

Algorithm Theoretical Basis Document

for

Decorrelation Stretch

Version 2.2 August 15, 1996

Ronald E. Alley
Jet Propulsion Laboratory
4800 Oak Grove Drive
Pasadena, CA 91109

The decorrelation stretch is a process that is used to enhance (stretch) the color differences found in a color image. The method used to do this includes the removal of the inter-channel correlation found in the input pixels; hence, the term "decorrelation stretch".

The purpose of this document is to explain:

- 1) the conditions that normally appear in multispectral data that indicate that color enhancement is needed,
- 2) how a decorrelation stretch addresses those needs,
- 3) how a decorrelation stretch works, i.e., computationally, what steps are performed,
- and 4) what the limitations are to this approach.

Introduction

The Advanced Spaceborne Thermal Emission and Reflectance Radiometer (ASTER) is a high-spatial-resolution multispectral imaging device scheduled to fly in Earth orbit starting in mid-1998, on the first platform of NASA's Earth Observing System (EOS-AM1). The instrument will have three bands in the visible and near infrared (VNIR) spectral range (0.5 to 1.0 μm) with 15 meter spatial resolution, six bands in the short-wave infrared (SWIR) spectral range (1.0 to 2.5 μm) with 30 meter spatial resolution, and five bands in the thermal infrared (TIR) spectral range (8 to 12 μm), with 90 meter spatial resolution (Kahle, et al., 1991; Yamaguchi, et al., 1993). An additional backward viewing telescope with a single band in the near infrared with 15 meter spatial resolution will provide the capability, when combined with the nadir viewing elements, for same-orbit stereo data. The instrument is being provided by the Japanese Government under the Ministry of Trade and Industry (MITI). The ASTER project is implemented through the Earth Resources Satellite Data Analysis Center (ERSDAC) and the Japan Resources Observation System Organization (JAROS), which are nonprofit organizations under the control of MITI. JAROS is responsible for the design and development of the ASTER instrument, which will be carried out by the Nippon Electric Company (NEC), the Mitsubishi Electric Corporation (MELCO), Fujitsu, and Hitachi under contracts with JAROS. The ASTER science team is an international team of Japanese, American, French, and Australian scientists. The team participates in the definition of the scientific requirements for ASTER, in the development of algorithms for

data reduction and analysis, and in calibration, validation, and mission planning.

The decorrelation stretch algorithm that is described by this document, produces a color image, suitably color contrast enhanced and uncorrelated in the output channels.

Overview

Experimental Objective

The decorrelation stretch is proposed as a processing technique to be used for browse data products at reduced resolution, and also as a standard data product for image analysis. It is anticipated that the requesters of the standard data product would be using this product primarily as a data visualization tool, since the primary purpose of the technique is to enhance the visual interpretability of the image data.

There are several qualities that are desirable for a product that is to be used as a visual tool. First and foremost, the product should show as much discrimination of scene variations as is practical. That is, if the data supports differentiation within the scene of different units (of vegetation types, geologic units, land uses, etc.), then that distinction should be apparent in the image. Second, it is desirable for colors and brightnesses to remain consistent among images acquired at different times or at different locations. That is, it is desirable for a unit to be identifiable by its color signature. A third desirable trait in a product used as a visual tool is for consistency with the other tools and knowledge available to the investigator/interpreter. When possible, the product should resemble previously available satellite image products, aerial photographs, maps, or scale models. Finally (and related to the previous point), a standard data product should adhere to the conventions used by other data products, to the extent possible.

It should be noted that each of these desired qualities is in conflict with each of the others, to a certain extent. As a result, compromises are needed to resolve competing goals. What constitutes an ideal image product is a function of the relative weight given to each of the desired qualities, and, as such, often a matter of personal preference.

The method chosen here, the decorrelation stretch, places its greatest emphasis upon a high degree of scene variation. The method used to generate the decorrelation stretch, however, can produce a family of transformations, each with an equal degree (to the extent this can be measured statistically) of scene variation. Among this family of transformations, this algorithm chooses the one that is expected to contain the least difference from the input image. This has the effect of both minimizing scene dependent effects (for the sake of consistency over place and time) and making the color-sense of the image to be consistent

with image products that have not been processed with a color transformation.

