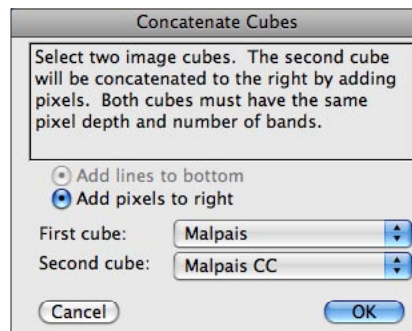
Concatenate Images

Two images may be combined side to side or top to bottom to form a new single image by choosing **Utilities -> Concatenate -> Images, Bands**. The images don't have to be the same spatial size, the output will be sized to the minimum dimension of the two. However, they must have the same pixel depth, i.e., both gray or true color. If the images both have [attached data](#) then the attached data pixel size and type (e.g., unsigned short, short, float, long) must also match. The bands of a cube are also candidates for concatenation. To concatenate an entire cube see the [Concatenate Cubes](#) description in the following subheading.



Concatenate Cubes

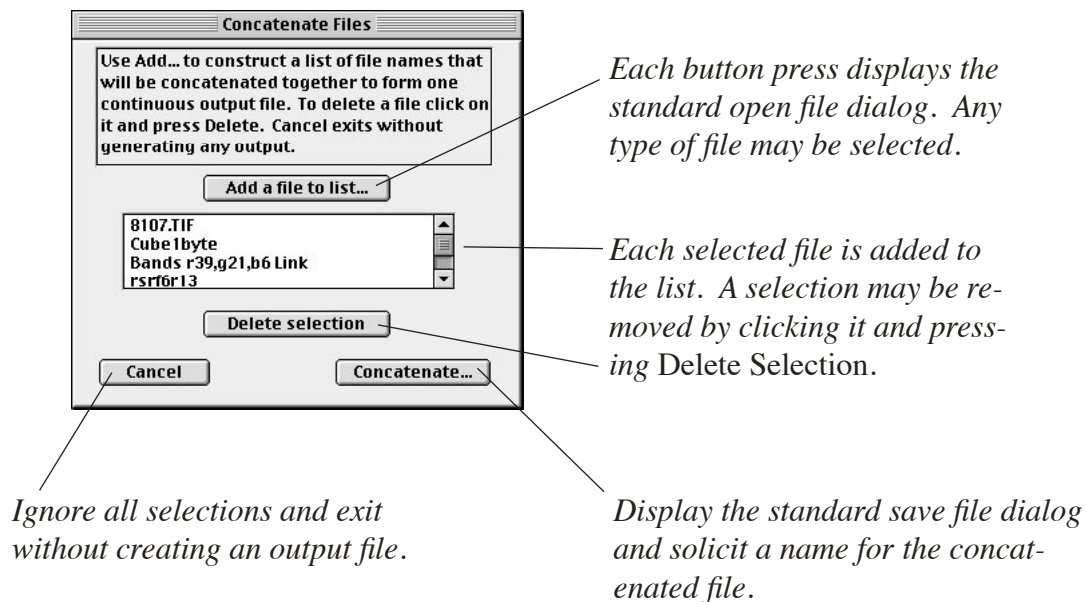
Two in-memory image cubes can be joined side-by-side if they have the same pixel depth and the same number of bands. In order to concatenate them by adding their lines they must first be [saved](#) as cube files and then processed using menu [Utilities -> Concatenate Files](#).



Concatenate Files

This utility allows you to make one continuous file from two or more files by reading each file and sequentially concatenating their contents to create a single new output file, e.g. combining individual Landsat file bands into a contiguous BSQ (Band Sequential) cube. The source files are not altered in any way. Compare this with the previously described **Concatenate Images**, which operates only on images loaded for display.

You start the process by selecting menu **Utilities -> Concatenate Files** which brings up the following dialog:



Changing a Window Title

It is often convenient to be able to change the title (name) of a window without having to first do a **Save as...** . An example is when printing where the title can be changed to something more descriptive and then listed below the output (see section [Printing](#) and note below). To perform the change activate the window and choose **Utilities -> Change Window Title**. This brings up the following dialog.

If checked and window is a cube then the loaded bands will be renumbered starting at 1.



The current window name, Frames 5-8r being changed to a more descriptive title.

Note: Window titles may consist of up to 63 characters, while file names are limited to 31 characters. You will be alerted when saving an image window with too many characters.

Apply Transformation

The active image cube file can be transformed (Utilities -> Apply Transformation) into another cube file by multiplying each spectral vector (signature) by a linear transformation in the form of an M x N matrix. The matrix must be an ASCII file (File->Save As Text and answering affirmative to Matrix) or such as those optionally saved via menu items Image -> Convert Image -> Principal Components, Functions -> Statistics -> Band Correlations and Functions -> F(Signatures). The format consists of a single line containing the dimensions of the matrix, (# rows tab # columns) followed by the M rows of the matrix with each of the N column elements separated by a space or tab and ended with a return. An example of a 4x3 transformation matrix is:

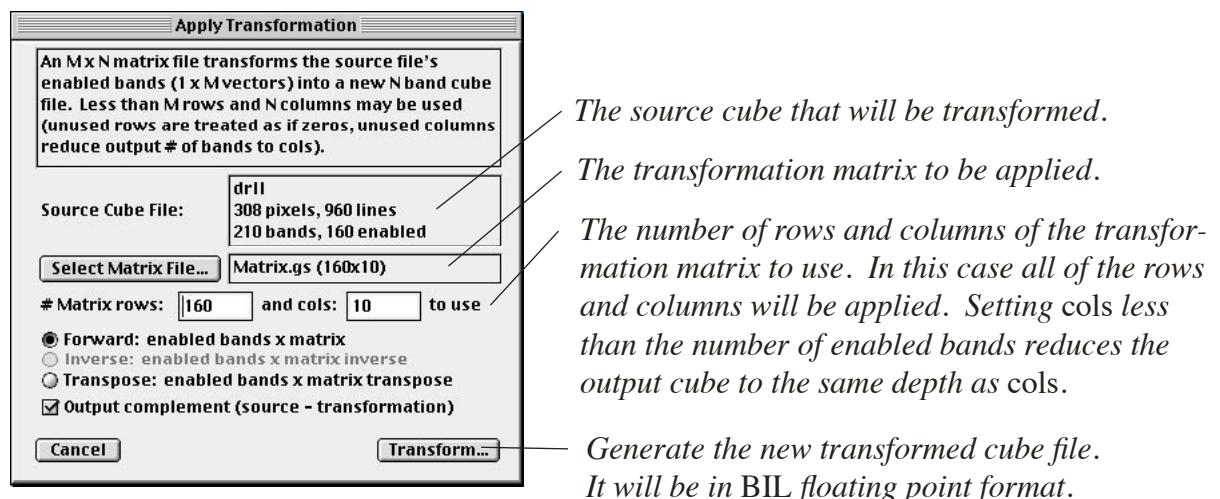
```
4 3
4.6864242694e-01  -6.9847956895e-01  2.5492193504e-01
5.0310142892e-01  4.5755700310e-01  7.0505069625e-01
5.0999322932e-01  -2.8597312694e-01  -3.8238290256e-01
5.1688502971e-01  4.7009281141e-01  -5.4009410160e-01
```

Algebraically, the following occurs:

$AP = B$ for the forward transformation
 $AP^{-1} = B$ for the inverse transformation
 $AP^T = B$ for the transpose transformation, where
A is each 1 by M spectral signature in the source cube
P is the M by N transformation matrix
B is each resulting 1 by N new spectral signature
 P^{-1} is the inverse of matrix P, P^T is the transpose of P

If the Output complement option is checked then $A - B$ is formed.

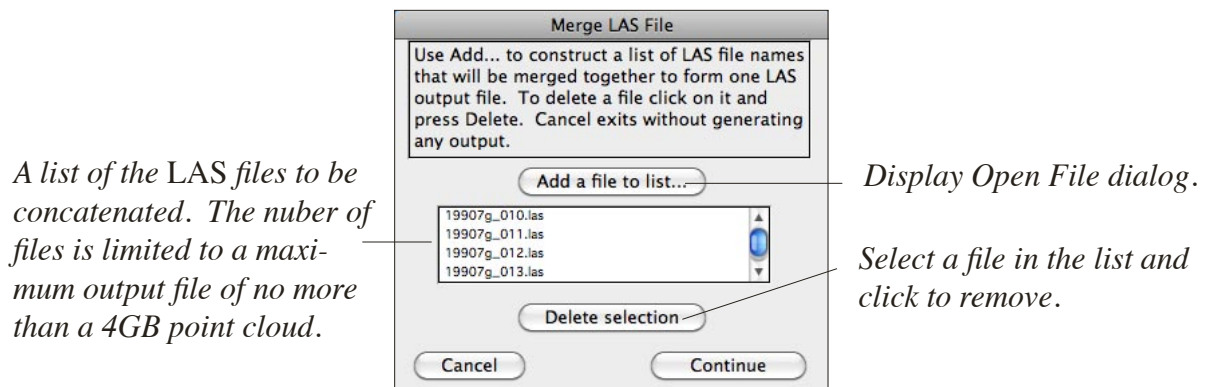
You can use less than M rows or N columns of P, the unused rows are treated as if they are zero and unused columns reduce the output vector dimension to 1 x cols. This allows you to construct an output cube using just the first significant vectors (columns) of a matrix file (see sections [Image to Principal Components](#) and [F\(Signatures\)](#)).



