

Saturation And Value Modulation (SVM): A New Method For Integrating Color And Grayscale Imagery

By David W. Viljoen¹ and Jeff R. Harris²

Geological Survey of Canada

615 Booth St.

Ottawa, ON, K1A 0E9

Telephone: (613) 995-1207

Fax: (613) 995-9273

e-mail: viljoen@nrcan.gc.ca; harris@nrcan.gc.ca

ABSTRACT

Algorithms for integrating color imagery with grayscale imagery have long been an important feature of many remote sensing (RS) image analysis and geographic information systems (GIS). Traditional methods for data integration include Red-Green-Blue (RGB) /Hue-Saturation-Value (HSV) transformation and RGB modulation. However, these techniques are either inflexible or present a compromise between the quality of the color and the contribution of the shading. Furthermore, these techniques can also result in serious color distortions. Layer transparency is another popular technique for integrating data that is available in most RS and GIS software packages. However, optimal integration of color and grayscale imagery is difficult to achieve using this method.

We briefly review the shortcomings of these traditional image integration methods and introduce a new method (Saturation-Value-Modulation [SVM]) for raster image integration developed by David Viljoen at the Geological Survey of Canada. SVM is flexible and does not compromise the color or grayscale components of the resulting integrated image. The general concepts behind this algorithm as well as the five parameters used to control the integration process are discussed. Various examples of how SVM can be used to integrate various geoscience data are also presented. Finally, we provide a brief overview of the ESRI ArcGIS implementation of SVM, though we do not include a detailed presentation of the actual Visual Basic code or the algorithm.

The ArcGIS map document (MXD) that contains the VBA (Visual Basic for Applications) code is available for download for those who wish to use SVM.

INTRODUCTION

There are two primary reasons for integrating a color image with a grayscale image. The first is to provide visual enhancement of a single dataset by combining differ-

ent characteristics. For example, a color image of a digital elevation model (DEM) can be integrated with a grayscale image of the shaded relief DEM (Figure 1). The second is to visualize the relationship between two very different types of data. For example, gamma ray spectrometer data can be combined with Landsat Thematic Mapper band 7 (Figure 2). Many methods have been developed to integrate imagery in remote sensing image analysis and geographic information systems. It is instructive to review a few of these methods to appreciate some of the advantages of the SVM method.

Modern remote sensing software and GIS often have a layer transparency feature that facilitates the integration of data and allows the user to increase or decrease the transparency of one layer to reveal the layer that would otherwise be hidden. The advantage of this method is that it is instantaneous, as it does not involve pixel-by-pixel computations and color transformations. This method can be used with two or more color images or a color image and a grayscale image. The resulting integrated image is a weighted interpolation of the colors of the contributing images. The disadvantage of this method in integrating color and grayscale imagery is that the resulting integrated image compromises either the color or the shading (Figure 3).

Remote sensing and GIS software often have tools for performing transformations between Red-Green-Blue (RGB) and Hue-Saturation-Value (HSV) color models. Figure 4 graphically shows the components of the HSV model where hue is the dominant wavelength of the color, saturation is the presence or absence of color, and value is the brightness and darkness. Color transformations involve pixel-by-pixel conversion of RGB color components into equivalent HSV components. Integration of color and grayscale imagery is achieved by replacing the value component (V) with the values from the grayscale image (Figure 5). One of the problems with this technique is that the value component is often important in defining colors in the color image (Harris et al. 1990, 1994). That