In order to display multispectral image data that lies outside the range of human vision, an artificial assignment must be made for representing the data in color. This is normally done by displaying three selected channels as the red, green, and blue components of an additive color picture. The convention used by this algorithm selects as its default color assignment the channel with the shortest wavelength to be the blue component, and the longest wavelength channel to be the red component. This is consistent with the visible spectrum, where the order is blue-green-red (short wavelength to long), and is also the most commonly found ordering in existing image products.

Historical Perspective

The decorrelation stretch originated as an adaptation and extension of two closely related data transformation techniques: the principal component transformation, and the Karhunen-Loeve transformation. The Karhunen-Loeve transformation (Loeve, 1955) is a linear transformation ("rotation") in multi-dimensional space, with the transformation vectors being defined as the eigenvectors of the covariance matrix of the original data. The principal component transformation is similar, except that the transformation vectors are derived from the correlation matrix rather than the covariance matrix. When the variances of the individual input variables are the same, the results of these two transformations are identical. While the definitions of these two transformations, as presented above, are distinct, the term "principal component image" is frequently used in image processing literature to describe the results of either of these transformations.

The principal component transformation provides a straightforward method of removing the correlation from multispectral image data, and has the added benefit of providing a method of reducing the dimensionality of multispectral images. Its primary drawback arises from the fact that its color assignments are not readily associated to any physical quantity, and may be dramatically scene and processing dependent.

In 1973, M. M. Taylor pointed out that once the multispectral data have been decorrelated by a coordinate system rotation, and the variances equalized in this new coordinate space, any additional rotation of the new coordinate space would not re-introduce correlation. He went on to compare several alternative additional rotations for a return to a more interpretable coordinate system.

In 1978, Soha and Schwartz proposed that for most remote sensing applications, the simple inverse rotation to the original color space was the rotation most suitable for image interpretation. It is this technique that has been given the name "decorrelation stretch", despite

the fact that all of the above techniques involve contrast stretching and yield uncorrelated outputs.

During the past decade, the decorrelation stretch has found increased usage, primarily for the newer multispectral imaging systems that have closely spaced channels in the spectral domain, and hence, extremely high interchannel correlation. This interchannel correlation is frequently observed to be larger than 0.95 for Thermal Infrared Multispectral Scanner (TIMS) data in channels comparable to the ASTER thermal channels, and in the Airborne Visible InfraRed Imaging Spectrometer (AVIRIS) channels comparable to the ASTER SWIR channels. The correlation coefficients among the ASTER VNIR channels are expected to be somewhat lower than SWIR and TIR channels on average, but even here values greater than 0.9 should not be rare.

The decorrelation stretch is presently the standard hardcopy product for all TIMS data acquired by Ames Research Center. It is also the standard TIMS browse product, available via the EROS Data Center (EDC).

Instrument Characteristics

The three multispectral telescope subsystems (VNIR, SWIR, and TIR) each produce images with different pixel spacings (15m, 30m, and 90m). In a typical scene, the correlation between channels of different subsystems can be expected to be much lower than between channels within a single subsystem. Applying the decorrelation stretch to input images of differing spatial resolutions will tend to produce distracting processing artifacts in the output image, and the benefits of a decorrelation stretch are reduced when the input channels are less highly correlated. As a result, the decorrelation stretch product is appropriate only for three input channels originating from a single telescope system.

Browse product images will be generated from a standard set of VNIR, SWIR and TIR channels. These channels are anticipated to be Bands 1, 2, and 3 for VNIR, Bands 5, 7, and 9 for SWIR, and Bands 10, 12, and 14 for TIR. An alternate set of channels may replace this selection, if warranted by the relative quality of the individual channels. An investigator may request a decorrelation stretch standard product from any three channels that have the same pixel spacing, but cannot receive a decorrelation stretch from channels that have different pixel spacings (i.e., resolution).