Merge LAS Files

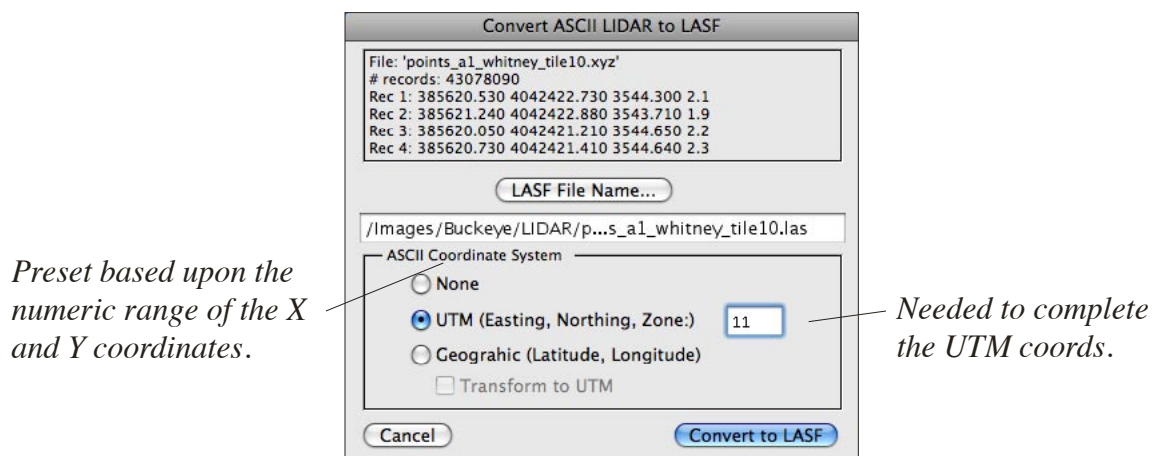
The point cloud data in several LAS files can be merged as a single point cloud in a newly created LAS file. All of the source files must have a common geometry. Specifically, they must be in the same UTM zone with coordinates expressed in meters. Selecting menu **Utilities -> Merge LAS Files** displays a dialog similar to the dialog shown in [Concatenate Files](#). Here, the files are LAS whereas any file type can be selected in [Concatenate Files](#).

The **Continue** button then produces a **Save As** dialog requesting a new file name for the output. Also generated is a list of the source files and an outline depicting how the source file's extents (X,Y max,min) are geometrically positioned to each other. The latter is exactly the same type of diagram image produced in [Mosaic](#) that includes the name of each image color coded within its corresponding outline.



ASCII LIDAR to LASF


A text file (e.g., [ASCII LIDAR](#)) that represents LIDAR data can be directly converted into the LAS file format using menu **Utilities -> ASCII LIDAR to LASF**. Typically, each line of the text file consists of X, Y, Z and intensity return values. The intensity value may be absent. The X and Y values are usually UTM Easting and Northing requiring that the user supply the UTM zone (see below). However, the program also accepts X and Y as representing Latitude and Longitude (note their order) that can be converted to UTM, in which case, the zone is automatically derived.



Applications

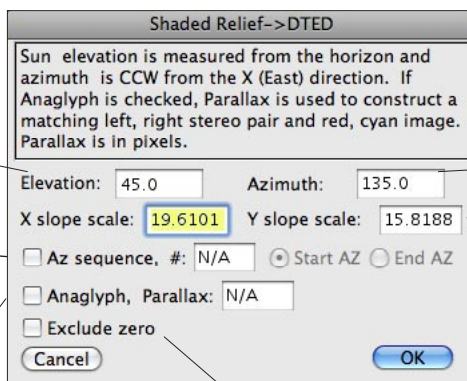
Shaded Relief

An image can be rendered into a synthetic image (or sequence of images) that depicts the effect of uniformly illuminating it with a parallel light source, e.g. the Sun. Each gray value, or its associated (attached) data (see sections [Files, Data...](#), [Raw](#), [Multiband](#), [NITF Files](#), [TIFF Files](#)), is interpreted as an elevation that intersects and reflects the light source as a function of the angle of incidence between the light and the slope of the elevation. Choosing Applications -> Shaded Relief... when a gray image is the front most window brings up the Shaded Relief dialog. Checkbox **Exclude zero** overrides any global pixel exclusion values (see section [Pixel Exclusion](#)).



Elevation angle above the data plane of the light source.

If Az sequence is checked then a synthetic cube of # values between Start Az and End Az is generated. Choosing this option disables Anaglyph.



Heading angle measured CCW from the east or positive X direction.


The denominator used in computing the slope in the X and Y directions. Initial values are computed from the average slope within the image. The lesser these values are the greater the shadowing effect.

If checked then ignore zero values when computing slopes.

If Anaglyph is checked then a red:cyan stereo anaglyphic image is also generated using the specified Parallax value. See section [Stereo Mate](#) for an example.



Gray image representation of data used to generate the Shaded Relief image shown to the right using the values defined in the above dialog.



Resulting Shaded Relief with the light coming from the upper left 45 degrees above the "horizon". The fine vertical striping is part of the data and not an artifact. The source image's data is also copied to this image.

Contour

A non-color image can be displayed as an ordered sequence of planes that intersect the gray values or attached data within the image. The result is a same size color image consisting of connected line segments (contours) at the points of intersection. In addition, the elevation at each line can be shown dynamically by enabling the [Info](#) window and placing the cursor over the contour line. A contour image can be saved as a polyline [ESRI Shapefile](#) using menu File -> Save As... . It can also be saved as a color [TIFF](#) file.

Applications	Image
Shaded Relief...	
Contour...	
Stereo	▶
Perspective...	
Radiance	▶
Lens Distortion...	

Choosing menu Applications -> Contour... displays the dialog shown below.

The contours can represent either the attached data or the displayed gray values.

The Starting and Ending contour respond to any [Pixel Exclusion](#) setting.

Limits the spatial area to any [selection rectangle](#) or the full image. The output can be scaled by a factor of 2 and the contour color is selectable.

Contour DTED

Enter the start, end, placement and output.
Contours may represent window gray values or any attached data. Use menu Functions -> Overlay to combine the output with the source image.

Contour using

☒ Attached data min: 970 max: 1516 ☐ Window gray values min: 13 max: 255

☒ Exclude zero

Limits

Starting contour: 970

Ending contour: 1516

Placement

☒ Interval (spacing): 60.6667 ☐ Number of contours: N/A

Output

☐ Use selection rectangle

☒ 1 x ☐ 2 x Color (rgb): 0 255 0

☒ Form overlay image

Cancel OK


The minimum and maximum values are shown and change as a function of the Use selection rectangle checkbox.

The starting and ending limits of the contours.


You can choose either the spacing between contours or the total number of contours to be drawn.

Presently, set to green.

Superimpose contours on a copy of source image else only output contours.



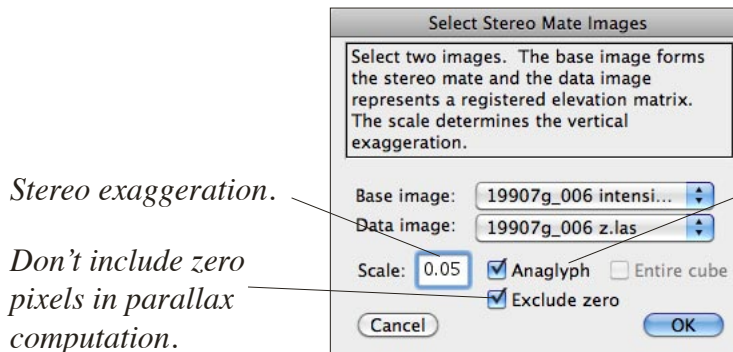
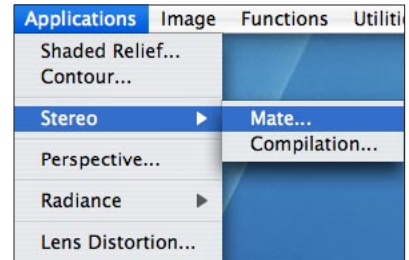
The source image representing elevation.



The resulting contours via the dialog shown.

Stereo Mate

A Base image along with a Data image representing elevation that is spatially registered to the Base image can be used to form a new version of the Base image with elevation induced displacements by menu Applications -> Stereo -> Mate. If Anaglyph is checked then the resulting color image can be viewed in stereo with red and cyan glasses. If the Base image is color the Stereo Mate image will also be color and if placed along side of each other as narrow strips you can view them with a pocket stereoscope.




Stereo exaggeration.

Don't include zero pixels in parallax computation.

Generate a red-cyan anaglyphic image.


If the Base image is a cube band then generate a stereo mate for each band.



An anaglyphic stereo image generated with red being the stereo mate and green and blue being the Base image. View via red and cyan glasses.

Stereo Compilation

Two overlapping images that have associated Rational Polynomial Camera ([RPC](#)) models can be stereo triangulated to extract geographic or UTM coordinates using menu **Applications -> Stereo -> Compilation**. In addition, stereo compilation can be performed to generate an elevation matrix and orthographic and anaglyphic images. The RPC's relate normalized latitude or northing (U) , longitude or easting (V) and height (W) to normalized image coords as follows:



Define cubic polynomials Px1 and Qy1:

$$\begin{aligned} Px1(U,V,W) = & a0 + a1*V + a2*U + a3*W + a4*V*U + a5*V*W + a6*U*W + a7*V*V \\ & + 8*U*U + a9*W*W + a10*U*V*W + a11*V*V*V + a12*V*U*U \\ & + a13*V*W*W + a14*V*V*U + a15*U*U*U + a16*U*W*W \\ & + a17*V*V*W + a18*U*U*W + a19*W*W*W \end{aligned}$$

$$\begin{aligned} Qx1(U,V,W) = & b0 + b1*V + b2*U + b3*W + b4*V*U + b5*V*W + b6*U*W + b7*V*V \\ & + b8*U*U + b9*W*W + b10*U*V*W + b11*V*V*V + b12*V*U*U \\ & + b13*V*W*W + b14*V*V*U + b15*U*U*U + b16*U*W*W \\ & + b17*V*V*W + b18*U*U*W + b19*W*W*W \end{aligned}$$

Define similar polynomials Py1, Qy1 with their own sets of coefficients giving a total of 80 coefficients for one image.

Then the normalized image coords (X1, Y1) for image 1 are computed as:

$$X1 = Px1(U,V,W)/Qx1(U,V,W) \quad Y1 = Py1(U,V,W)/Qy1(U,V,W)$$

with a similar set of equations for image 2 with a corresponding set of RPC polynomials.

$$X2 = Px2(U,V,W)/Qx2(U,V,W) \quad Y2 = Py2(U,V,W)/Qy2(U,V,W)$$

Therefore, measuring the same (conjugate) point on both images yields 2 pairs of image coordinates (4 equations in 3 unknowns) which are solved via least squares to produce the normalized latitude or northing (U), longitude or easting (V) and height (W). Visit www.spacemaster.com for source information pertaining to [RPC](#) models. Also, search for “Rational Polynomial Camera” in www.google.com.

The steps to perform stereo triangulation are delineated below.

