



Combining Rasters



with
TNTmips®

Before Getting Started

This booklet introduces the Raster Combination operations available in TNTmips®. These operations allow you to apply mathematical or logical expressions to combine or transform a set of two or more raster objects, or to derive altered values from a single raster object. Raster combination operations can serve a wide variety of purposes in the preparation, processing, and analysis of raster data.

Prerequisite Skills This booklet assumes that you have completed the exercises in the tutorials *Displaying Geospatial Data* and *TNT Product Concepts*. Those exercises introduce essential skills and basic techniques that are not covered again here. Please consult those booklets for any review you need. You should use the Display Process (Main / Display) to view the input and output objects for these exercises.

Sample Data The exercises presented in this booklet use sample data that is distributed with the TNT products. If you do not have access to a TNT products DVD, you can download the data from MicroImages' web site. In particular, this booklet uses sample files in the COMBRAST data collection.

More Documentation This booklet is intended only as an introduction to combining raster images. Details of the processes described here can be found in a variety of tutorial booklets, Technical Guides, and Quick Guides, which are all available from MicroImages' web site.

TNTmips® Pro and TNTmips Free TNTmips (the Map and Image Processing System) comes in three versions: the professional version of TNTmips (TNTmips Pro), the low-cost TNTmips Basic version, and the TNTmips Free version. All versions run exactly the same code from the TNT products DVD and have nearly the same features. If you did not purchase the professional version (which requires a software license key) or TNTmips Basic, then TNTmips operates in TNTmips Free mode. All the exercises can be completed in TNTmips Free using the sample geodata provided.

Randall B. Smith, Ph.D., 21 April 2011
©MicroImages, Inc., 1997-2011

You can print or read this booklet in color from MicroImages' web site. The web site is also your source for the newest Getting Started booklets on other topics. You can download an installation guide, sample data, and the latest version of TNTmips Free.

<http://www.microimages.com>

Welcome to Combining Rasters

The Raster Combination operations in TNTmips (available on the Image / Combine menu cascade) provide a varied and powerful set of tools for correcting, enhancing, transforming, and analyzing raster data. Most of the operations require a set of raster objects that are **coregistered** (have the same line and column dimensions, geographic extents, and cell size). Each raster object in the set is used as a variable in a mathematical or logical expression, which is applied on a cell-by-cell basis to the input raster values to produce a new set of output rasters. Other operations use mathematical or logical functions to derive new values from a single input raster.

Many of the Raster Combination procedures are designed for processing and enhancing multispectral imagery (scanned color-infrared airphotos, frame-grabbed aerial video, or satellite imagery). The exercises in this booklet are organized into groups that illustrate the steps in several common processing or analysis tasks. These tasks include creating several types of index rasters, repairing a “data dropout” in one band of a raster set, and processing scenes acquired on different dates to correct for differences in illumination conditions (image normalization). By performing these tasks you will become familiar with many of the Raster Combination operations.

The exercises use a set of Landsat Thematic Mapper (TM) scenes of cropland adjacent to the town of Hanford, California. The three scenes, acquired in April, June, and October of 1993, have been georeferenced, resampled to the same projection, and trimmed to the same raster size and geographic extents.



Pages 4-7 introduce the Multiresolution Fusion process for pan-sharpening and the Automatic Combination process. The exercises on pages 8-11 sample several of the Predefined Raster Combination operations, with a focus on computing index values. Page 12 introduces the Principal Components transformation for reducing the dimensionality of raster sets. The Decorrelation Stretch process for enhancing multiraster datasets is discussed on page 13. Pages 14-19 lead you through the task of repairing a data dropout in one raster of a multiband set. Pages 20-22 demonstrate key steps in normalizing multitemporal datasets.



Landsat TM scene of Hanford, CA area for April 27, 1993. TM band 5 (middle infrared) is displayed as red, TM4 (near infrared) as green, and TM3 (red) as blue. Vegetated areas are indicated by green colors and bare soil by pink and magenta.

Multiresolution Fusion (Pan-Sharpening)

STEPS


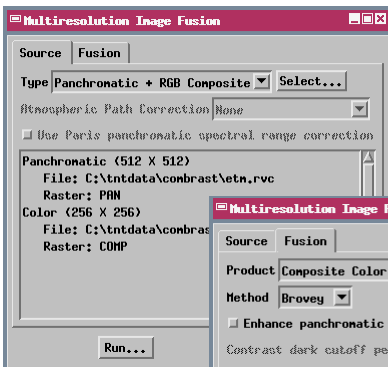
- choose Image / Combine / Multiresolution Fusion from the TNTmips menu
- in the Multiresolution Image Fusion window, choose Panchromatic + RGB Composite from the Type menu
- press[Select]
- in the Select Objects window choose PAN from the ETM Project File for the Panchromatic raster and COMP as the Color raster and press [OK]
- choose Composite Color from the Product menu on the Fusion tabbed panel
- choose Brovey from the Method menu
- turn off the Enhance panchromatic contrast toggle
- press [Run], create a new Project File FUSION, and accept the default output raster name

The Multiresolution Image Fusion process allows you to enhance the spatial resolution of a multispectral image (composite or separate bands) by combining it with a higher-resolution panchromatic (pan) band, a procedure commonly referred to as pan-sharpening. The multispectral and pan images must have matching spatial extents and the cell size of the multispectral image must be an integer multiple of the pan cell size.

The raster objects used in this exercise are extracts of a Landsat 7 image (cell sizes: 30-m multispectral and 15-m pan). The Landsat 7 panchromatic band includes green, red, and near-infrared wavelengths. These three multispectral bands were combined as blue, green, and red color components in the Color Conversion process (Raster / Combine / Convert Color) with contrast-enhancement to form the color-infrared composite image that is used here.