Algorithm Description

Mathematical Objective

The decorrelation stretch program provides a technique to enhance the color separation in images with high interchannel correlation. If one views the pixels from 3 channels of an ASTER scene as 3-vectors, this is done by first finding the linear transformation that results in removing the correlation among the vectors in the transformed space. This is an eigenvector problem, and can be thought of as a rotation of the coordinate system of the original vector space. Within this rotated space, each component is rescaled (contrast stretched) by Normalizing the variances of the vectors. Then the rotation that returns the vectors to the original coordinate system is applied. In practice, both of the rotations and the variance Normalization step can be described by matrix and vector operations, and can be combined into a single mathematical operation that operates on the input image and produces the decorrelation stretched output.

The net effect of the process is to obtain an output image whose pixels are well distributed among all possible colors, while preserving the relative sense of hue, saturation, and intensity of the input image.

Description of Algorithm

First, a grid of pixels is selected that is expected to form a representative sample of all "good" surface pixels found within the scene. Any pixel that had been marked as "bad data" by prior processing is excluded from the grid. Similarly, any pixel where the surface is hidden by clouds (as determined by a separate cloud detection algorithm) is also removed from the grid. This grid is initially formed by subsampling the image at regular intervals. If subsampling results in an inadequate number of usable pixels, all "good" pixels are included in the grid. At the present time, the initial subsampling is programmed to be every third sample of every third line, and the threshold for an adequate number of pixels is set to 1,000. These values may be changed in later versions, if the constraints of quality or speed require it.

For user requested standard (not browse) products, the requester may optionally specify a rectangular subarea for statistics gathering, rather than use the entire image. This option would permit a greater diversity of color within the region of interest, at the expense of overall scene diversity.

Using the sampled pixels, the nine sums that are needed to calculate the covariance matrix for the three channels are accumulated. These sums are:

For $l=1,3$; $m=1,l$, and sampling n pixels ,

$$SUMX_{l,m} = \sum_{k=1}^n P_{k,l} * P_{k,m}$$

$$SUM_l = \sum_{k=1}^n P_{k,l}$$

where $P_{k,l}$ is the value of the k th pixel for Channel l

The covariance and the correlation matrices are computed, using the following formulas:

The elements of the covariance matrix are computed as follows:

$$Cov_{l,m} = \frac{1}{n-1} \left[SUMX_{l,m} - \frac{1}{n} * SUM_l * SUM_m \right]$$

For the correlation matrix elements:

$$Corr_{l,m} = \frac{Cov_{l,m}}{(Cov_{l,l} * Cov_{m,m})^{1/2}}$$

The eigenvectors and eigenvalues of the system described by the correlation matrix (or, optionally, of the covariance matrix) are computed. The matrix of eigenvectors is referred to as the rotation matrix, \mathbf{R} , in subsequent steps.

The "stretching vector" (or Normalization vector), \mathbf{s} , is formed by taking the reciprocal of the square root of each element in the eigenvalue vector, and multiplying it by the desired standard deviation for the output image channels. For true Normalization, the desired standard deviation would be one, but in order to yield output values in the appropriate range for eight bit pixels (i.e., byte data) a higher target value is used. Currently the target standard deviation is set to 50.

The final transformation matrix, \mathbf{T} , is composed from the rotation matrix and the stretching vector. This is done by the following matrix multiplication:

$$\mathbf{T} = \mathbf{R}^t \mathbf{s} \mathbf{R}$$

Prior to doing the transformation upon the image, this transformation is applied to a vector of the means of the input channels. The result is used to compute the offsets needed to reposition the output image values to the 0 to 255 dynamic range of eight bit data. For each pixel in the scene, the output pixel vector (3 valued) is computed by applying the final transformation matrix, and then the offset vector.