- Each image must have RPC data. This is needed to provide.
- Select two images from the dialog **Image 1,2** popups.
- Load the RPC's for each image via the **Load RPC** button, if necessary.
- Follow the dialog instructions. Alternately shift-click on each conjugate point.
Latitude, longitude and height will appear in the list window.
- Use button **List** to create a formatted list of the results.
- If a solution isn't possible “Invalid RPC computation” will appear in place of the coords.

When at least one point has been triangulated you can perform a stereo compilation surrounding that point or any chosen selection rectangle appearing on Image 1. Currently, Stereo Compilation is still under development and there will be changes and refinements to the dialog interfaces in the future. The Stereo Compilation dialog is shown below and the controlling details are shown in the Parameters and Recompile dialogs on the following pages by clicking the respective buttons.

The example shown here has an associated [RPC](#) file defined in UTM coordinates (North-ing, Easting). If the coordinates are given in geographics (Latitude, Longitude) then the dialog labels and values will change appropriately.

Select 2 images from the available images shown in the popup menus. If there is a text file with the same base name as the image name concatenated with [_rpc.txt](#) then it will be loaded. Else, you can load any [RPC file](#) via menu Edit -> Options -> Attach RPC File. If an image is true color then select the components to be used as the basis for all operations.

Image area surrounding the selected list point. Use Scale and Angle to orient for proper stereo parallax viewing.

Manually introduce X and Y parallax into Image 2's computed image coordinates.

Interactively offset the Northing, Easting and Height of the given selected list point.

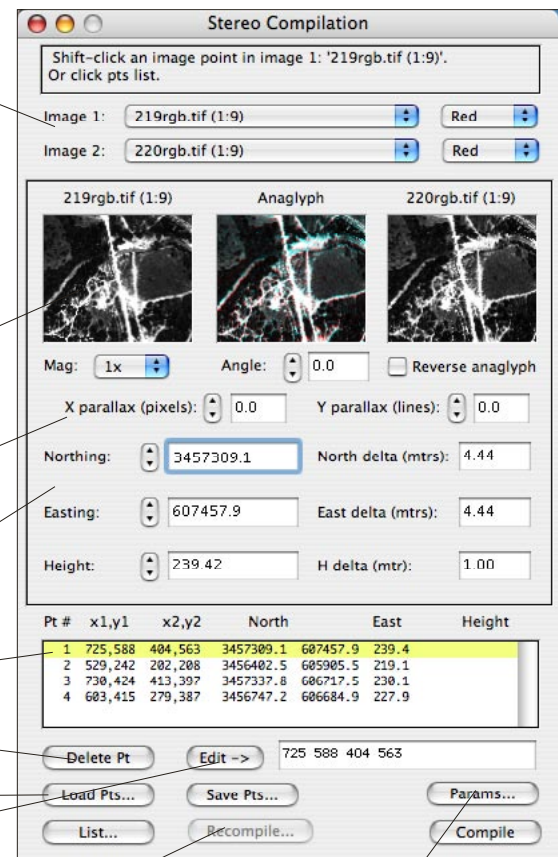
Select a point in the list and use Del button to delete.

Load/Save list points.

Generate a separate text window with formatted results.

Displays a dialog listing previous compilations that can be recompiled using various criteria.

The Params... button will be enabled after a point is selected from the list.



The following pages contain the stereo compilation results of a user selected area along with the specific compilation parameters. In this source image Northing is to the right and Easting is downward. Stereo compilation outputs are always oriented so that Northing (Latitude) is upward and Easting (Longitude, east direction) is to the right. Therefore these outputs appear rotated 90 degrees counter-clockwise relative to the image area. See section [Rational Polynomial Coefficients Files](#) for a listing of one of the two RPC files used.

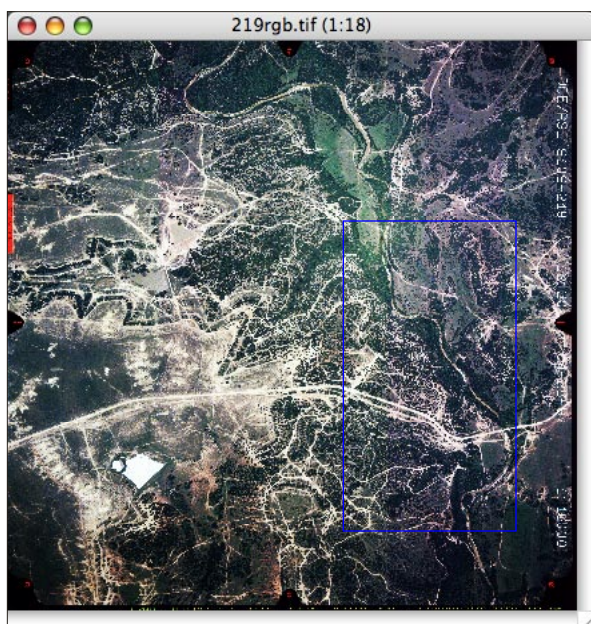



Image 1 of a stereo pair showing the selected compilation area after choosing Selection rectangle in Stereo Compilation Parameters.




Compilation area shown in blue resulting from user selection rectangle.



Elevation data image.



Ortho image of selected area in Image 1.



Anaglyphic stereo image.

Image Scale:

Lists the relationship between a 1 pixel spacing in the image and the corresponding spacing, in meters, on the ground. This is determined by the values in the RPC file and is a nominal value for the entire image.

Output Area:

User defined values that set the output size of the various images.

Pixels, Lines - Defaults to 128 by 128, changes automatically if **Selection rect** is chosen, below. This will be the size of **Image Outputs**. The orientation is always North (Latitude) is upward (lines) and Easting (Longitude) is to the right (pixels).

Output scale - This is a relative number. A value of 1 represents an output pixel equal to the Ground Sample Distance (GSD). Increasing the scale increases the area that is compiled so that an output pixel represents a greater distance on the ground. A smaller scale essentially magnifies the area to be compiled. Note: **Output scale factor** does not change the output size.

Centered around... Uses an area of **Pixels, Lines** centered on the selection point as the nominal compilation region.

Selection rect... Uses the area defined by a selection rectangle drawn upon **Image 1**, if a selection rectangle exists, as the compilation region.

Image Outputs:

Each output image is the size determined by **Pixels**, **Lines**. All images are 32 bit float except the anaglyphic image which is true color (24 bits).

DEM - Digital elevations. Needed for the [Recompile Image](#) dialog.

Ortho image - An orthographic image for each source image area.

Anaglyphic image - A red/cyan stereo image using **Image 1** ortho.

Correlation image - The linear correlation values for each elevation. Also, needed for the [Recompile Image](#) dialog.

Kernel image - The final size of the correlation kernel for each elevation.

Corr vs height - A floating point cube consisting of $\text{height}(\text{max}-\text{min})/\text{delta}$ bands of correlation. One correlation for each height plane.

Height vs kernel - A floating point cube consisting of $\text{kernel}(\text{max}-\text{min})/2+1$ bands of elevation. One elevation for each kernel size.

Height vs parallax - A floating point cube is constructed. Each plane represents a specific x,y parallax offset. If this options is selected the dialog below will appear when the **Compile** button in the main **Stereo Compilation** dialog is chosen. Enter the starting and ending x,y parallaxes and the number of steps. The compilation will proceed and generate a DEM and correlation for each combination of x,y offsets. The result is assembled into a cube. These offsets are independent of those selected in the main **Stereo Compilation** dialog which are not used. Menu [Functions -> Statistics -> Band Correlations](#) can be used to analyze the results.

Height vs Parallax

'Height vs parallax' is checked in the Parameters dialog. A sequence of compilations will be generated based on the x and y parallax offsets entered below and assembled into a cube. The number of planes is the product of the steps.

Starting x parallax:	-1	Starting y parallax:	-1
Ending x parallax:	1	Ending y parallax:	1
Number of x steps:	3	Number of y steps:	3

Cancel OK

Correlation Parameters:

These determine the correlation size (kernel), height range and various limits.

Height limit - Height determination is limited to this min, max range.

Height resolution - Change in height between correlations for a given kernel size.

Min, max kernel sizes - The minimum and maximum size correlation areas.

Height delta limit - Each change in kernel size must produce a change in height less than or equal to this value to continue to a smaller kernel size.

Corr lower limit - Stop decreasing the kernel when the correlation coefficient for a specific height determination is less than this value.


The procedure used in determining a height value follows this algorithmic scenario:

- a. Set latitude and longitude to upper left corner of output area.
- b. Choose a kernel size (starting with **Max kernel size**).
- c. Set height to minimum.
- d. Compute image's pixel values using [RPC](#) equations for each point within the the kernel and add/subtract any X and Y parallax entered in the main **Stereo Compilation** dialog. Parallax correction is applied only to **Image 2**. If **Parameter** dialog option **Height vs parallax cube** is checked then it will substitute its own dynamic parallax values.
- e. Compute the linear correlation coefficient between the 2 arrays of pixels and save. One height produces one correlation value for this kernel size and position.
- f. Increment the height by **Height resolution** and repeat steps d. and e. until the maximum height is reached. Now have an array of correlation values.
- g. Find the maximum correlation value and the corresponding height (same index). This produces an elevation for this kernel size.
- h. If the correlation value that generated this height is greater or equal to **Correlation lower limit** and the difference between this height and the height produced by the previous (larger) kernel is within **Height delta limit** then decrement the kernel size and repeat steps c. thru g. When these conditions are false or the kernel size equals **Min kernel size** then the last height satisfying the conditions is associated with the current latitude and longitude. Generally, you want the smallest kernel size that produces a valid correlation and slowly changing height vs kernel size.
- i. Increment the longitude or latitude an amount determined by **Output scale** and repeat all steps until the output area is covered.

Notes:

It very important to set the **Height limit range** so that it encompasses the expected range of elevations within the selected compilation areas. Not doing so will result in limited values in the DEM image and erroneous ortho and anaglyphic outputs. View the DEM histogram (menu [Functions -> Plot -> Histogram](#)) and check that the histogram is wholly contained in the plot. The counts should taper off to zero at each end. If not, expand the height limits in the [Stereo Compilation Parameter](#) dialog and rerun the compilation.

Each DEM output image has a title that identifies the compilation [parameters](#) that generated it. The first part is simply “DEM” followed by the number (in sequence) of the output. This number matches the correlation, ortho and anaglyph outputs so that they can be identified as a group. Each ortho contains either “a” or “b” for Image 1 or 2 respectively. If a DEM is [recompiled](#) then the source DEM number is affixed to the new DEM number and appears within parentheses. Shown below are examples of a compilation and its recompilation.