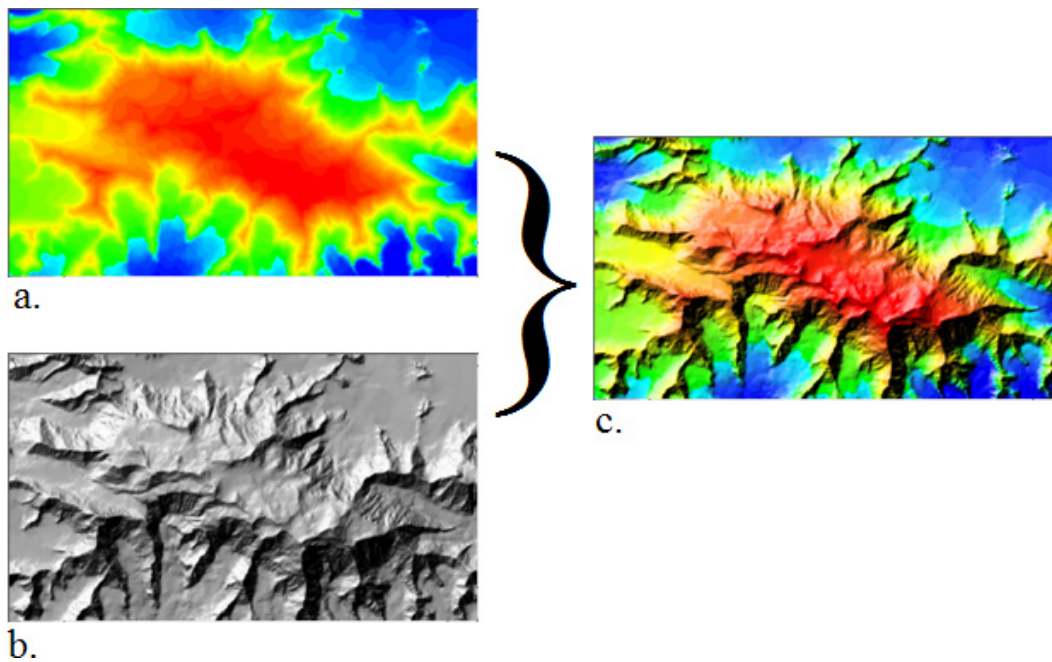


Figure 1. a. Color image of a digital elevation model (DEM) of Mt. Logan. b. Shaded relief of Mt. Logan DEM. c. Integrated image using Saturation-Value-Modulation (SVM) method.




Figure 2. a. Ternary gamma ray spectrometer color composite image K-Th-U (RGB) – imagery supplied by E. Schetselaar – ITC b. Landsat TM 7 c. Integrated image using Saturation-Value-Modulation (SVM) method.




Figure 3. Layer transparency feature of ArcGIS.




Figure 4. Hue-Saturation-Value color model.




Figure 5. Traditional Value Replacement Method of image integration.

is, often the difference between lighter and darker colors in an image is higher and lower numbers representing the value component of the colors. If the value components are modulated or replaced, then the difference between the colors will be changed or eliminated. Figure 6 shows how replacing the value component of a dark and lighter green results in an image where the dark and lighter greens cannot be differentiated. In this example, the only difference between the two greens in the original color image is in the value component. A second problem with this technique is that the original colors can be corrupted when the value component is replaced (Harris et al., 1990, 1994). For example, yellow can appear as dirty green in the integrated image, and red can appear brown. Another problem with this method is that, if the saturation of the colors is low, then replacing the value component results in an image where the saturation of the color is further reduced (Figure 7).

Another traditional method of image integration involves pixel-by-pixel multiplication of the RGB components by the grayscale values scaled between 0 and 1 (Figure 8). The main problem with this method is that the scaled values of nearly all pixels in a grayscale image are less than one, so the colors in the resulting integrated image are darker than the original image. This scaled values problem can also corrupt the apparent hue of the color. For example, a yellow might appear to be some kind of green (Figure 9).

Both value replacement and RGB modulation methods offer very little flexibility on how the integration is performed, and there are few or no parameters that can be used to control the result of the calculations.

Unlike layer transparency, the SVM method is not interactive and involves pixel-by-pixel computations similar to those associated with the value replacement and RGB modulation methods. However, the SVM method provides more flexibility on the integration process which results in

integrated images that are superior to those produced by traditional methods.

OVERVIEW OF THE SATURATION-VALUE-MODULATION (SVM) METHOD

The saturation and value color components of an illuminated object change with the angle of incidence. For example, Figure 10 shows a cylinder illuminated from the right side. The colors of the parts facing the source of illumination have a lower saturation and a higher value component, whereas the parts facing away from the source of illumination have a higher saturation and lower value. In the area around the “cutoff” line (Figure 10), the saturation and value components will be that of the natural color of the object.

In Figure 11, the “shade value” (x-axis) is zero on those surfaces that face away from the source of illumination; the highest values (e.g., 255 or 100) will be assigned to surfaces that face toward the source of illumination. Multiplier curves that range between 0 and 1 can be used to modulate the saturation and value color components depending on the shade value. The value of a pixel in the grey-scale image (Shade value) defines a vertical line that intersects the saturation and value multiplier curves. The points of intersection of this vertical line and multiplier curves are the saturation and value multipliers ranging from 0 to 1 (Figure 11). The saturation and value components of the color image at the same pixel location are multiplied by their respective multipliers. The resulting modulated saturation and value components are integrated with the original hue component to create the SVM image in HSV coordinate space. The final step is a transformation of the hue and modulated saturation and value components to RGB coordinates for display purposes.