Portion of the original multispectral image (left) and the pan-sharpened result (below).



NOTE: For best fusion results, the composite color image (or set of separate RGB bands) should cover the same spectral range as the panchromatic image band, as in this exercise.

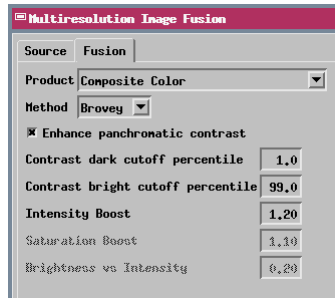
Enhance Pan Contrast in Image Fusion

The fusion process uses raw panchromatic raster values, unmodified by any saved contrast table, to create the pan-sharpened product. However, the *Enhance panchromatic contrast* toggle on the Fusion panel activates settings that can be used to apply mild contrast enhancement to the selected panchromatic raster during image fusion. You can set percentile values for dark and bright cutoffs to apply a linear contrast enhancement. The Intensity Boost setting is a multiplier applied to the panchromatic values; boost settings greater than 1.0 (the default is 1.30) result in brightening across the whole range of panchromatic values. If you cannot achieve the desired brightness and contrast in the pan-sharpened image by these means, you may want to create a contrast table for the pan raster and use the Apply Contrast process (Raster / Filter / Apply Contrast) to create a new pan raster containing contrast-enhanced values for use in Multiresolution Fusion.

The Product menu on the Fusion tabbed panel allows you to choose the form of the pan-sharpened product. When your source type is Panchromatic + RGB (either Composite or Separated), the product choices include Composite Color and either 8-bit or 16-bit RGB separates (set of three grayscale rasters). The separates choices allow you to apply additional contrast enhancement when you create RGB displays of the pan-sharpened separates. If you choose Unspecified for the source type, you can select a pan raster and three or more grayscale bands to produce one sharpened output raster for each input grayscale raster.

STEPS

- turn on the *Enhance panchromatic contrast* toggle on the Fusion panel
- set the Intensity Boost value to 1.20
- keep the default dark and bright cutoff percentile values (1.0 and 99.0, respectively)
- press [Run] and direct the output raster to the FUSION Project File



Part of the pan-sharpened image with enhanced panchromatic contrast



NOTE: when using RGB bands for the color input, raster values are not contrast-adjusted by any saved contrast table. For best color in the Panchromatic + RGB Separates mode, use contrast-enhanced RGB bands produced by the Apply Contrast process.

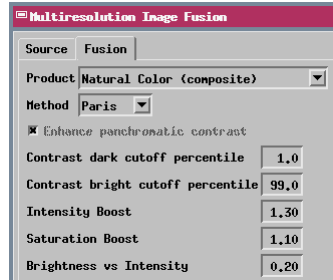

Calibrated Image Fusion

STEPS

- choose IKONOS (Pan + MS) from the Type menu on the Source panel
- press[Select]
- navigate to the IKX Project File and select PAN for Panchromatic, BL for Blue, GL for Green, RL for Red, and NA for Near Infrared
- turn on the toggle *Use Paris panchromatic spectral range correction*
- on the Fusion panel choose Natural Color (composite) from the Product menu
- choose Paris from the Method menu

The QuickBird and IKONOS satellites each provide blue, green, red, and near-infrared multispectral bands and a pan band that spans this entire spectral range with four times the spatial resolution. The Multiresolution Fusion process provides specially-designed pan-sharpening methods for these images. Use the Paris panchromatic spectral range correction (Source panel) and the Paris fusion method (Fusion panel) to produce natural color or color-infrared pan-sharpened composites or RGB raster sets with the proper color balance, brightness, and spectral calibration.

The Paris fusion method automatically incorporates pan contrast enhancement. It also includes the option to increase the color saturation of the product (via the Saturation Boost setting) and to vary the relative contributions of the two color models used to form the color fusion result (via the Brightness vs Intensity setting).



- set the pan contrast settings as shown in the illustration
- press [Run] and direct the result to the FUSION Project File



RGB display of IKONOS RL, GL, and BL bands (natural color).




Natural color result from fusion process.


Automatic Raster Combination

The Automatic Raster Combination process can be used to compare and analyze corresponding cell values in a set of two or more input rasters. The process is best suited for sets of categorical rasters in which the raster values identify spatial classes of some type. The process tabulates all of the combinations of input cell values for corresponding cell locations, then assigns each unique combination to a distinct arbitrary cell value in a new combination raster. An attached CELLVALUES table identifies the actual combinations: it has a record for each cell value in the combination raster and fields containing the corresponding source raster cell values and the count of cells having that combination.

In this example the input rasters show land use for an area for two different years; the values in each raster are 4-digit land use class codes. Each source raster includes 62 classes, and the process finds 154 unique combinations of these classes. The cell counts in the CELLVALUES table can be used to analyze the proportion of the area that experienced overall or particular changes in land use.



Land use raster for 1986




Land use raster for 1995

The Transfer Colors option creates a set of color palettes for the combination raster; each palette reproduces the distribution of class colors for one source raster.

STEPS

- press [Exit] on the Multiresolution Image Fusion window
- select Image / Combine / Automatic from the TNTmips menu
- press [Rasters] on the Automatic Raster Combination window
- select rasters LC1986 and LC1995 from the LCRAST Project File
- press [Run]



VALUE	COUNT	LC1986	LC1995
13	1852	1200	1200
60	44	1200	1700
11	1376	1300	1300
69	19	1300	1700
70	210	1300	1400
74	51	1300	1200
111	39	1300	4410
14	6879	1400	1400
75	116	1400	1700
84	571	1461	1461
35	18	1500	1500
29	109	1600	1600

154 of 154 records show

CELLVALUES table for the combination raster showing combinations of land use classes. Table has been sorted on the LC1986 field and is shown with row colors interleaved by group.