Elevation output image. The matching correlation, orthos and anaglyph images are titled:

Rho1;
Ortho1a c
Ortho1b c
Anaglyph1a c

where “c” stands for cubic resampling.

A recompilation of the DEM image shown on the left. The “(1)” in DEM2(1) indicates that is a recompilation of image DEM1. See [Stereo Recompilation](#) for details.

Stereo Recompilation:

A DEM output image can be recompiled as long as it was generated during the current [Stereo Compilation](#) dialog. This is done by choosing the DEM from a popup list and selecting the criteria that determine which of the heights in the DEM will be recompiled. Each height in the DEM is tested against the selected parameters and a new height is generated using the specified kernel size. Generally, a larger size kernel will produce an average height with a smaller deviation from the DEM's mean and a higher correlation. Note: a recompiled DEM can again be recompiled.

Shown below is the **Recompile DEM** dialog. The current options are still being developed and like the basic [Stereo Compilation](#) dialog they will change.

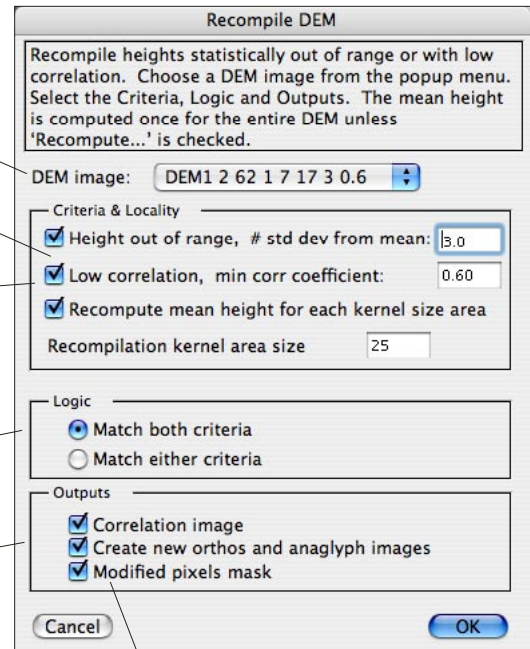
A list of previously compiled DEM output images.

Height and correlation criteria to decide which DEM values to recompile.

If checked then the mean and std are computed for each kernel size area surrounding the tested DEM pixel instead of one computation for the entire DEM.

If more than one criteria then how they are logically combined.

What new outputs are created in addition to a new DEM. If a matching correlation image was not generated when the selected source DEM was created then a new Correlation image is not created.



Generate a red mask image displaying each corresponding pixel that meets the given criteria and is recompiled.

Parameters:

Each height pixel in the DEM is tested against the chosen **Criteria** using the parameters listed below and if out of range that pixel is recompiled with the **Recompilation kernel**.




std dev from mean - Recompile each height $< \text{mean} - (\#) \times \text{std}$ or
height $> \text{mean} + (\#) \times \text{std}$.

min corr coefficient - Recompile all heights for which the corresponding correlation for that DEM was less than the given lower limit.
Note: in order to use this **Criteria** the chosen DEM image must have been compiled with the **Correlation image** option in the [Stereo Compilation Parameters](#) dialog.

Recompilation kernel - The size of the kernel to use when recompiling a pixel that meets the criteria.

Perspective

A data image of elevations can be reconstructed as if viewed from almost any perspective by applying basic photogrammetry. If there is a corresponding **Base image**, i.e., each feature pixel is one-to-one with an elevation then the features themselves can be draped across the elevations to yield a different perspective. The process is initiated by selecting **Perspective...** from the **Applications** menu which displays the following dialog. All values may be interactively modified.



- Base Image:** Feature image that is transformed into the perspective view. The **Data image** may also be used if no **Base image** exists.
- Data image:** An image whose pixels are interpreted as elevations. See section [Raw](#).
- Output Data...** If checked then a perspective view of the selected **Data image** will also be generated else only the **Base image** is constructed.
- Starting pos.:** The coordinates of the perspective or view point relative to the **Base image**. The origin is the upper left pixel with positive X in the right direction and positive Y moving downward. Z is above the image. After selecting the **Base image**, X,Y,Z are preset to above the middle of the **Base image** at a height of the maximum of X and Y. You can readily edit these values. Note that X,Y,Z is a point, not a direction, that is determined by Omega, Phi and Kappa.
- Ending pos.:** If **Generate a sequence**, below, is checked then this is the ending view point of a uniform set of view points and angles that will be used to generate a synthetic cube movie.
- Generate a sequence:** If checked then # equally spaced view points are constructed and assembled into a cube.

Omega, Phi, Kappa: The rotations about the X, Y, and Z axes respectively in degrees. If **Link Omega & Phi** is checked then the perspective orientation is fixed pointing to the center of the **Base image**. Therefore, a **Sequence** will produce a fly-by perspective with an invariant center.


Link Omega & Phi to X,Y,Z: If checked then **Omega** and **Phi** are automatically computed as a function of **X** and **Y** so that the perspective view is pointed at the center of the **Base image**. You can check this box to force a centered view, click an **Outline** button and then uncheck it to modify **Omega** and **Phi** manually for a more desirable view. See the **Outline** buttons described below.

Data scaling: The multiplicative factor applied to the Data. The smaller the number the “flatter” the perspective appears. Usually, a value of .05 is a good starting point. The larger the scaling the more exaggerated the effect of elevation and the longer the computation time.


Focal length: How far behind the perspective point the imaging plane lies. The greater the focal length the larger the nominal magnification. A small **Focal length** is analogous to a wide angle lens. The initial value is set equal to the **Z** coordinate and produces a one-to-one scale. Changing the **Focal length** does not change the output image dimensions.

Reset: Set the perspective back to the initial starting values as if viewing the image from a point directly over the center with no rotations.

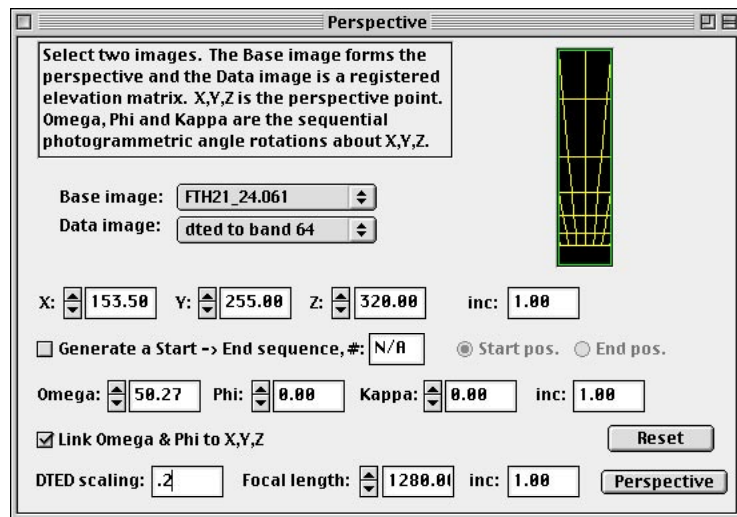
Perspective: Generate the perspective image(s).




A shaded relief image generated from elevation data is used as the Base image and provides the features that will be viewed from a different perspective as shown to the right.



The reconstructed perspective image using the photogrammetric parameters entered into the previously described user dialog.




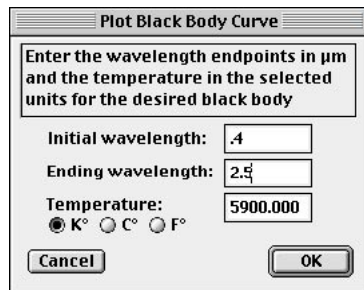

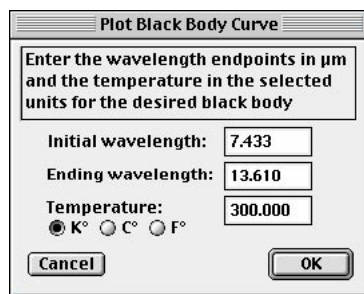

The above dialog was used to select and generate the perspective image shown here. Both images were then rotated 180 degrees to produce the final views shown below. Note that Link has been selected to force pointing at the center of the source image.



Radiance



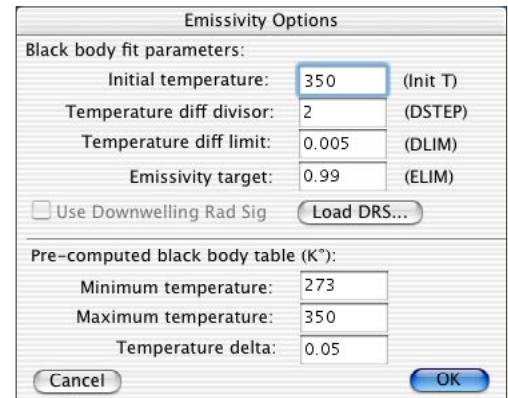
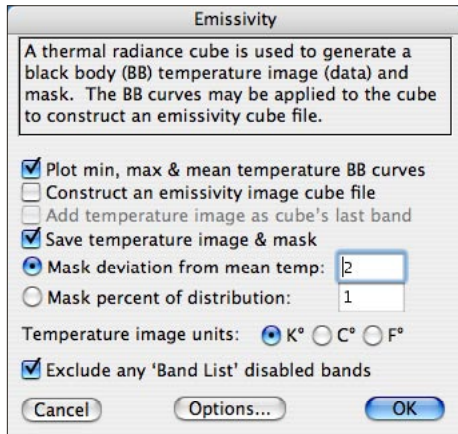
Applications -> Radiance -> Emissivity processes thermal images, generally in the range of 3 to 18 microns, that are in radiance units of $\text{Wm}^{-2}\mu\text{m}^{-1}$. Structurally, there is no distinction between a thermal image cube and a visible to SWIR cube, only the wavelength definition. HyperCube allows you to generate a temperature data image and optionally, an emissivity cube. The temperature is derived by comparing the radiance of each image cube pixel with the radiance of a series of black body curves. The closest match, via a user specified tolerance, produces the temperature for that pixel. There is also a utility, Plot BB Curve, that can display a black body radiance curve for any given temperature.

Shown below are the various user dialogs. Note that the routines require wavelengths to be in microns (μm) but that the temperature can be in centigrade, Kelvins or Fahrenheit.