Figure 12 presents a schematic of the SVM method. Not shown is the transformation of the color image from




Figure 6. Loss and corruption of color with value replacement method. The three circled areas have three shades of green differentiated only by the value component as shown in the table. Replacing the value component of these green areas with values in the shaded relief image results in a loss of the shades of green. Replacing the value component in yellow and red areas results in colors that appear “dirty”.




Figure 7. Low saturation colors become lower with value replacement method.

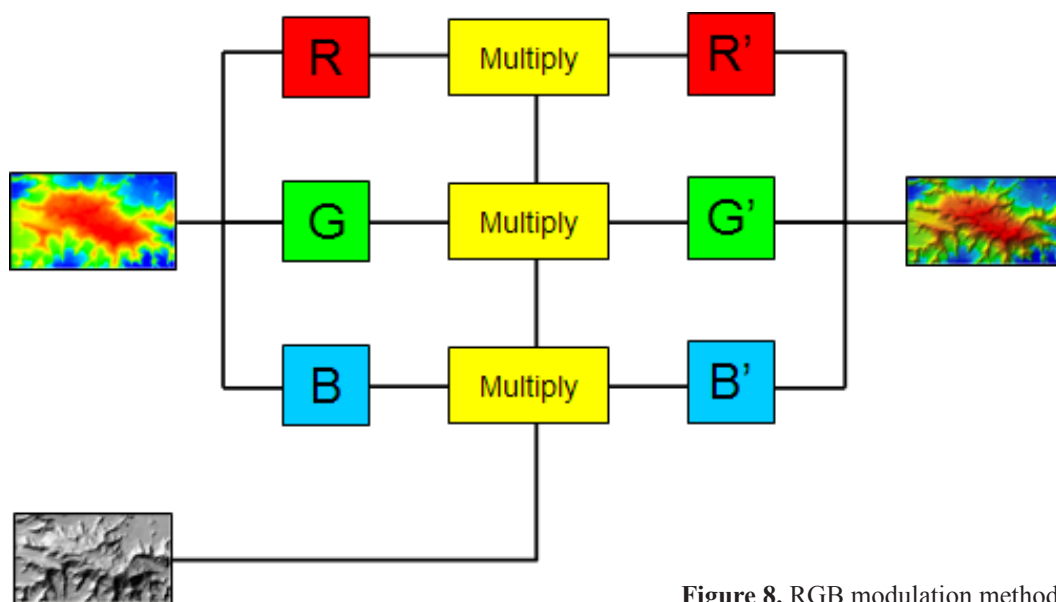


Figure 8. RGB modulation method of image integration.




Figure 9. a. Color DEM image of northern Manitoba. b. Shaded relief of northern Manitoba DEM. c. Integrated image using RGB modulation method. The box highlights an area where light green, yellow, and light red pixels have been transformed to darker colors since all pixels have been multiplied by a value from the shaded relief image of less than 1. In fact, virtually all pixels will be multiplied by multipliers less than 1 which creates a darker overall integrated image product.




Figure 10. A rainbow colored cylinder illustrates the Saturation-Value-Modulation (SVM) Concept. For surfaces facing the illumination source, the colors have full value but lower saturation. For surfaces facing away, the colors have full saturation but lower value. For surfaces at the “cutoff”, the color has full saturation and value.

RGB to HSV components. The schematic shows how the saturation (S) and value (V) components are modulated by multipliers that are determined by the pixel value in the grayscale image. The modulated saturation (S_m) and value (V_m) are used with the original hue component and transformed to an RGB composite image file that can be displayed in remote sensing software or a GIS. The multiplier curves and the parameters that define their shape are key elements of the SVM method.

SVM Parameters

There are five SVM parameters necessary to define the shape of the saturation and value multiplier curves. Together, they provide the ability to control various characteristics of the resulting integrated image.