Subtract the Path Radiance Value

STEPS

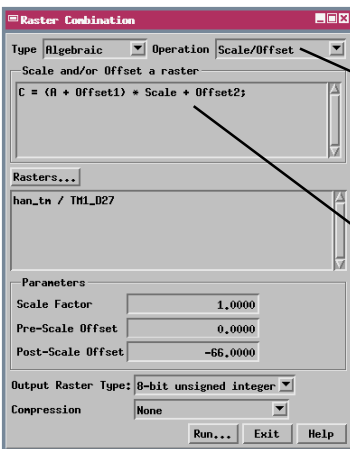
- choose Image / Combine / Predefined from the TNTmips menu
- in the Raster Combination window choose Algebraic from the Type menu
- choose Scale/Offset from the Operation menu
- press [Rasters] and use the standard File / Object Selection procedure to select raster object TM1_D27 from the APR27 folder in the HAN_TM Project File
- set the Post-Scale Offset value to -66 (negative)
- press [Run], create a new Project File COMBRAS, and accept the default output raster name

Path-radiance corrected bands for the Hanford TM scenes can be found in the HAN_PR Project File.

When we study satellite images of the Earth's surface, we often want to use the brightness values recorded by the sensor in various spectral bands to differentiate and identify different kinds of surface materials. However, before you use satellite imagery in other Raster Combination operations, you should use the Scale/Offset operation to subtract a path radiance value from each band.


Path radiance is light scattered back to the satellite sensor by gas molecules and dust in the atmosphere. The scattered light adds to the brightness signal received from the Earth's surface. The amount of scattering (and thus the path radiance value) is greatest for blue light (TM band 1) and decreases with wavelength, becoming almost negligible for the longer middle-infrared wavelengths (TM band 7). Path radiance also varies between scenes taken on different dates. Several techniques for estimating path radiance values are described by Sabins (1997, p. 261) and Jensen (1996, p. 116).

The Scale/Offset operation allows you to multiply raster values by a scaling factor and to add separate pre-scaling and post-scaling offset values. To subtract a path radiance value, keep the Scale Factor set to 1.00 and enter a negative offset value.



Select Scale/Offset from the Operation option button.

This text field shows the mathematical expression used in the operation.



Result of subtracting 66 from each cell in TM1_D27 to correct for path radiance.

Compute a Band Ratio

A simple but very useful raster combination for multispectral imagery is the ratio of one spectral band to another. Use the Algebraic/Divide operation to create simple band-ratio images.


In an individual spectral band, identical surface materials can have different brightness values from place to place because illumination varies with slope angle, slope direction (aspect), and shadowing by adjacent features. Because these illumination effects have about the same magnitude regardless of wavelength, they tend to cancel out when one band is divided by another. A band ratio image therefore emphasizes the inherent properties of the surface materials. In a grayscale display of a ratio image, the darkest and lightest tones identify areas with the greatest difference in reflectance for the two spectral bands. Areas with similar reflectance appear in intermediate gray tones.

Different band ratio combinations emphasize particular types of surface materials. The ratio of near-infrared (TM4) to red (TM3) is widely used as a vegetation index. Green vegetation exhibits a large and unique difference in reflectance in these two bands, appearing bright in the near-infrared and dark in the red band. In a 4/3 ratio image the areas with a full plant canopy appear very bright, and the ratio value decreases as the proportion of plant cover to soil background decreases.


With 8-bit input rasters, the possible range of output ratio values is 1/255 to 255 (excluding 0 input values). Using the commonly available multispectral bands, useful ratio values generally range from 0.4 to 15.0. Use a floating point (decimal value) raster format for the output to preserve the full range of ratio values without requiring rescaling.

STEPS

- in the Raster Combination window select Divide from the Operation option button
- press [Rasters], navigate to the APR27 folder in the HAN_PR Project File, and select object TM4_D27_P for A and TM3_D27_P for B



- choose 32-bit floating-point from the Output Raster Type menu
- press [Run] and direct the output raster to the COMBRAS Project File



TM 4/3 ratio image for April Hanford scene displayed with auto-normalized contrast enhancement. Bright areas show actively growing crops.

Compute a Normalized Difference Index

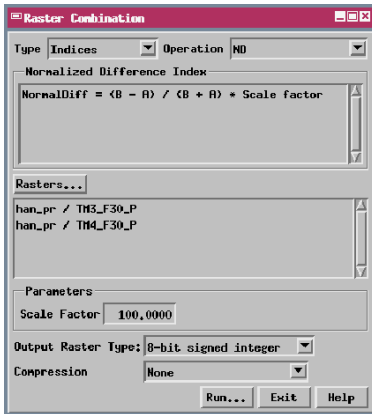
STEPS

- in the Raster Combination window choose Indices from the Type menu; the Operation menu defaults to ND
- press [Rasters], navigate to the JUN30 folder in the HAN_PR Project File, and select raster object TM3_F30_P for A and TM4_F30_P for B
- check that the Scale Factor value is set to 100
- press [Run] and direct the output raster to the COMBRAS Project File

Simple band ratio images have several disadvantages. Sensor noise (which is usually specific to particular bands) is accentuated by the ratio calculation. The ratio result also includes both fractional values (A/B is less than 1 when $A < B$) and values greater than 1 (for $A > B$). If these values are scaled to the 8-bit data range (0 to 255) by multiplying by a constant scale factor, the lower range (for $A < B$) is compressed and the upper range is expanded.


The normalized difference index is a variant form of ratio calculation that alleviates these problems. The difference between corresponding raster values in two spectral bands ($B-A$) is “normalized” by dividing by their sum ($B+A$). As with a simple band ratio, brightness variations due to topography are largely eliminated. Output values can vary from -1 to +1, so the data range is symmetrical around 0 ($B = A$). This range is easily scaled to the 8-bit signed integer data range (-127 to +128). The default scale factor of 100 produces a possible range of output values of -100 to +100.