Menu Applications -> Radiance -> Plot BB Curve dialogs, on the left, and their corresponding black body radiance curves on the right for the given temperature and wavelength ranges.

Menu Applications -> Radiance -> Emissivity displays the following dialogs. The default values shown were used to produce the displayed outputs from a thermal image cube.



The output image at left is a mask showing the warmest, red, and coolest, blue, regions within the thermal image cube, i.e., those temperatures 2 or more deviations from the mean. At right is a shaded relief of the temperature image that has been similarly coded. The coolest region coincides with a creek. Below is the histogram of the temperature data image in units of degrees Kelvin.





Menu item Applications -> Radiance -> ISAC performs the “In Scene Atmospheric Compensation” (ISAC) algorithm on a thermal hyperspectral image which as in the Emissivity function must be in radiance units of $\text{Wm}^{-2}\mu\text{m}^{-1}$. This function attempts to remove the effects of the atmosphere and is described in publication: **JGR-Atmospheres: Young, Johnson and Hackwell, “An in-scene Method for Atmospheric Compensation of Thermal Hyperspectral Data”, 21 December 2002.**

The current options available are shown below.


A synthetic cube is generated where each plane is an sdtol mask dependent on a value from the sequence. No other output is produced.

Only bands between these limits will be used in computations.

If checked, any bands that are disabled in the Windows->Band List and will also be excluded from computation.


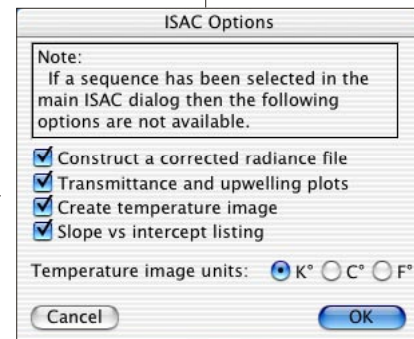


Source image




Black body output mask for an sdtol value of .0045.

Various possible output products that may be selectively enabled if sequence has not been selected in the main dialog.



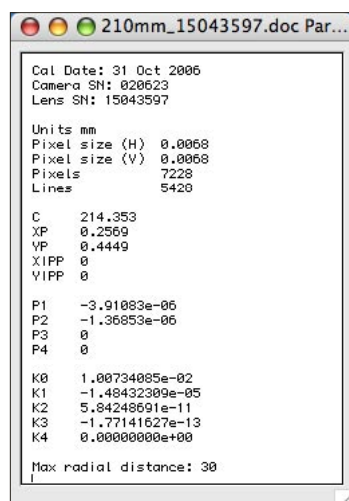

Graphs of transmittance and upwelling versus wavelength respectively.



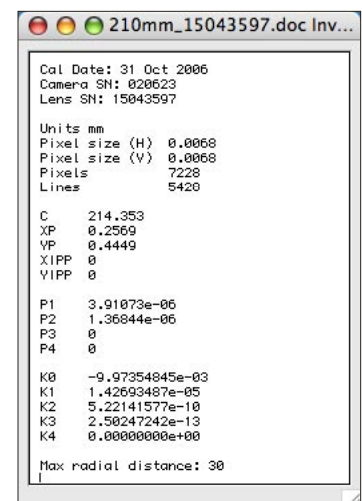
Lens Distortion

A digitally acquired image can be geometrically corrected for lens distortion by selecting menu Applications -> Lens Distortion. HyperCube uses the parameters as defined by the USGS/EROS Report of Camera Calibration (visit web site: <http://erosproject.cr.usgs.gov/osl>). Note that the USGS parameters define corrections, i.e., the uncorrected image coordinates are transformed via these parameters to produce an x and y delta that is added to the measured distorted coordinates to yield a distortion free x and y coordinate. HyperCube uses the USGS parameters to generate new parameters that correspond to the inverse transformation that maps a distortion free output space into the distorted space. Besides generating a corrected image in memory you can also read, correct and store an entire directory of [TIFF](#) images.

Shown below is the main dialog window. It contains controls to load, view, edit and plot the USGS reports. You can select an in-memory image for correction as well the source and destination directories for batch processing.




The result of selecting the List button. The extracted parameters from the calibration file are shown on the left. The computed inverse parameters are displayed to the left. Parameter names corresponds to those in the USGS




Selecting Plot from the Lens Distortion dialog generates the plots shown below for the loaded USGS Report. Those on the left display the radial and decentering correction deltas that are added to a measured x,y coordinate to produce a corrected x,y pair. The right hand column of plots are generated by computing the coefficients of the inverse equations defined by the Report. Although the Report uses focal plane millimeters, axes are shown in pixels. In this case, for a focal plane array that is 7228 pixels by 5428 lines with a pixel and line spacing of .0068 mm and a maximum valid radius of 30 mm from the principal point.


Each plot represents a correction (delta) that is added to an image coordinate to produce a new x,y coordinate. The left column corrections are applied to the distorted image to produce an distortion free position for that pixel value. The right column is the inverse: given a distortionless coordinate, where is the distorted pixel (coordinates) that belongs there. The inverse transformation is used to create the new distortion free output image. It will be the same size as the input image. Any boundary voids that may arise are replicated from the source image.




Radial distortion correction.




Radial inverse correction.




X decentering correction.



X decentering inverse.



Y decentering correction.



Y decentering inverse.

The **Edit** button on the **Lens Distortion** dialog displays each parameter in an edit field as shown in the dialog below. These values are very specific to the equations in which they are used and are in millimeters (HyperCube uses **Pixel size** to convert everything to digital values). Specifically, the USGS Report uses them as follows (more details can be found at <http://erosproject.cr.usgs.gov/osl>):

$$\text{Xbar} = (X - \text{XP}) \quad \text{Ybar} = (Y - \text{YP}), \quad \text{where } X, Y \text{ are measured coords} \\ \text{and } \text{XP}, \text{YP} \text{ is the principal point}$$

$$R = \sqrt{\text{Xbar}^2 + \text{Ybar}^2}, \quad \text{where } R \text{ is the radius in mm}$$

Radial correction:

$$\text{Xc} = \text{Xbar} * (\text{K0} + \text{K1} * R^2 + \text{K2} * R^4 + \text{K3} * R^6 + \text{K4} * R^8), \\ \text{Yc} = \text{Ybar} * (\text{K0} + \text{K1} * R^2 + \text{K2} * R^4 + \text{K3} * R^6 + \text{K4} * R^8),$$

Decentering correction:

$$\text{Xd} = (1 + \text{P3} * R^2 + \text{P4} * R^4) * (\text{P1} * (R^2 + 2 * \text{Xbar}^2) + 2 * \text{P2} * \text{Xbar} * \text{Ybar}), \\ \text{Yd} = (1 + \text{P3} * R^2 + \text{P4} * R^4) * (\text{P2} * (R^2 + 2 * \text{Ybar}^2) + 2 * \text{P1} * \text{Xbar} * \text{Ybar}),$$

Final new coordinates:

$$\text{Xcorrected} = X + \text{Xc} + \text{Xd} \\ \text{Ycorrected} = Y + \text{Yc} + \text{Yd}.$$

Note that **Focal length** and **Max radial distance** are not used in the equations. The later parameter defines the limit of the corrections accuracy.

Lens Distortion Parameters

Edit/modify the USGS lens distortion parameters. All values must correspond to mm measurements as defined in the Lens Calibration Report: (visit <http://erosproject.cr.usgs.gov/osl>).

Pixel size (H mm):	<input type="text" value="0.0068"/>	Pixels:	<input type="text" value="7228"/>
Pixel size (V mm):	<input type="text" value="0.0068"/>	Lines:	<input type="text" value="5428"/>
XP:	<input type="text" value="0.2569"/>	XIPP:	<input type="text" value="0"/>
YP:	<input type="text" value="0.4449"/>	YIPP:	<input type="text" value="0"/>
P1:	<input type="text" value="-3.91083e-06"/>	P3:	<input type="text" value="0"/>
P2:	<input type="text" value="-1.36853e-06"/>	P4:	<input type="text" value="0"/>
K0:	<input type="text" value="1.00734085e-02"/>	K3:	<input type="text" value="-1.77141627e-13"/>
K1:	<input type="text" value="-1.48432309e-05"/>	K4:	<input type="text" value="0.00000000e+00"/>
K2:	<input type="text" value="5.84248691e-11"/>		
Focal length:	<input type="text" value="214.353"/>	Max radial dist:	<input type="text" value="30"/>

The result of selecting Edit in the Lens Distortion dialog. Edited parameters are checked for blunders and consistency when OK is clicked.

Miscellaneous

Icons

HyperCube has a family of 8 Finder icons. Five of these are regular Mac data types and represent TEXT, JPEG, PICT, TIFF and Sun Raster Format files. The other three represent the HyperCube application, unformatted binary data and PhotoShop™ convolution kernels as shown in the figures below.



HyperCube application icon.



ASCII text file, e.g., *.hdr, *.wvl.



Image files saved in JPEG, NITF, Pict, SRF, and TIFF formats.



Unformatted image file saved as Raw.



PhotoShop™ convolution kernel of type 8BCK.


Double clicking an image icon will automatically launch HyperCube and load the corresponding image. Similarly, a *.wvl or *.hdr icon can be double clicked as well.

The Windows version does not have an icon family, only the application has an icon. Any HyperCube documents must be opened within the application.

Error Messages

When HyperCube detects an internal error such as running out of memory it generally can recover enough to allow you to at least save and print. There is one serious type of error dialog that may be presented indicating that an immediate exit is warranted. But if you have generated some result that just must be saved or printed then HyperCube can be forced to continue at risk of an immediate crash. Generally, the exit path should be taken especially if it is followed by another dire warning.

This error message will not appear in the Windows version.



Capturing the Screen and Listing Pixels

The contents of the main display screen may be captured at any time other than a modal dialog (i.e. those demanding attention) by simultaneously pressing the command 1(one) keys (Ctrl 1 in the Windows version). This creates a new true color window of the entire display that may be saved like any other pict image. Note that a full screen color window requires approximately 2MB for 832 pixels by 624 lines. HyperCube checks for sufficient memory and warns if it is insufficient.

Additionally, in the Windows version you can automatically copy the entire active window to the clipboard by depressing Ctrl Print Screen. Then choose File -> New... followed by Edit -> Paste to place it in a window where it can be saved.