Grayscale Value Cutoff (CutOff)

In cases where the pixel values in a grayscale image are lower than the “grayscale value cutoff” (see Figure 11), the value component (V) of the color image will be modulated, and the saturation component will be equal to the saturation in the original color image (i.e. saturation multiplier equals one). For grayscale pixel values greater

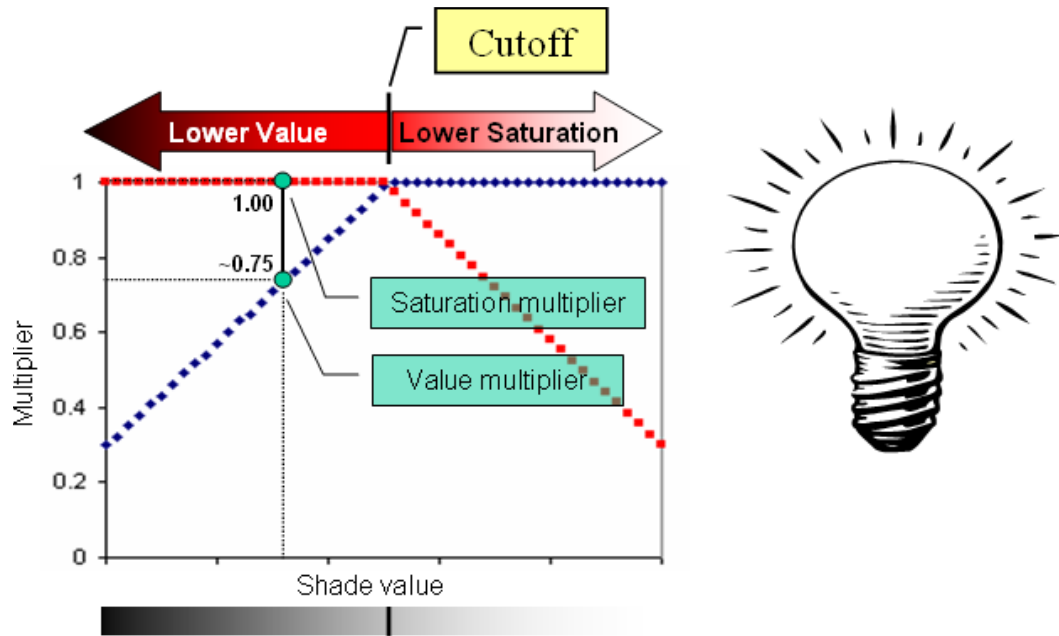


Figure 11. SVM saturation and value multiplier curves are used to model the lower saturations for surfaces that face towards an illumination source and lower color values for surfaces that face away from an illumination source. Surfaces that neither face towards nor away from the illumination source (e.g. horizontal surfaces in a digital elevation model) will have minimal or no change to their original color values.




Figure 12. Schematic of the SVM method.

than the cutoff, the saturation component (S) will be modulated, and the value component will equal the value in the original color image (i.e., multiplier equals one). At the cutoff value, the value and saturation multipliers are equal to one, so the color in the integrated image will be the same as the original color image at the cutoff value.

A general SVM rule of thumb is to maximize the number of pixels in the integrated image that have the same color as the original color image. This means that a cutoff value that maximizes the number of pixels with saturation and value multipliers of 1 should be selected. In most cases, this cutoff is represented by the peak in the grayscale image histogram.

For a typical shaded relief digital elevation model, the peak in the histogram coincides with pixels that represent horizontal surfaces. For shaded relief images, the cutoff can therefore be computed from the following equation:

$$\text{Cutoff} = 255 * \sin(A)$$

This assumes a range of values in the grayscale image is 255 and A is the illumination source altitude in degrees (0-90). For example, the peak of histogram of a shaded relief DEM with an illumination angle of 45 degrees will be approximately 180 (Figure 13).

Figure 14 shows the impact of changing the cutoff on the integrated image. Lowering the cutoff below the optimal value of 180—the peak in the grayscale histogram—results in an image with lower overall saturation (i.e., washed out colors) than the original color image. Increasing the cutoff value above the cutoff results in an image with lower overall value (i.e., darker).

Minimum Value Multiplier (Vmin)

The minimum value multiplier, which can vary between 0 and 1, determines how dark the pixels will be in areas where the grayscale pixel values are low. For example, if Vmin is 0, then for grayscale pixel values of 0 the value component of the color image will be multiplied by 0. Any color with a value component of 0 is black. This means that color pixels that have the same location as grayscale pixel values of 0 will be black in the integrated image. As Vmin increases, these same pixels will become brighter. A Vmin value of 1 will result in no value modulation. In this case, the pixels in the integrated image that have the same location as grayscale pixels with values lower than the cutoff will be the same as the original color image.