The Normalized Difference Vegetation Index (NDVI) is a widely used index for green vegetation cover. It assigns the near-infrared channel to B and the



The Indices procedures include a number of generic operations to extract biophysical properties such as vegetation cover and soil brightness from multispectral datasets.

red channel to A. Higher values for NDVI (brighter tones) indicate a greater proportion of ground cover by green vegetation.




TM 4/3 ratio and NDVI images for the Hanford TM scenes can be found in the HAN_INDXX Project File.

Hanford NDVI image for June 30, 1993. Bright tones indicate fields with actively growing crops.

Compute TM Tasseled Cap Indices

NDVI is a derived value that expresses the variation in a real biophysical property of the surface: the degree of cover by green vegetation. Another approach to deriving index values is to project the spectral values from the full set of available bands to a new set of orthogonal coordinate axes in spectral space. The new axes are chosen to correspond as closely as possible to important biophysical properties.

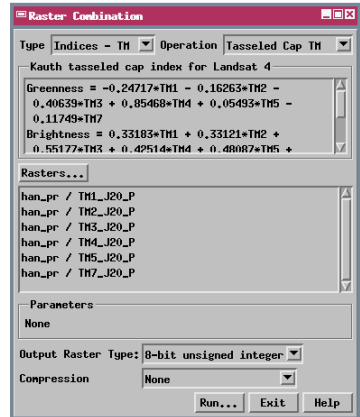
The Tasseled Cap transformation for the six nonthermal Landsat TM bands (1-5 and 7) computes three index values: Greenness (a vegetation index), Brightness (an index of soil brightness), and Wetness (an index of soil moisture content). Most of the variability in soil and vegetation conditions contained in the six TM bands is expressed in these three dimensions. Each of the indices is computed cell by cell as a weighted sum (linear combination) of the input band values. The computation has the form $\text{Index Value} = a\text{TM1} + b\text{TM2} + \dots + f\text{TM7}$. The weighting coefficients a through f are predefined for each index, and are shown in the expression panel of the Raster Combination window.




RGB display of Tasseled Cap index rasters for October Hanford scene with Brightness = Red, Greenness = Green, and Wetness = Blue. Index rasters for all dates can be found in the HAN_INDXX Project File.

STEPS

- in the Raster Combination window choose Indices - TM from the Type menu; the Operation menu defaults to Tasseled Cap TM
- press [Rasters], navigate to the oct20 folder in the HAN_PR Project File, and select object TM1_J20_P for TM1, TM2_J20_P for TM2, and so on through TM7
- press [Run] and direct the output Greenness, Brightness, and Wetness rasters to the COMBRAST Project File



The Tasseled Cap transformation is named for the shape of the point distribution in a plot of Greenness versus Brightness for typical vegetated regions.


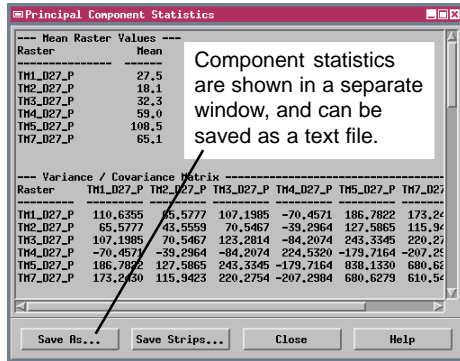



Compute Principal Components

STEPS

- ☑ press [Exit] on the Raster Combination window
- ☑ choose Image / Combine / Principal Components from the TNTmips menu
- ☑ press [Rasters] in the Principal Components window, navigate to the APR27 folder in the HAN_PC Project File, and select objects TM1_D27_P through TM7_D27_P
- ☑ press [Run] and direct the six output principal component rasters to the COMBRAS Project File
- ☑ press [Exit] to close the Principal Components window

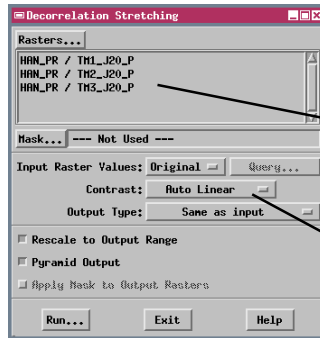
The Principal Components process is another means of transforming a set of spectral bands to derive a reduced number of parameters (**dimensional reduction**) and increase interpretability. Like the Tasseled Cap transformation, the Principal Components process maps each set of input raster values to a new set of orthogonal coordinate axes. However, the coefficients used in the Principal Components transformation are derived from a statistical analysis of the input rasters. The first principal component is the direction of greatest spread (variance) in the data. The second component is the direction perpendicular to the first with the next largest variance. The other components are determined by the requirement of mutually perpendicular axes. This procedure produces a set of uncorrelated output rasters.



First principal component image from the April Hanford TM scene (see the HAN_PC Project File). For typical TM scenes, over 95 percent of the spectral variability in the six nonthermal bands is expressed in the first three principal components. These derived rasters can be interpreted directly or used as input for other processes, such as Automatic Classification, in which the reduction in the number of input rasters greatly speeds processing. Dimensional reduction is possible with common multispectral datasets because certain groups of spectral bands (visible light bands, for example) show a high correlation in their responses to similar surface materials.

Apply Decorrelation Stretch

Decorrelation stretching is a process that enhances the color display of highly correlated raster sets, such as the first three Landsat Thematic Mapper bands. The process performs a principal components transformation on the set of input bands, applies a contrast stretch to the components, then reverses the transformation. When the output rasters are displayed in RGB, hue and intensity are usually similar to the original image, but the color saturation is greatly increased. This enhancement exaggerates the differences in spectral properties between surface materials to a greater degree than is possible using conventional contrast enhancement of the original bands. As a result, you can more easily discriminate subtle variations in surface materials using the decorrelated raster set.