Pressing key combination: command 4 on the Mac and Ctrl 4 on the PC when the cursor is positioned over an active (front most) image window will create one or more text windows containing a listing of pixel values surrounding the cursor. The display values and the underlying pixel (e.g., 16 bit, float) values are listed in a 15 by 15 array. True color images produce a separate list window for each color.

Saving and Copying Scrollable Lists

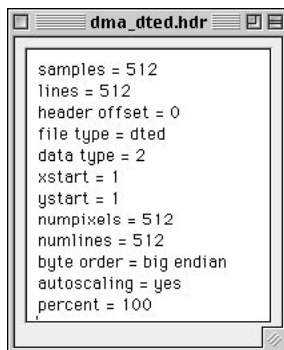
Any window which has a scrollable list, e.g. spectral plot, band correlations, that does not contain a specific **Save** button may still be written to a file by selecting **Save** from the **Edit** menu. This will bring up the standard file save dialog. The data will be formatted as tab delimited ASCII text and can be reread by any application that reads text. Also, this same scrollable list may be copied to the clipboard as tab delimited, c/r terminated lines of text by choosing the Edit -> Copy menu and then pasted into another application.

Loading/Saving Text Files

Any small text file can be opened and displayed in a scrollable list. This is especially useful for examining header files. The file is opened by choosing menu File -> Open As and **Text** from the format dialog popup. It can then be edited (**Cut**, **Copy** and **Paste**) and resaved. The up and down arrow keys and the page up, page down keys are used to navigate through the text. Choosing File -> New... and selecting the **Text** option will open a blank **Text** window that may be typed into and saved as text.


An image or any selection may be saved as an ASCII text file by choosing menu File -> **Save As** and **Text** as the format (and optionally as a matrix, see section [Apply Transfor-](#)
[mation](#)). The pixels are tab delimited with each line followed by a return. The size of the resulting file may be 4 or 5 times larger than the source.

Note: Any embedded zeros within a text file are automatically removed when the file is read. Carriage returns on the Windows version are followed by a newline character.



*Using menu File->Open As, Text and selecting a header file, *.hdr, results in the window at left. This may be edited and saved, functioning as a simple word processor. Dragging the cursor within the window's content area will automatically scroll exposing additional lines as will the arrow and page keys. Text will either wrap-around or limit itself to a line if a carriage return occurs within the width of the window. The window may be resized as necessary.*


HyperCube Menus



The Windows menu will contain Cascade and Tile in the Windows version but not the Mac version.


The availability of a specific menu function is directly keyed to the current active window. If an item is not enabled and you think it should be, it's most likely because it is not applicable to the front most window or is ambiguous. For example, **Embed Wavelength** is enabled only when a cube window is front, likewise for **Classify**. **Fourier** and other operators can only be used when an image window is active and **Convert Image** is very dependent upon the active window's bit depth.

Some menu functions produce operators that become attached to the window that is active. **Color Map**, **Gray Map**, **Filter**, **Flicker Cube**, and **Fourier** are examples of attached operators. When the parent window is closed the operator window disappears.



Windows version main Menu Bar.

New... Create a new empty image window that can be used to paste the current clipboard. You are presented with the following dialog options which are preset according to the contents of the clipboard.



Pixels Width of window

Lines Height of window

Text Create a new Text window

Image Create a new Image window.

Gray scale 8 bit gray window

True color 32 bit full color window

Open... Presents the standard Mac open file dialog showing recognized files.

Open As... A modified open file dialog which includes [TIFF](#) and other image files.

Load Sel... Available when an overview is front most. Brings up the last dialog used to open the source image with the requested size fields set to correspond to any selection rectangle that appears on the overview. See section [Overviews](#).

Close Close and dismiss the currently front most window. Any modified window first presents an opportunity to save it.

Close All Closes all open windows.

Save Save the current active image window as a Pict file.

Save As... Present a modified save dialog allowing you to save in Pict and other type formats.


Page Setup... Bring up the standard printer setup dialog for the chosen printer.

Print... Display the print dialog and print the contents of the front window.

Quit/Exit Exit HyperCube. A warning is given to save various results.

Edit


- Undo** Not implemented in the present HyperCube version.
- Cut** Also not implemented as yet.
- Copy** Copy the current front window to the clipboard. If a selection rectangle is present then copy just its contents.
- Paste** Paste the contents of the clipboard to the front window. You can position it before embedding.
- Clear** De-toggles enabled items in Library list, Classify points list and Cube points.
- Select All** Toggles on all items in a Library list or Classify points list.
- Duplicate** Make an exact copy of the current front window contents but not its functionality. Equivalent to performing: [Copy, New and Paste](#).
- Send to Back** Place the current front window behind all other HyperCube windows.
- Options** A submenu with the following sub-functions.




- Sound** Toggle alert sound on/off.
- Detailed Messages** Toggle image file header info on/off.
- List Allocation** Display a snapshot of internal structures.
- [0,0] Image Origin** Change the upper left coordinate origin from [1,1] to [0,0]

Labels

A submenu specifying the font, size and position of dynamic labels (see [Dynamic Displays and Plots](#)).



Mag/Rot Interp... Interpolation for magnify and rotate (see [Image](#) menu).



Color Selector... Display the selection point to color correspondence [editing](#) window.

List Clipboard... What's currently in the clipboard.

Pixel Exclusion... Displays a dialog specifying certain pixel values that are excluded from histogram computations. See section [Pixel Exclusion](#).

Image to Ref... Display a dialog defining the correspondence between an image coordinate and a reference coordinate.

Geo/UTM Grids... Overlay geographic or UTM lines on a [referenced image](#).

Attach RPC ... Present an Open File dialog to select a Rational Polynomial Coefficients ([RPC](#)) file and attach it to the active image.


Create Geo RPC... Generate Rational Polynomial Coefficients ([RPC](#)) for an image containing a Geographic or UTM [reference](#).

Geo Coords <--> UTM Coords...


Convert a referenced Geographic image to a [referenced](#) UTM coordinate system and vice versa. If the image contains Rational Polynomial Coefficients ([RPC](#)) then also convert them.

Applications

Shaded Relief... Solicit the elevation and azimuth of an artificial light source and generate a reflective relief of an image or its attached data.



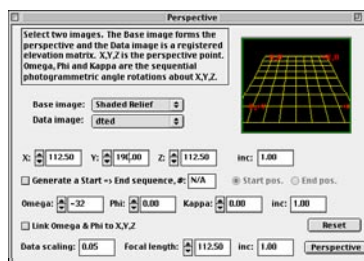
Contour... Generate contour lines from an image and display the results.




Stereo... Produce an anaglyphic red and cyan stereo image from elevations. Generate elevations and orthographic images.

Mate...
Compilation...

Perspective... Generate a perspective view given a base image and its matching digital elevations.




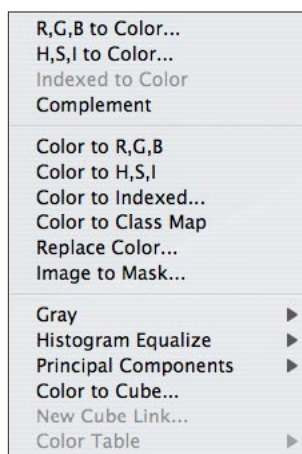
Radiance... Construct a new radiance image cube from an existing thermal radiance cube using various techniques.



Lens Distortion... Apply USGS/EROS Report of Camera Calibration to a digitally acquired image.

Image

Convert Image A submenu used to change the structure of an image.



R,G,B to Color...	Select 8 bit images to form a color image.
H,S,I to Color...	Select 8 bit hue, saturation and intensity to form a color image.
Indexed to Color Complement	Changes 8 bit pseudo color to true color.
Color to R,G,B	Complement a color image.
Color to H,S,I	Separate a color image into planes.
Color to Indexed	Separate a color image into hue, saturation and intensity.
Replace Color...	Change true color to pseudo color.
Image to Mask...	Substitute one color with another.
Gray	Change a specific color (gray) to a mask.
Hist Equalize...	Convert color to gray image submenu.
Principal Comp...	Perform various equalizations.
Plane to Cube	Convert color to principal components.
Color to Cube	Use a gray image as cube base. This menu item toggles with Color to Cube.
New Cube Link...	Convert true color image into a 3 band image cube.
Color Table	Change the present cube color composite link.
	Change present 8 bit pseudo color table.

Attach A submenu that associates one image's characteristic with another.

Data Copy data from one image to a second image.

Class Map Make a virtual class map using a class map.

Remove Data Discard any attached 16 or 32 bit data.
Also used to convert 48 bit color to 24 bits.

Add to Cube... Presents a dialog list from which to choose images for a new cube.

Zoom in Magnify the current image window by 2.

Zoom out Demagnify the current image window by 2.

Flip A submenu allowing images and cubes to be reversed and reflected.




Left -> Right Reverse the pixels within a line.

Top -> Bottom Reverse the order of the lines.

Band Order Change the ordering of cube bands.

Rotate Multiple 90 degree rotations of image and cube windows.



90 cw Rotate image or cube 90 degrees clockwise.


90 ccw 90 degrees counter clockwise.

180 cw 180 degree rotation.

General User input of rotation.

Shift

Left and right pixel and line cyclic shifts of image windows.



Left, Right, Up, Down Direction of shift.

Even Just even numbered line.

Overwrite If toggled, results overwrites source.

Enter Shift Allows user specified shift value.

Color Map...

Display interactive dialog of color balance controls for color images.

Gray Map...


Change the contrast and brightness of any image or cube window.

Data Map...

Select the way in which data values are mapped for display.

Filter

A submenu of static and interactive spatial domain image filtering.



Static 5x5...

Display a static 5 by 5 kernel convolution static dialog.

Dynamic 5x5...

Display an interactive 5x5 kernel convolution dialog.

Magnify ...

Generates a new magnified/minified view of the image or selection rectangle by the factor chosen in the submenu. Certain plots (e.g., data histograms, eigenvalues) can also be magnified.

Delete Band

The current cube face band is deleted from the image cube display. The cube will be disconnected from any file reference.

Disconnect...

The current image cube is disconnected from any file reference. It becomes an in-memory only image cube.


Embed Wave...

Write the wavelength into each band of the loaded cube. Useful before Save Bands function.

Cube Color Composite

A submenu enabled when a cube image is front.