Practical Considerations

Computational/Programming/Procedural Considerations

As detailed previously, the decorrelation stretch is conceptually described as four separate operations: statistics gathering, forward rotation, contrast stretching, and back rotation. In practice, the final three steps may be combined into a single step. As a result, there is no need for passing through the data more than twice: once (sampling only a fraction of the image) for statistics gathering, and a second time to apply the transformation to the input pixels to form the output image.

When requesting this data product, the investigator has several decisions to make regarding the processing of the data. First, three channels must be selected for input into the decorrelation stretch. Any three channels are valid, provided that all three come from the same subsystem (VNIR, SWIR, or TIR). Next, the matrix used to compute the eigenvalues and eigenvectors must be chosen. Either the correlation matrix or the covariance matrix may be used. The correlation matrix (the default choice) is preferred if the user wishes to give the three channels equal weight, regardless of the relative scene variance among the channels. The covariance matrix is preferred if the user wishes to weight the input channels proportionally to each channel's scene variance. When the three input channels have equal scene variances, there will be no difference between the outputs from these two options. The target output pixel means and standard deviations (normally set to 127.5 and 50.0) may also be adjusted for a given scene, at the data requester's discretion. Finally, the investigator may choose to select a rectangular subarea of the scene to be used for the gathering of statistics. Normally, the statistics are sampled from the entire scene.

Calibration and Validation

Since the output of the decorrelation stretch does not correspond to data in physical units, no calibration is necessary.

For the same reason, validation in the traditional calibration/validation sense is not needed. Prior to launch, the operational version of the computer program to produce this product will be tested to be certain that its results conform to the algorithm

specifications and are in agreement with alternate implementations of the algorithm. Immediately after launch, a number of the browse product images will be examined to insure that the algorithm is functioning correctly. They will be examined visually for satisfactory color enhancement, and statistically tested to verify uncorrelated output channels with the proper scene means and variances.

Quality Control and Diagnostics

A report file and an error log are generated when the decorrelation stretch product is produced. The possible contents of the error log, which should normally be empty, are discussed in the Exception Handling section. The contents of the report file (to be delivered with the image product and included as part of the image header data) include the number of pixels sampled for the gathering of statistics, the channel means and standard deviations, the channel covariance and correlation matrices, the principal component eigenvectors and eigenvalues, and the coefficients used to generate the decorrelation stretched image.

The information contained in the report file provides several indicators that can be used by the requester to assess the quality of the inputs to the decorrelation stretch, and, to a certain extent, indicate the nature of the transformation and the likely quality of the results. In particular, the off-diagonal correlation and covariance matrix terms and the three eigenvalues will indicate the severity of the color contrast enhancement that has been applied.

Exception Handling

There are two exception conditions in input data that this algorithm checks for and acts upon. If a pixel has been flagged as "bad data" or "cloud" by the ASTER Scene Classification product, that pixel is excluded from the statistics gathering process.

Only one condition terminates execution of the decorrelation stretch program without producing an output image. If fewer than a threshold value (presently 1000) of usable pixels are found when gathering scene statistics, a message is sent to the error log, and the program terminates. A usable pixel is defined as any pixel that has not been classified as "cloud" or "bad data". Under these circumstances, there would be far too little surface data to provide a useful image of the surface, and no image product need be generated.

There are two additional exceptions that the program monitors. Neither of these conditions cause the program to terminate prior to completion. If an input channel variance is zero (that is, all pixels in one channel have the same value), or an eigenvalue is zero (which will occur when the value at each pixel in one channel is a linear combination of the pixel values of the other two channels), the algorithm will be unable to fill the entire color space. In these cases, a warning message is sent to the error log, but the program continues and runs to completion. Both of

these conditions may theoretically occur with valid data, but are far more likely to arise from an incorrect or corrupted input dataset. One such example would be the case where two of the three requested input channels are, in reality, data from the same channel.