STEPS

- choose Image / Combine / Decorrelation from the TNTmips menu
- press [Rasters] in the Decorrelation Stretching window, navigate to the OCT20 folder in the HAN_PR Project File, and select objects TM1_J20_P through TM3_J20_P
- press [Run] and direct the three output rasters to the COMBRAS Project File
- press [Exit] to close the Decorrelation Stretching window

You can use more than three input bands with the decorrelation process, but it provides the greatest enhancement for highly correlated bands.

You can choose from several contrast stretching methods on the Contrast option button.



Hanford bands TM3, TM2, and TM1 for the October scene displayed as Red, Green, and Blue (respectively), providing a “natural color” view of the scene.




Same band combination after decorrelation stretch. The increased saturation helps distinguish different crops (greens) and different soil types or conditions (shades of orange, red, brown, and lavender).

Threshold a Raster to Create a Mask

STEPS

- ☑ choose Image / Combine / Predefined from the TNTmips main menu
- ☑ in the Raster Combination window choose Logical from the Type menu; the Operation menu defaults to Threshold
- ☑ press [Rasters] and select raster object TM2 from the HAN_BD Project File
- ☑ change the Threshold Value to 2
- ☑ choose Above Threshold from the Output White menu
- ☑ press [Run], create a new Project File HAN_BD1, and change the output raster name to BD_MASK
- ☑ [Exit] the Raster Combination process

Raster object TM2 with area of bad data values (0 and 1, displayed as black).



Change the Threshold Value to 2.0


Select Above Threshold from the Output White option button to assign an output value of 1 for input cell values greater than 2.

This is the first of six exercises that lead you through the task of “repairing” an area of “bad data” in one raster of a multiband set. A sensor or data transmission malfunction can result in a loss of data for part of an image band. The result is one or more lines in the raster that contain spurious values (sometimes called a data dropout). This condition has been simulated for these exercises by editing band TM2 in the June Hanford TM raster set.

The first step is to use the Logical Threshold operation to create a binary raster to mask the dropout area in later processes. The mask should have a raster value of 0 for each cell in the dropout and a value of 1 for all other cells. The Threshold process uses a single input raster value as a boundary to determine the binary value of the output cells. You can choose to output a value of 1 (white) for all input cell values below the threshold, or the inverse. In this case there is a numeric break between the dropout values (0 and 1) and the smallest real data value (23), so a threshold value of 2 provides the necessary separation.

Select Logical from the Type menu.

Threshold is the default Logical operation.



Perform Multilinear Regression

Now use the Multilinear Regression process to determine the mathematical relationships between the TM rasters in the HAN_BD Project File. In the next exercise we will use these relationships to predict values for the dropout area in band TM2.

The Multilinear Regression process calculates a series of linear regression equations for a set of rasters. The analysis treats each input band in turn as the dependent variable, and the remaining bands as independent variables. Each of the resulting equations identifies the linear combination of remaining bands that best reproduces the raster values in the dependent band. The best fit is determined by minimizing the sum of the squared deviations between the predicted values and the real values.

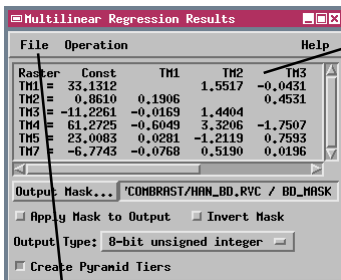
The regression relationships should be based only on the real data in each of the rasters. The spurious values in TM2 should not influence the calculations. Therefore we use the mask produced in the previous exercise to eliminate the dropout area from the calculations. If you were running the process on a larger set of input rasters, you could speed processing by using a subset of the cells to determine the regression equations. This subset is defined by setting a sampling interval in the line and column directions.

Output rasters for each of the exercises in this sequence can be found in the HAN_BD Project File. You can use these objects as input for the next step, or use the ones you place in your HAN_BD1 output file.

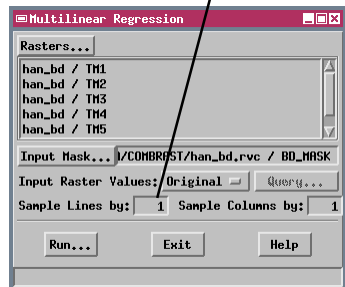
STEPS

- choose Image / Combine / Multilinear Regression from the TNTmips menu
- press the Rasters button in the Multilinear Regression window and select raster objects TM1 through TM7 from the HAN_BD Project File
- press [Input Mask] and select the BD_MASK object from the HAN_BD Project File
- press [Run]

You can set line and column sampling intervals by entering values in these text boxes.



Calculated values are shown in the Multilinear Regression Results window.




Keep the Multilinear Regression and Multilinear Regression Results windows open for use in the next exercise.

Options on the File menu allow you to save the regression results as an array file (for use in other processes) or as a text file.


Predict Values from Regression Results

STEPS


- in the Multilinear Regression Results window, turn on the Apply Mask to Output and the Invert Mask toggle buttons
- choose Predict from the Operation menu
- in the File / Object Selection window, highlight TM2 in the Selections list
- navigate to the HAN_BD1 Project File and press the New Object icon button 
- accept the default object name (P_TM2) in the New Object window and press [OK]
- press [OK] on the File / Object Selection window to begin processing
- when the process finishes, choose Close from the File menu on the Multilinear Regression Results window
- click [No] on the Verify window that asks if you want to save the results
- click [Exit] on the Multilinear Regression window

Once the regression relationships have been computed, we can use the Predict operation to calculate estimated values for any band from the values in the remaining rasters in the set. We need predicted values only for TM2, so we assign an output raster only for this band; no other predicted rasters are computed.