Generates a true color image using the values set in **Definitions...** to determine which bands of the cube file to use. The new color image is automatically linked to the cube (see sections [Color Composite](#) and [Classify Functions](#)).



Specific Wave

Use the 3 closest bands to the values set in **Definitions...**

Interval Wave

Integrate all bands lying in each interval to form the red, green and blue components.

Definitions...

Specify the red, green and blue wavelengths for the color composite.

Functions

Flicker Cube Generate a new cube window devoted to flickering between bands.

Flicker/Superimpose Display a dialog of image names from which to choose the sources for an image flicker or image superposition.

Statistics A submenu of statistical functions to be applied to the current cube.

Pair Correlations...
Band Correlations...

Pair Corr... Cross correlation computation between 2 images.

Band Corr... Cross correlation between successive cube bands.

Plot

A submenu of graphical operators.

Histogram ►
Spectra (Selected Points)
Spectra (Dynamic)
Scatter...
Profile
Scaling...

Histogram... Perform a pixel distribution histogram of the current image or cube.

Hist (Mask... Same as above using an image mask.

Spectra (Se... Display the spectral plot window for the selected active cube points.

Spectra (Dy... Dynamically display a spectral plot as the cursor moves over a cube.

Scatter... Generate a 2 dimensional plot of the pixel distribution between 2 images.

Profile Draw a gray level density plot of the active image profile line. If the image is color draw plots for red, green and blue.

Scaling... Selectable when a plot window is active. Allows you to change the plot axes.

Arithmetic... Open an interactive window allowing arithmetic expressions to be performed on images.

Overlay Color... Generate a composite image from a mask and base image.

Adjacency... A submenu for comparing and outputting spatially near pixels.

Classify... Open up the classification dialog window for the current image cube.

Calibration... Open the calibration dialog for the frontmost **Classify** window.

F(Signatures)... Compute eigenvector matrix of **Classify** signatures.

Fourier... Display an interactive dialog to perform frequency domain filtering upon the current 8 bit image window.

Mosaic Display interactive dialogs to select and control the merging of two or more images into a composite assemblage.

Warp... Display an interactive dialog to select and control the geometric transformation of one image to another.

Utilities

Reformat File... Open a dialog for reformatting the spatial and/or pixel type of a file.

Wrap JPEG... Create a NITF file containing an embedded JPEG (2000) file.

GeoTIFF, JPEG... Combine GeoTIFF geometry with JPEG to form NITF file.

ASCII to Library... Read and convert a columnar text file to a Library file.

Library to A submenu to convert the current front most library list window into an image cube file or an ASCII text file.

Image Cube File...
ASCII File...

Class Map Editor... Enabled when a classification image is front most. Permits changing the region color and name association.

Concatenate A submenu used to select two images that will be concatenated in either pixel or line order to form one new image.

Images, Bands...
Cubes...

Concatenate Files... Display a dialog allowing you to select multiples files to be concatenated into a single sequential new file.

Change Window Title... Solicit a new title for the current active window.

Apply Transformation... Multiply a cube file by an orthogonal transformation.

Merge LAS Files... Combine the point clouds of several LAS files into a new LAS point cloud file.

ASCII LIDAR to LASF Convert an ASCII LIDAR file to a LASF file.

Windows

Show Info Display a small dynamic window that shows various HyperCube details as the cursor passes over other windows.

Show Band A submenu showing wavelengths and annotation of multiband images.

List Display a listing of the wavelengths of the current image cube to allow editing.

Annotation Display a listing of the annotation associated with each band that you may edit.

Cascade Windows version only. Arranges all of the open windows in an overlapped diagonal order.

Tile Windows version only. Resizes and rearranges the open windows in a row by row pattern.

“Names” A dynamic list in order of creation of HyperCube windows. The front most window is preceded by a check mark.

Index

Symbols

- % 37,50
- (1) 177
- *.c2n file 147
- *.hdr 35,106,154,156,188
 - band list 106
- *.wav 154,156,188
- *.wvl 105
 - band list 105
- .ev 138
- .gs 138
- .pm 138
- .tfw
 - ESRI world file 26
- .wav 49,105
- .wvl 49,105,154,156,188
- [0,0] 63
- [0,1] 38,114
- [1,1] 63
- {constant,constant,constant} 85
- 180 cw 195
- 8BCK 188
- 90 ccw 195
- 90 cw 195

A

- Abs diff 121
- Add to Cube... 195
- Adjacency 91
- Affine 103
- Algorithms 121
- Alpha
 - PNG files 25
- Anaglyph 168
 - point cloud 19
 - SDF 30
- Anaglyphic
 - Stereo Compilation 173
- angle
 - profile 80
- Arithmetic 135,197
 - sequence 86
- Arithmetic functions
 - non standard 82
 - standard C 81
- Arithmetic operators 81
- Arrow keys
 - flicker offset 109
- ASCII 99
 - dialog 18
 - LIDAR 15
 - Saving images 189

- Attach Data 71
- Averaging
 - spectral plots 119
- Azimuth 168

B

- Bands
 - loading 153
- Band Annotation 107
- Band Correlations 74,197
 - sequential 74
 - specific 74
- Band List 105,121,151
 - revert 106
 - scaling 106
- Band Order 58,195
- Band Scrolling 108
- Binary 68
- BIP
 - FWF cube 30
- Black body 182
- Block size 160
- Brightness 53
- Buckeye 159
- Byte 153
- Byte order 36,156

C

- Calibration 150
- Cascade 198
- chunks
 - PNG files 25
- Classify 93,94
- Classify... 121,190,197
 - overlays 135
- Classify Each 128
- Classify Mean 128,140
- Classmap Editor
 - Color to Classmap 71
- Class map
 - virtual 145
- Class Map Editor
 - ESRI Shape Files 12
 - singularities 94
- Clear 52,114,192
- Clear Lib 114
- Close 191
- Close All 191
- Coefficients 51
 - Daubechies 132
 - Haar 132

- Color Balance 55
- Color Composite 126,130,196
 - color keywords 113
 - save hdr 113
- Color depth
 - reduction 71
 - Remove Data 71
- Color domain
 - histogram match 70
- Color Image 131
 - 16 bit components 26
- Color map
 - histogram 76
- Color Map... 55,190,196
- Color Mapping
 - hue, sat, int 55
- Color Mask 131
- Color overlay 67
- Color Selector 128,130,144
- Color separation 68
- Color Table 195
- Color to cube 111
- Color to Indexed 66,195
- Color to Library 144
- Color to Planes 64,195
- Columns 166
- Command-option-C 67,116
- Complex data
 - HDF 13
- Concatenate Files 164
- Concatenate Images 163
- Constructing a cube 111
- Contours 169
 - ESRI Shape file 12
- Contrast 50,53
- Control-click
 - profiles 80
- Convariance
 - noise 124
- Convert
 - 16 bits to 8 71
- Convert Image 195
- Convolution 83
- Convolver
 - 3x3 61
 - 5x5 61
 - offset 61
 - scale 61
- Coordinates 63
 - mosaic references 100
- Coordinate Permutation
 - heading,roll,pitch 29
- Copy 52,192
- Copy color 116
- Corner coordinates
 - image to reference 45

- Correlation 72,74
 - band statistics 74
 - pairs 72
- Corr coef 121
- Countable 130
- Covariance 74,122,124
 - band statistics 74
 - F(signatures) 138
 - Mahalanobis 123
 - Match Filter 123
 - pairs 72
 - principal components 65
- Crash 188
- Cross correlation 83
- Cut 192
- C language 81

D

- Daubechies 132
 - coefficients 83
- Delete 52,53
- Delete Band 112
- Density plot 80
- Deriv diff 121
- Discrete Wavelet Transform
 - image arithmetic 83
- Detailed Messages 8,24,26,76,192
 - cummulative histogram 76
 - ERDAS Imagine 11
 - HDF tags 13
 - histogram 76
 - JPEG 14
 - LASF 18
 - PNG files 25
- dft_nterms 95
- dft_phase 95
- Difference image 133
- Difference plot 134
- Discrete Wavelet Transform
 - compression 89
- Double clicking 105
- Dpi 155
- Dragging a selection 52
- Drag & Drop 120
- Duplicate 52,153,192
- Dynamic
 - filters 96
- Dynamic Displays 114

E

- Easting
 - ASCII LIDAR 167
- Edit 114
- Eigenvector 138
 - signatures 138

- Elevation 168
 - Stereo Compilation 173
- Embed Wavelengths 190,196
- Emissivity 182
- Enable Dynamic Display 146
- Energy
 - wavelets 84
- ENVI®
 - header file 42
 - wavelength file 49
- ERDAS Imagine® 10
- Error Messages 188
- ESRI 12
- ESRI Shape File
 - exporting 147
- Euclid distance 121
- Excel™ 74
- Export shapefile 149

F

- False alarm 140
- False hits 93
- Fast Fourier Transform 95
- Filter 196
- Filtering
 - 3x3 60
 - 5x5 60
 - convolution 60
 - denoise 129
 - frequency domain 95
 - interactive 60
 - spatial 60
 - wavelet 132
- Flicker
 - arrow keys 109
 - x,y offsets 109
- Flicker Cube 109,197
- Flicker Image 109
- Flip 58,195
- Floating point 36,153,156
- Fly-by 180
- Focal length 180
- Form Neighborhood... 127,140,142
- Fourier 95,121,190,197
 - high boost 96
 - low pass 96
 - no DC 121
 - phase 121
 - phase delta 121
 - power spectrum 96
- Frequency
 - histogram 137
- Frequency domain
 - functions 83
- Full Wave Form (FWF) 27

G

- Gains & Offsets 39
- Gains & offsets 157
- GeoTIFF 26
- GPS 27
 - SDF 27
- Gram-Schmidt 138
- Grayscale 195
- Gray Map... 53,190,196
- Green dot 141
- Ground truth 131,140

H

- H,S,I 64
- Haar 132
 - coefficients 83
- Harmonic 121
- HDF
 - Hierarchical Data Format 13
- Header 35
 - *.hdr files 42
 - saving 37,106
- Header file
 - image to reference 45
- heading
 - profile 80
- Height
 - direction 46
 - RPC 46
- HFA 11
- Hierarchical File Architecture 11
- Histogram
 - color mapping 55
 - cumulative distribution 76
 - Detailed Messages 76
 - exclusions 78
 - file 76,77
 - gray mapping 55
 - masked 77
 - of classifications 134
 - Plot 197
 - summary 76
 - window 76,77
- Histogram equalization 69,113
 - matching 69
 - uniform 69
- Hit rate 140
- Hotelling 65
- HSI
 - K-means 125
- Hue 64
 - color mapping 55
- Huffman tables
 - PNG files 25