Constraints, Limitations, and Assumptions

The decorrelation stretch algorithm is best suited to the case where the input data of all three channels have a joint distribution that is Gaussian (or near Gaussian) in form. Fortunately, the algorithm is fairly insensitive even to substantial deviations from the ideal. One should be aware, though, that if the distribution of the input pixels is strongly bimodal (or multimodal), the effectiveness of the decorrelation stretch is weakened, and there will be less diversity of color in this image than in other images. This can most easily occur when a substantial fraction of the scene consists of material that is unusually and very uniformly dark (or light), with the remainder of the scene composed of materials that are spectrally extremely different. Common examples of this would include scenes dominated by bodies of water, snow, or uniform vegetation cover, but with the remainder of the scene filled with materials that are in high contrast to the dominant material.

The decorrelation stretch is also limited in the sense that only three of the available channels are used to generate the output image. Information that is present only in a channel that has not been selected for input is lost to this product. Given the high degree of interchannel correlation, this will only rarely limit the quality of the output image. The channels selected for browse products will guarantee that if a given channel is not selected, at least one, and usually both of the adjacent channels will be selected.

Finally, the decorrelation stretch algorithm is a method of color enhancement that exploits whatever interchannel differences that may exist. Implicit in this technique is the assumption that the differences are real, and not noise or processing artifacts. The algorithm single-mindedly produces a color enhanced output; if noise is a major component of the scene variation, the algorithm will enhance those noise differences to produce an output that, while colorful, will be painfully noisy.

It is not practical, or within the scope of this product, to examine each scene for the presence of noise, to characterize the type of noise present, to determine the most effective remedy for such noise, and then perform noise removal. However, during the initial post-launch period, the decorrelation stretch product will be closely examined for the appearance of consistent, systematic noise. If such noise is found, and suitable corrective action can be determined (and the corrective action is

not applied earlier in the data processing stream), this algorithm will be modified to clean the input images of the offending noise prior to the statistics gathering step.

Browse Product Considerations

There appears to be some confusion, or at least lack of agreement, regarding the use of this algorithm to create both an on-demand standard data product and a browse product. The questions and issues in contention are primarily procedural. As such, the following discussion might be more appropriate in an implementation plan, rather than in a theoretical basis document. However, these issues came to light during the review of this document, and a discussion here is needed to address the concerns of the reviewers.

ASTER Level 2 data products are to be generated "on-demand" (when a user specifically requests the product for a scene), rather than "routinely" (the product generated for all data acquired). Typically, the browse product would be a data product extracted from the routine processing stream, formatted in a standard fashion, then stored in a manner for easy retrieval. Since there is no routine processing stream involving ASTER Level 2 data products, an alternative approach is needed.

The decorrelation stretch was selected as the product to be routinely (i.e., for all scenes gathered) generated and placed into the browse catalog because it best meets three criteria: utility, ease of generation, and robustness. First, it provides a good visual clue to the utility of the data. That is, the viewer can readily determine whether the region of interest in a scene is cloud-free, noise-free, and of a quality appropriate for the task envisioned for the data. Second, the decorrelation stretch is computationally relatively easy to generate. Third, there are very few conditions that would prevent the generation of this product.

For browse product generation, all ASTER scenes that have been processed through Level 1b, and are being archived at a U.S. DAAC will be decorrelation stretched, prior to (or concurrent with) the Level 1b data being archived. One product from each of the three subsystems (VNIR, SWIR, and TIR) will be generated for each of these scenes, using the defaults for all user options. The resulting images will be formatted, subsampled, and compressed in a manner to make them compatible for browse catalog usage. It is anticipated that the formatting, subsampling, and compression methods will be mandated by the custodian of the browse catalog, for the sake of consistency across the various EOS instruments. The browse format images will be delivered to the browse catalog, where they will permanently reside, but the decorrelation stretched images will not be saved in their original format and resolution.

When a user requests a decorrelation stretch product of previously acquired data, the product will be regenerated.