Figure 15 shows the impact of changing the value of Vmin. Note how the shadows become brighter as Vmin increases. Values greater than 0 and less than 0.4 are generally recommended.

Value Multiplier Exponent (Vexp)

The value multiplier exponent will increase or decrease the number of pixels that will have their value component multiplied by a value close to Vmin. Higher Vexp values mean that the multiplier will rise slowly from Vmin. Figure 16 shows how increases in Vexp increase the proportion of “dark pixels” in the integrated image. Values of 1 or less generally provide good results.

Minimum Saturation Multiplier (Smin)

The minimum saturation multiplier, which varies between 0 and 1, determines how much color there will be for pixels where the grayscale values are high. For example, if Smin is 0, then for grayscale pixel values of 255 (the maximum in the image), the saturation multiplier will be 0 and the saturation of the color in the integrated image will be 0. In this case the color in the integrated image will have “no” color and will typically be white or light gray. As Smin increases, these same pixels will have higher saturation and more color. If Smin is 1, then there will be no saturation modulation and the pixels above the cutoff will appear to be “flat”. Figure 17 shows the impact of increasing Smin from 0 to 0.6. Smin values between 0 and .4 are recommended.

Saturation Multiplier Exponent (Sexp)

The saturation multiplier exponent will increase or decrease the number of pixels that will have their saturation component multiplied by a value close to Smin. Higher Sexp values mean that more pixels will be multiplied by a multiplier close to Smin. Figure 18 shows how higher values of Sexp decrease the proportion of “washed out pixels” in the integrated image. Generally values between 1 and 3 provide good results.

ARCGIS IMPLEMENTATION OF SVM

The SVM method was implemented as a Visual Basic for Applications (VBA) application in ESRI's ArcMap application. It works with ArcGIS (ArcView, ArcEditor, or ArcInfo) and does not require any special ESRI extensions (e.g., Spatial Analyst).

The VBA implementation allows the user to set each of the five SVM parameters described above. The input images can be one of the following:

- 8- or 16-bit image with an associated CLR file
- 8-bit TIF image (colors embedded in file)
- 32-bit TIF image (color defined by RGB channels)
- 3 separate 8-bit RGB images




Figure 13. Histogram of pixels values for shaded relief of Mt. Logan (Figure 1). Peak is approximately $255 * \sin(45) = 180 =$ recommended SVM cutoff value.




Figure 14. SVM Parameters – Grayscale value cutoff (CutOff).

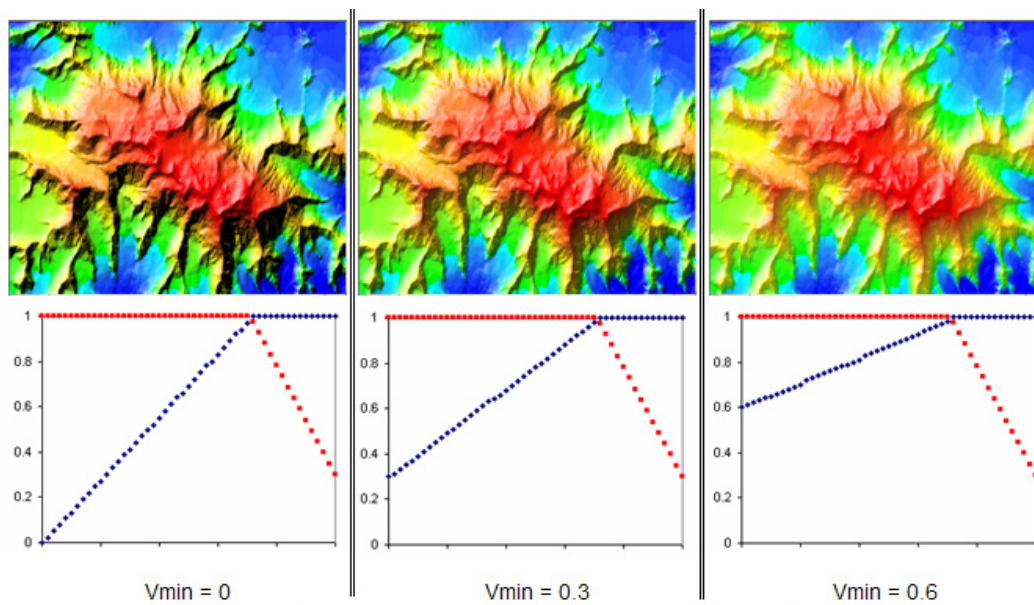


Figure 15. SVM Parameters – Minimum value multiplier (V_{min}).