Later we will want to merge the predicted TM2 values for the dropout area with the real TM2 values for the remainder of the image. The steps leading to this merger are easier if predicted values are assigned only to the dropout area, and another known value to the remaining portion of the predicted TM2 raster. To do this we have the Predict operation apply the bad data mask to the output raster after predicted values are calculated. Since the mask contains 0 values for cells in the dropout area and values of 1 for remaining cells, we need to invert the mask for this operation. Then the predicted values for the dropout area are multiplied by 1, and passed unchanged to the output raster, while other cells are assigned the default null value (255).



Output P_TM2 raster with predicted raster values for the dropout area. Masked areas are assigned a value of 255, but are shown in cyan in this illustration.



Use Replace All Operation

In order to merge predicted and actual rasters most simply, the cells for which we have real data in raster TM2 should have a value of 0 in the predicted raster. Unfortunately, the Predict operation assigned a value of 255 to these cells. We can use the Replace All operation to remedy this situation.



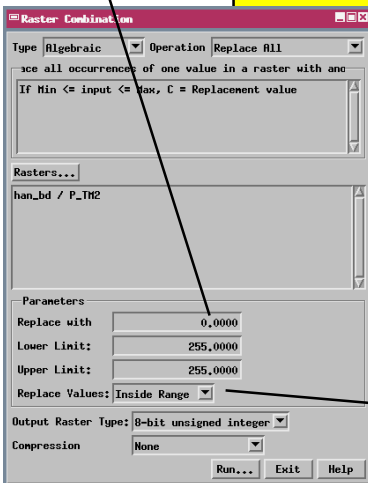
The Replace All operation replaces a single input value or range of values with another designated value. In this case we replace the value 255 in the input raster with 0. However, 255 was designated as the null value in the raster produced by the Predict operation in the previous exercise. Before cells with this value can be processed in this exercise, it is necessary to use the Project File Maintenance process (Tools / Manage Project Files) to edit the object information for the predicted TM2 raster and turn off the Has Null Value toggle button. (This has already been done for object P_TM2 in the HAN_BD Project File.)

STEPS

- choose Image / Combine / Predefined from the TNTmips menu
- choose Algebraic from the Type menu
- choose Replace All from the Operation menu
- press [Rasters] and select object P_TM2 from the HAN_BD Project File
- change the Lower Limit value and the Upper Limit value to 255
- press [Run], open the HAN_BD Project File, and change the output raster name to P_TM2_0

Keep the Raster Combination window open for the next exercise.

Keep the default value of 0.00 for the Replace With setting.



Output raster P_TM2_0, in which the raster value 255 in the input has been replaced by the value 0 (shown in yellow in this illustration).

The default Inside Range option replaces all input values within the range of the Lower Limit and Upper Limit. In this case both limits are identical, and only the value 255 is replaced.

Multiply Rasters to Apply a Mask

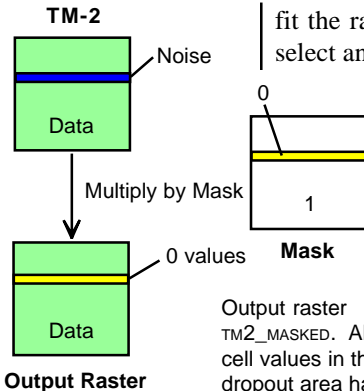
STEPS

- ☑ in the Raster Combination window, choose Multiply from the Operation menu
- ☑ press [Rasters] and select objects TM2 and BD_MASK from the HAN_BD Project File
- ☑ press [Run], open the HAN_BD1 Project File, and name the output raster TM2_MASKED


In previous exercises we used a binary mask raster (BD_MASK) to control the input and output of values in the Multilinear Regression process. Before we merge the predicted raster values from the previous exercise with the real data values in raster TM2, we need to make a masked version of the TM2 raster in which all values in the dropout area are set to 0. (Recall that the dropout area includes values of 0 and 1). We can accomplish this by using the Multiply operation to multiply each value in TM2 by the

corresponding value in BD_MASK. Since the mask has a value of 0 for all cells in the dropout area, the multiplication will produce 0 values for the entire dropout area. Cells outside the dropout area have a mask value of 1, so the multiplication passes the real data values to the output masked raster with no change.

You can multiply any types of grayscale raster object in the Multiply operation. The operation also includes a Scale Factor with a default value of 1.00. If you multiply two or more nonbinary rasters, you may need to adjust the Scale Factor to rescale the data to fit the range of the selected output raster type, or select an output raster type with higher bit-depth.



Output raster TM2_MASKED. All cell values in the dropout area have a value of 0.



Keep the Raster Combination window open for the next exercise.



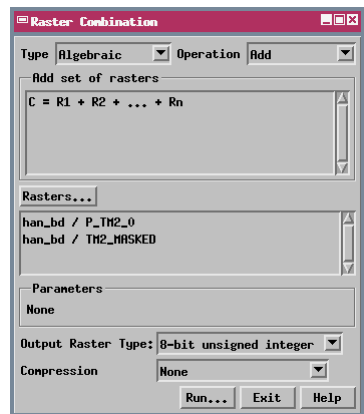
Add Rasters

We can now use a final algebraic operation to merge the masked TM2 raster (`TM2_MASKED`) with the raster that contains predicted values for the dropout area (`P_TM2_0`). The dropout area has 0 values in `TM2_MASKED` and predicted values in `P_TM2_0`. Conversely, the remaining area has real data values in `TM2_MASKED` and 0 values in `P_TM2_0`. So a simple cell-by-cell addition of raster values will merge the predicted values for the dropout area with the real data without changing any of the raster values.