References

- Gillespie, A. R., Kahle, A. B., and Walker, R. E. (1986), Color enhancement of highly correlated images. I. Decorrelation and HSI contrast stretches, *Remote Sensing of Environment*, 20:209-235.
- Kahle, A. B., Palluconi, F. D., Hook, S. J., Realmuto, V. J., and Bothwell, G. (1991). The advanced Spaceborne thermal emission and reflectance radiometer (ASTER), *International Journal of Imaging Systems and Technology*, vol. 3, pp. 144-156.
- Loeve, M. (1955). *Probability Theory*, D. van Nostrand Company, Princeton, N.J.
- Press, W. H., Flannery, B. P., Teukolsky, S. A., Vetterling, W. T. (1986), *Numerical Recipes*, Cambridge University Press, Cambridge, U.K., 818p.
- Rothery, D. A. (1987), Decorrelation stretching as an aid to image interpretation, *International Journal of Remote Sensing*, vol. 8, pp. 1253-1254.
- Soha, J. M., and Schwartz, A. A. (1978), Multispectral histogram normalization contrast enhancement, *Proceedings of the 5th Canadian Symposium on Remote Sensing*, Victoria, BC, Canada, pp. 86-93.
- Taylor, M. M. (1973), Principal components color display of ERTS imagery, *Third Earth Resources Technology Satellite-1 Symposium*, 10-14, December, NASA SP-351, Vol. 1, Section B, pp. 1877-1897.
- Yamaguichi, Y., Tsu, H., and Fujisada, H. (1993). A scientific basis of ASTER instrument design. *SPIE Proceedings*, pp 150-160

Appendix

Included here are three examples of the decorrelation stretch, applied to TIMS data. Figures 1 and 2 employ TIMS channels 1, 3, and 5, which correspond to ASTER TIR channels 10, 12, and 13. Figure 3 uses TIMS channels 2, 3, and 5, which correspond to ASTER TIR channels 11, 12, and 13. In each case, the image on the left side of the figure uses a channel-by-channel color contrast enhancement, while the right side image displays the same data after the application of a decorrelation stretch.

Figure 1 is a TIMS scene of the Maricopa County Agricultural Center, located south of Phoenix, Arizona. The data were acquired on September 7, 1991. The varying colors of the agricultural fields in the right hand image are indicators of the stages of cultivation of the fields, from plowed fields to mature plants. The brightness indicates the temperature of a particular location, which for plowed fields is related to soil moisture, and for vegetated areas is related to plant dryness and/or stress.

Figure 2 is from the north slope of the Mauna Loa Volcano on the island of Hawaii. The data were acquired on November 1, 1985. The distinctive colors in the right hand image are the result of a sequence of lava flows over time. The colors relate to the flows' relative age, roughness, and degree of alteration from weathering and volcanic gases. Several of these flows may be readily distinguished from each other in the decorrelation stretched image, while their differences are not apparent in the left hand image.

Figure 3 contains the Puu Oo vent of the East Rift Zone of the Kilauea Volcano near the southeastern coast on the island of Hawaii. The data were acquired on September 30, 1988. The sulfur dioxide contained in the vented gases produces a characteristic yellow-orange track in the decorrelation stretched image, extending down and to the right from the Puu Oo vent. A second, fainter track of another sulfur dioxide containing cloud can be seen extending from a nearby vent just above the top center of the image to the right side of the scene.

The correlation matrices for these three scenes are reported in the table that follows.

The final page contains a chart that displays the ASTER Standard Product interdependencies.




Figure 1





Figure 2




Figure 3

Table 1

	Band 1	Band 3	Band 5	
1	1.000			Maricopa,
3	0.902	1.000		Arizona
5	0.910	0.734	1.000	September 7, 1991
	Band 1	Band 3	Band 5	
1	1.000			Mauna Loa,
3	0.964	1.000		Hawaii
5	0.985	0.961	1.000	November 1, 1985
	Band 2	Band 3	Band 5	
2	1.000			Puu Oo Vent
3	0.962	1.000		Hawaii
5	0.980	0.960	1.000	September 30, 1988

ASTER Product Inter-Dependencies



¹Produces a cloud mask that is incorporated into other products

²Computed simultaneously with Radiance at Surface

³Refers to a database of DEM data regardless of the source