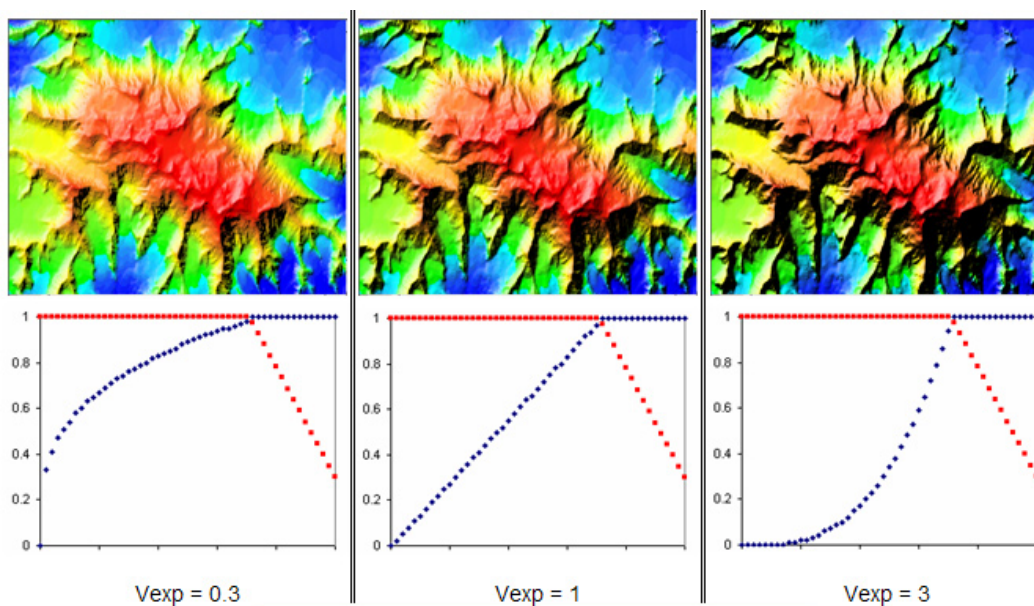


Figure 16. SVM Parameters – Value multiplier exponent (V_{exp}).

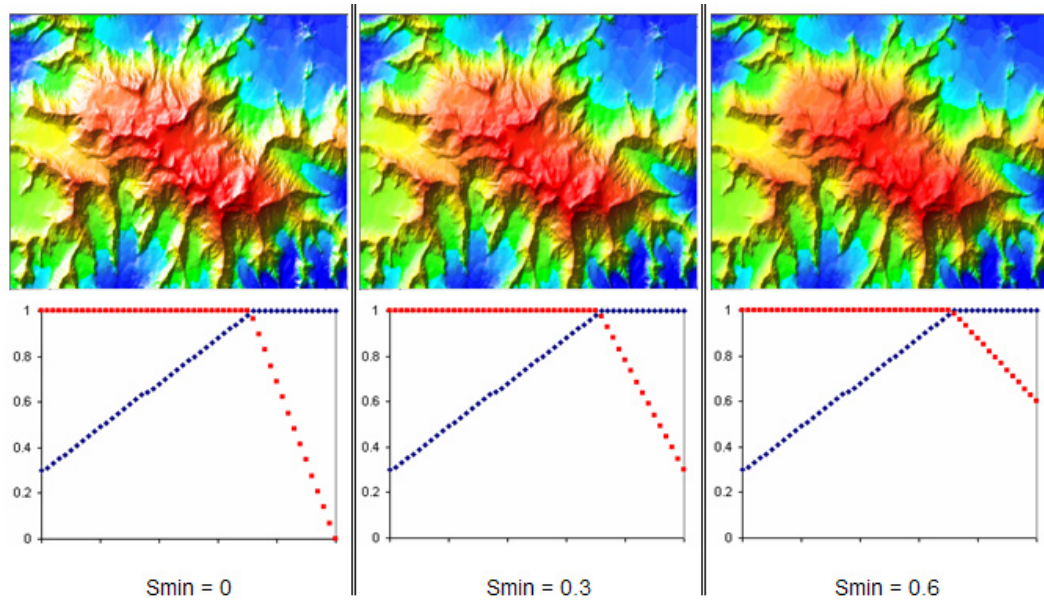


Figure 17. SVM Parameters – Minimum saturation multiplier (S_{min}).