The Add operation performs this cell-by-cell addition. You can input any number and type of raster objects. There is no option to automatically scale the result, however, so in some cases you may need to choose an output raster type with greater bit-depth than the input objects to ensure that output values do not exceed the range of the output raster type. (For an 8-bit output raster, any calculated value greater than 255 is assigned a value of 255).

STEPS

- ☑ choose Add from the Operation menu in the Raster Combination window
- ☑ press [Rasters] and select objects `P_TM2_0` and `TM2_MASKED` from the `HAN_BD` Project File
- ☑ press [Run], open the `HAN_BD1` Project File, and change the output raster name to `TM2_FIXED`



Raster object `TM2_FIXED`, with predicted values for the dropout area merged with real values for the remaining portion of the raster.

Keep the Raster Combination window open for the next exercise.

Use Logical Range to Create a Mask

STEPS

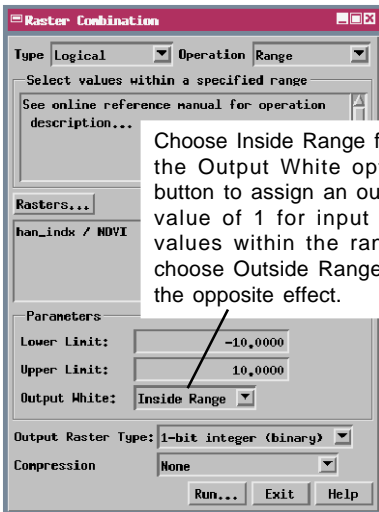
- ☑ choose Logical from the Type menu in the Raster Combination window
- ☑ choose Range from the Operation menu
- ☑ press [Rasters], navigate to the APR27 folder in the HAN_INDXX Project File, and select object NDVI
- ☑ set the Lower Limit parameter value to -10, and the Upper Limit value to 10
- ☑ press [Run] and direct the output raster to the COMBRASST Project File

NOTE: Determine limit values in the Display process by comparing the NDVI raster to an RGB display of TM 5-4-3, in which vegetated areas appear green and wet areas appear dark blue. Some trial-and-error testing may be required to find the best values.


Satellite scenes of the same area acquired on different dates may have different brightness values for identical surface materials because of varying sensor and illumination conditions. These variations can complicate analysis and interpretation of a time-series of images. The next three exercises present key steps in adjusting (normalizing) the April and October Hanford TM scenes to match the conditions of the June scene.

First we use the Logical Range operation to create a mask from the April NDVI raster. The Range operation creates a binary raster with values determined by a continuous range of values in the input raster. The input range is defined by Lower Limit and Upper Limit parameter values. In this exercise the mask is designed to mask out cells whose brightness is likely to change through time: vegetated areas (high NDVI) and areas of water or wet soil (low NDVI). The Lower Limit value is chosen to separate wet from dry conditions, and the Upper Limit value to separate nonvegetated and vegetated cells. The operation assigns an output value of 1 for each input cell with an NDVI value inside the

specified range.



Choose Inside Range from the Output White option button to assign an output value of 1 for input cell values within the range; choose Outside Range for the opposite effect.



Mask raster produced by the Range operation displayed over the grayscale April NDVI raster. Zero values in the mask are displayed in green, identifying cells with brightness influenced by vegetation or wet conditions.

Merge Masks using Logical AND

The mask created for the April scene in the previous exercise has a value of 1 for cells in nonvegetated, nonwet areas. We need to identify areas that remain in this state for all three image dates. Such areas (called **pseudoinvariant** features) have about the same spectral properties in each scene and can therefore be used to determine the brightness adjustments required to normalize the three scenes. Further information on scene normalization using pseudoinvariant features can be found in Jensen (1996, p. 116-121) and in Schott and others (1988).

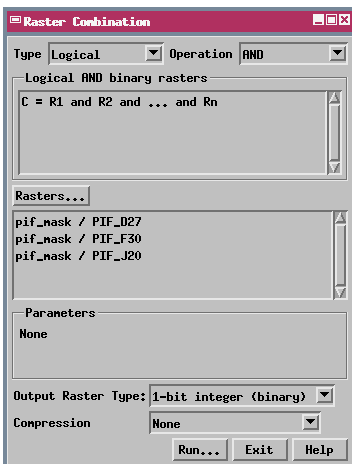
We can use the Logical AND operation to create a pseudoinvariant features mask (PIF mask) from NDVI masks created for each date. The AND operation creates a binary raster from two or more input binary rasters. An output cell is assigned a value of 1 only if the corresponding cell value in *each* input raster has a value of 1. The resulting PIF mask therefore has a value of 1 only for cells that were nonvegetated and nonwet on all three dates. The PIF mask can be applied to each TM band using the Multiply operation (as in a previous exercise) to create a pseudoinvariant features raster set for each date.

STEPS


- choose AND from the Operation menu
- press [Rasters] and select objects PIF_D27, PIF_F30, and PIF_J20 from the PIF_MASK Project File
- press [Run] and direct the output raster to the COMBRAST Project File

The Logical OR operation also creates a binary raster from two or more input binary rasters. An output cell is assigned a value of 1 if the corresponding value in *any* of the input rasters is 1.

The output pseudoinvariant features mask is included as object PIF_MASK in the PIF_MASK Project File. Pseudoinvariant features raster sets for each date can be found in the HAN_PIF Project File.



Keep the Raster Combination window open for the next exercise.



PIF mask created by applying the Logical AND operation to the three input masks. Values of 0 are shown in yellow, and 1 in blue. Only a small proportion of the cells represent pseudoinvariant features.