I

- Icons
 - appl 188
 - Pict 188
 - Raw 188
 - Text 188
- Icords
 - keyword 44
- Image Arithmetic 81
 - ** operator in 81
 - abs,...,atan2 81
 - remapping 85
 - shift_left... 82
- Image cube
 - adding to an 111
 - constructing an 111
- Image Shift 72
- Image to Mask 68
- Image to Reference 149
 - change in origin 45
- Indexed to Color 66
- Info
 - band statistics 74
- Info Window 63,114
 - color mapping 55
 - cursor position 63
 - NITF coordinates 24
 - Show Info 63
 - spectral coords 114
- Intensified color
 - classify 133
- Intensity
 - color mapping 55
- interlace
 - PNG files 25
- Interleave 38,153,160
- Interpolate 30
- Interpolation 103
 - mosaic references 101
 - point cloud 20
- Interval Wavelengths 113
- ISAC 184

J

- JPEG 24,26,66
 - baseline 14
 - compression 14
 - embed 158
 - markers 14
- JPEG DCT 26,159

K

- K0... 187
- Keywords
 - definitions 43
- Kmeans 121,125
 - classes 125
 - class color 125
 - iterations 125
 - options 125
 - plots 125

L

- LASF Files 15
- Latitude 44
 - Stereo Compilation 172
 - stereo triangulation 171
- Latitude/Longitude
 - Conversion 47
- Left 59
- Left -> Right 58,195
- Lens distortion
 - decentering 187
 - focal length 187
 - pixel size 187
 - principle point 187
 - radius 187
- Lib:File
 - match lists 144
- Library
 - NEFDB 117
 - USATEC 117
- Library... 127,143,160
 - List 160
- Library to Cube 160
- Lib Key 144
- LIDAR 15
 - Z values 16
- Linear 53
 - correlation 72
- Link 71
- Linked image 126,196
- Link to Cube 64
- List allocation 192
- Load
 - Bands 112
- Loading
 - by double clicking 105
- Longitude
 - Stereo Compilation 172
- longitude 44
- lossless
 - PNG files 25
- LZW 26

M

- Magnify
 - bilinear 62
 - cubic 62
 - minification 196
 - popup menu 62
 - replicate 62
 - selection rectangle 62
- Magnify Selection 196
- Mag Interpolation 193
- Mahalanobis
 - classifier 121
 - classify 121,123
 - distance 121,123
 - list 121,123
 - parameters 129
- Mask 67,68,90,195
 - creating 124
- Mask image 30
- Matched filter 122
 - abundance 121,122
 - mask 121
- Match Criteria 130
 - Countable 130
- Matrix 166
 - eigenvector 138
 - orthonormal 138
 - projection 138
 - save 166,189
- Mean signature 133
- Menus 190
- Merge LAS Files 167
- Meters per pixel 80
- Method 143
- More Options 133
- Mosaic 97
 - control 99
 - options 99
- Multiband 35

N

- NEFDB 117
- New 53,191
- NITF
 - coordinates 24
 - mono 156
 - Overviews 41
 - reformatting 156
- Northing
 - ASCII LIDAR 167
- NTSC
 - gray() 82

O

- Omega, Phi, Kappa 179
 - point cloud 20
- Open... 191
- Open As... 8,160,191
- Operator 55
- Options 192
- Orthogonal 97,103
- Orthonormal 138
- OS X
 - Pict pseudo color 66
 - TIFF 8
- Outline 98
- Overlay 90
- Overlay Color 93
- Overview
 - Load Selection 191
 - point cloud 19
 - TIFF 24,26
- overview
 - LASF 15
 - LIDAR 15
- Overviews
 - scale factor 40
 - scaling 40
 - selection rectangle 40
 - starting band 40

P

- P1... 187
- Page Setup... 155,191
- Pair Correlations... 72,197
- Palette 116
- palette
 - PNG files 25
- Paste 52,192
- Percentage 130
- Perspective 179
- PhotoShop™ 61,188
- Pict 52,188
- Pixel Exclusion
 - contour 169
 - filters 60
 - histograms 76
 - shaded relief 168
- Pixel size
 - lens distortion 187
- Pixel type 36,153
- Planes to Color... 64
- Plane to plane 97
- Plot 143,197
 - profile 80
 - scaling 104
 - scatter 197
 - selection 104
 - spectra 114,120

- Plotting Signatures 150
- Plot Scaling 114
- Plot transform
 - histogram match 70
- PNG
 - Portable Network Graphics 25
- Point Cloud 19
 - isometric 22
 - perspective 22
 - profile 22
- Polynomials
 - RPC 171
- Power spectrum 83
 - compression 96
- Premiere™ 154
- Principal Components 65
 - Apply Transformation 166
 - signatures 138
- Print... 191
 - spectra 119
 - windows 155
- Print Image 155
- Print Window 155
- Profile 80
 - Control-click 80
- profile line
 - point cloud 23
- Projection 138
- Projection matrix 138
- Pseudo color 66

Q

- QuickTime™ 154
- Quit 191

R

- Radiance 182,184
- Radius 91
- Raw 35,153,188
- Recompilation 175
- Redo stats 124
- Ref 150
- Reference coords 42
- Reformatting 156
 - Multiband 156
 - NITF 156
 - Raw 156
- Reformat Cube File... 154,156,160,198
- Remapping 51
- Remap bands 51
- Remove
 - attached data 71
- Replace Color 67
- Residuals
 - example 102
 - warp 102

- Riegl
 - full wave form
 - SDF files 27
- Right 59
- ROC 140
- ROC curve 128
- ROC Mask 131
- Rotations 58,195
- Rows 166
- RPC
 - Create 47
 - format 48
 - geographic coords 171
 - height direction 46
 - Lat/Long Grids 46
 - TIFF files 26
 - Wrap JPEG 158
- Rubberband 54

S

- Sample Data Files 6
- Saturation 51
 - color mapping 55
- Save
 - Bands 112
 - images 153
 - Image Arithmetic 191
 - Overviews 41
 - signature 128
- Save As... 191
- SBET
 - file 29
- Scale 38,114
 - gains & offsets 157
 - spectra 119
- Scaling
 - plot 104
 - selection 104
 - wavelength 106
- Scatter 79,197
- Screen capture 189
- sdfifc32.dll 27
- sdfifc64.dll 27
- Search
 - classmap 147
- Search Domain 127,143
- Selection point 116
- Selection rectangle 52,53,55,65,72
 - data mapping 56
 - plots 104
- Sequence 131
- Sequential
 - correlation 74

- Shaded Relief 168
- Shape Files 12
- Sharp <--> Blur 60
- Shift-click 102,114
 - classify neighborhood 126
 - classify selection 126
 - profile 80
- Shifts 59,196
 - All 59
 - down 82
 - Even 59
 - Left 59
 - left 82
 - line 73
 - offsets 73
 - pixel 73
 - Right 59
 - right 82
 - up 82
- Short 153
- Show Band List 49,105,111,198
 - cube construction 112
 - disable all 105
 - enable all 105
- Show Info 198
- Sigma 39
- Signature 127,143
 - height
 - width
- Signed 36,156
- sliders
 - filtering 95
- Slope
 - x 82
 - xy 82
 - y 82
- Sobel
 - 3x3 subset 61
 - 5x5 61
 - xy gradient 61
- Sound 192
- Source color
 - natural 133
 - saturated 133
- Specific Wavelengths 113
- Spectra... 160,197
- Spectral Libraries
 - averaging 119
 - example of 118
 - format of 117
 - plotting 119
- Spectral Plots 114,126
 - examples 115
- Spike suppression 17
- SRF 8,9,153,188
 - Sun Raster Files 9

- Statistics 50,197
- Stats Region 124
- Stereo
 - anaglyph 173
 - compilation 173
 - elevations 173
 - ortho 173
 - recompilation 178
- Stereo glasses 170
- Stereo Triangulation 171
- Sun 9
- Superimpose 110
- Superimposing Signatures 120,143
- Synthetic cube 131,136

T

- TARGA 9
- TEC 117
- Temperature 182
- Template 122
- Text
 - Saving 189
- Thermal 182
- Threshold 128
- TIFF 153,188
 - Overviews 41
- Tile 198
- Toggle button 120
- Top -> Bottom 58,195
- Transformation 103
 - matrix 166
- transparency
 - PNG files 25
- True color 66
 - composite color 113,196
- Type II Error 140
- Type I Error 140

U

- Underflow 51
- Undo 192
- Unsigned 36,156
- Use lowest elevation 17
- Use source spectra color
 - color selection 116
 - image products 135
 - library searches 144
- USGS/EROS
 - Lens distortion 185
- UTM 26,45
 - ASCII LIDAR 167
 - conversion 47
 - easting 44
 - northing 44
 - reference coords 44
 - Stereo Compilation 172

WGS84 45
zone 44

V

Vector 94
Vector angle 122
Vector plot 149
Version 6
Video 59

W

Warp 97,197
 affine 102
 non-linear 102
 orthogonal 102
 residuals 102
Water absorption 74,115
Wavelength
 *.wav 105
 .wvl 49
 bandlist 105
 header 49
 library 161
 Save *.wav 106
 scaling 106
 Wavelength File 49
Wavelet
 pre-filter 132
Window
 SDF files 9
Windows version
 band scrolling 108
 carriage returns 189
 Cascade 190
 Classmap Editor 147
 control-double-click 80
 control key 6
 Copy color 53
 fatal error message 188
 HyperCube version 6
 icons 188
 main menu 190
 Open As Bands 112
 Open File dialog 8
 printing 155
 profile plots 80
 Replace Color 67
 screen capture 189
 screen print 189
 SDF files 27
 selection rectangle 52
 system requirements 6
 TIFF 8
 Tile 190
Window vs File 65,74

Wm-2 μ m-1 182,184
 ISAC 184
Wrap JPEG 158

X

X,Y
 coordinates 17
 LIDAR 17
X,Y,Z
 columns 15

Z

Zeros 189
Zone 44
 negative 44,45
 Southern Hemisphere 44
Zoom 52
Zoom in 195
Zoom out 195