Figure 18. SVM Parameters – Saturation multiplier exponent (S_{exp}).

The CLR file contains space delimited values for pixel value and RGB components. An example record from a CLR file might be:

18 244 64 120

This means that pixel values of 18 in the color image have RGB coordinates of 244, 64, and 120 respectively.

The output from SVM is a 3-band RGB Band Interleaved by Line (BIL) image which is easily imported or used directly by remote sensing software. It can also be easily exported to TIF or ESRI Grid format in ArcGIS.

The ArcMap document (MXD) that contains the VBA code can be downloaded from the SVM FTP site along with sample data (Viljoen, 2006).

APPLICATION OF SVM IN GEOSCIENCE

The SVM method has broad application to geoscience studies that require integration of a color and grayscale image. Many geological mapping applications, for example, require interpretation of various types of remotely sensed and geophysical data. The integration of these data types often provides images that offer a unique perspective of the Earth's surface, which enables the interpretation of many geological features that, without integration, would have been impossible to make. Furthermore, the relationships often evident in the resulting integrated imagery offer a unique interpretation

tool. The following three examples highlight the value of integrating different geoscience data and demonstrate the advantages of the SVM method over traditional integration methods.

Integrating Geological Map Units with Shaded Relief Aeromagnetics

The magnetic characteristics of rocks at, and below, the Earth's surface often reflect mappable variations in lithologies. The magnetic characteristics of rocks are measured with aeromagnetic sensors, and these measurements are often processed into colorful images that represent total field, vertical gradient, and other derivative products. Integration of colored geological units with shaded relief versions of total field aeromagnetics can provide an image that is extremely useful for geological mapping.

Geological units are usually represented by vector polygons in a GIS and, given that SVM is entirely a raster-based method, these vector polygons must be rasterized and have the same projection, pixel resolution, and map extent as the total field shaded relief image. Detailed step-by-step instructions on how to use the ArcGIS version of SVM for this kind of integration are available for download (Grant and others, 2006).

Figure 19 shows the result of using SVM to integrate rasterized geological units with shaded relief total field aeromagnetics. As can be seen, rock units, after SVM integration, are characterized by different magnetic signatures on the integrated image. This image




Figure 19. Comparing image integration methods – a. Rasterized geological map units. b. Shaded relief aeromagnetic image for the same area as the geology map. c. Integrated image using the value replacement image integration method. d. Integrated image using the SVM method. e. Integrated image using RGB modulation method.

can greatly assist mapping endeavours, as the rock units can be modified based on variations in their magnetic signature. With respect to the integration methods, note how the integrated image produced by the SVM method retains the original colors of the geology map compared to the other two traditional methods. The small squares highlight areas where color corruption and loss are evident in the integrated image produced by the other methods.

Integrating a Ternary U-Th-K RGB Gamma Ray Spectrometer Image with a Shaded Relief Digital Elevation Model

Integration of gamma ray spectrometer data, which measures the emission of the elements U, Th and K from the Earth's surface with a shaded DEM, can also provide a very useful image for geological mapping. Variations in the above radioelements often reflect different rock units as well as areas of potential mineralization. However, since the gamma ray data comprise three channels (U, K, Th) that are often correlated, the color variations in ternary images are often low. The integration of ternary imagery presents a challenge to all methods of image integration. The Value replacement method, for example, often results in almost complete color loss as shown in Figure 20. As with the previous example, the

SVM method results in an integrated image that retains the original color to a much greater degree than the other methods.

Pan Sharpening Landsat Multi-spectral With A Panchromatic Image

Pan sharpening is the term given to an image processing technique that uses higher resolution (smaller pixels) grayscale imagery to improve the visualization of lower resolution (large pixels) color or color composite images. Many satellite and airborne sensor systems have multi-spectral channels and panchromatic channels. Landsat 7, for example, provides 6 channels of 30 meter pixel resolution for portions of the visible, near infrared and short-wave infrared of the electromagnetic spectrum. It also contains a panchromatic channel that covers the entire visible part of the electromagnetic spectrum with 15 meter pixel resolution. Integrating the higher resolution panchromatic channel with a RGB color composite of the lower resolution visible channels results in a sharper color composite image as shown in Figure 21. The boxes in the images are areas that highlight the differences between these image integration methods. If the accuracy of the original colors is important, than SVM is the preferred method. If the original colors are not important, then value replacement is another pan-sharpening option.