Apply Scale and Offset

STEPS

- choose Algebraic from the Type menu in the Raster Combination window
- choose Scale/Offset from the Operation menu
- press [Rasters], navigate to the APR27 folder in the HAN_PR Project File, and select object TM1_d27_P
- set the Scale Factor parameter value to 1.279
- set the Post-Scale Offset value to 0.08
- press [Run] and direct the output raster to the COMBRAS Project File

Before performing the linear regression, the Has Null Value option was turned off for rasters in the HAN_PIF Project File. This includes the 0 values in the calculation; these values should be equivalent on different dates for bands with path radiance correction applied.

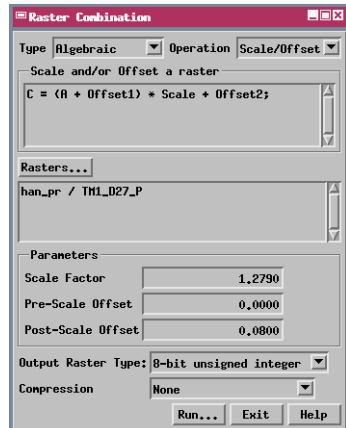
The pseudoinvariant features in each TM band for each date exhibit a range of brightness values. If we compare the same band on two dates (TM1 for April and June, for example), there is a strong linear relationship between the sets of brightness values. By using the pair of PIF-masked rasters in the Multilinear Regression process (discussed on page 15), we can derive linear regression equations describing this relationship. The equation that expresses June values as a function of April values (displayed in the Multilinear Regression Results window) yields a constant and a coefficient that can be used to adjust the April TM1 values to match the June scene conditions. (Each TM band pair yields a unique constant and coefficient).



We use the Scale/Offset operation to rescale the April TM1 raster (path-radiance corrected) to match the June scene. The coefficient from the June-April regression equation is used as the Scale Factor, and the constant as the Post-Scale Offset.



Landsat TM scene for April normalized to June illumination conditions. RGB band assignments are TM5, TM4, and TM3, respectively.



Normalized versions of the Hanford TM scenes can be found in the HAN_NORM Project File.

What Next?

This booklet has provided a brief introduction to Raster Combination operations in TNTmips. Although the exercises in this booklet focus on sample applications of Raster Combination operations to satellite imagery, there are many additional applications. These operations can be applied to any type of spatially-varying numerical data stored in raster form. For example, raster data on crop type, soil conditions, and past crop yield could be used to calculate required fertilizer application rates on a cell-by-cell basis for precision crop management. Raster values of soil erodibility, slope conditions, rainfall, and vegetative cover could be used as input to the Universal Soil Loss equation to calculate soil erosion rates. Raster combination operations can thus be very useful in spatial modeling and raster-based GIS analysis. You can also perform similar computations using rasters and attributed vector objects in the GeoFormula process, which is described in the tutorial booklet *Using Geospatial Formulas*.

References

The following references are good places to start if you want additional information on raster combination operations and their applications to remote sensing and spatial modeling:

- Bonham-Carter, Graeme F. (1994). *Geographic Information Systems for Geoscientists: Modelling with GIS*. New York: Pergamon. 398 pp.
- Burrough, P. A. (1986). *Principles of Geographical Information Systems for Land Resources Assessment*. Chapter 5, *Methods of Data Analysis and Spatial Modelling*. Oxford: Clarendon Press. pp. 81-102.
- Jensen, John R. (1996). *Introductory Digital Image Processing: a Remote Sensing Perspective* (2nd ed.). Chapter 6, *Image Preprocessing: Radiometric and Geometric Correction*, and Chapter 7, *Image Enhancement*. New York: John Wiley and Sons. 316 pp.
- Sabins, Floyd F. (1997). *Remote Sensing: Principles and Interpretation* (3rd ed.). Chapter 8, *Digital Image Processing*. New York: W. H. Freeman. 494 pp.
- Schott, John R., Salvaggio, Carl, and Volchok, William J. (1988). *Radiometric Scene Normalization Using Pseudoinvariant Features*. *Remote Sensing of Environment*, 26, 1-16.

Advanced Software for Geospatial Analysis

MicroImages, Inc. publishes a complete line of professional software for advanced geospatial data visualization, analysis, and publishing. Contact us or visit our web site for detailed product information.

TNTmips Pro TNTmips Pro is a professional system for fully integrated GIS, image analysis, CAD, TIN, desktop cartography, and geospatial database management.

TNTmips Basic TNTmips Basic is a low-cost version of TNTmips for small projects.

TNTmips Free TNTmips Free is a free version of TNTmips for students and professionals with small projects. You can download TNTmips Free from MicroImages' web site.

TNTedit TNTedit provides interactive tools to create, georeference, and edit vector, image, CAD, TIN, and relational database project materials in a wide variety of formats.

TNTview TNTview has the same powerful display features as TNTmips and is perfect for those who do not need the technical processing and preparation features of TNTmips.

TNTatlas TNTatlas lets you publish and distribute your spatial project materials on CD or DVD at low cost. TNTatlas CDs/DVDs can be used on any popular computing platform.

Index

AND (logical operation).....	21	normalization of multirate images.....	3,20-22
Add (algebraic operation).....	19	normalized difference index.....	10
automatic combination.....	7	pan-sharpening.....	4-6
band ratio.....	9,10	path radiance.....	8
brightness index.....	11	pseudoinvariant features.....	21
contrast stretch.....	13	principal components transformation.....	12
decorrelation stretch.....	13	range (logical operation).....	20
dimensional reduction.....	12	ratio index.....	9,10
divide (algebraic operation).....	9	replace all (logical operation).....	17
greenness index.....	11	saturation.....	13
mask raster.....	14,16,18,20	scale/offset (algebraic operation).....	8,22
multilinear regression.....	15,16	tasseled cap indices.....	11
multiply (algebraic operation).....	18	threshold (logical operation).....	14
multiresolution fusion.....	4-6	vegetation index.....	10,11
NDVI.....	10,11,20,21	wetness index.....	11



MicroImages, Inc.

www.microimages.com