Figure 20. Comparing image integration methods – a. Ternary gamma ray spectrometer color composite image K-Th-U (RGB) composite image. b. A shaded relief digital elevation model of the same area as the K-Th-U image. c. Integrated image using the value replacement image integration method. d. Integrated image using the SVM method. e. Integrated image using RGB modulation method.

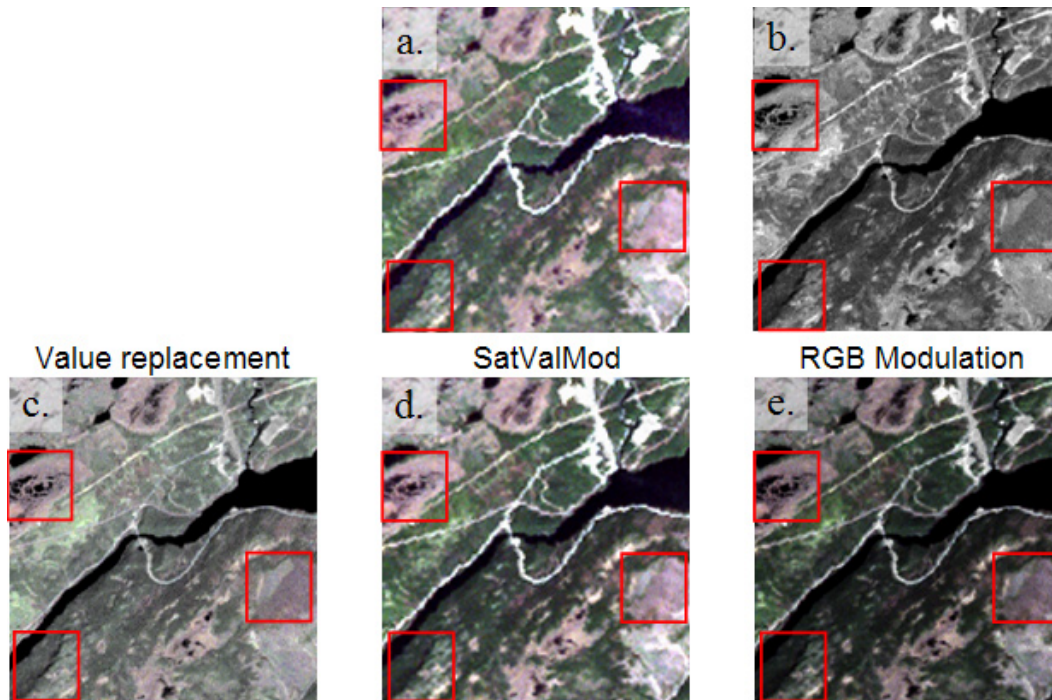


Figure 21. Comparing image integration methods. a. Landsat 7 RGB color composite image of Thematic Mapper (TM) bands 3, 2, and 1 respectively (30 meter pixels). b. Landsat TM Band 8 (15 meter pixels) c. Pan sharpened image using the value replacement image integration method. d. Pan sharpened image using the SVM method. e. Pan sharpened image using RGB modulation method.

CONCLUSIONS

The Saturation-Value-Modulation (SatValMod or SVM) method is based on the real-world concept of darker colors for surfaces in shadow (lower value component) and color loss (lower saturation component) on illuminated surfaces. This real-world concept is implemented in SVM as a pair of multiplier curves that modulate the saturation and value components of colors in the color image. These modulated saturation and value components are combined with the original hues in the color image and transformed to red-green-blue components for display.

The SVM method of integrating color and grayscale imagery provides superior results over many other integration methods because there is no compromise between color and shading, as is the case for layer transparency. In addition, there is no distortion of colors, as can result from value replacement, RGB modulation, and other methods. Unlike the other integration methods, which provide little or no control over the integration process, SVM uses five different parameters that provide a great deal of control over the characteristics of the resulting integrated image.

The use of color is important in conveying geoscience information such as geological units, geophysical

properties, radiometric characteristics, and many others. Visualizing the relationships between these data is made possible through image integration techniques. The SVM method is a superior algorithm for integrating color and grayscale imagery, which results in integrated images that preserve the original color and grayscale characteristics of the input imagery. SVM is freely available from the SVM FTP site listed in Viljoen (2006).

REFERENCES

- Harris, J.R., Murray, R., and Hirose, T., 1990, IHS Transform for the Integration of Radar Imagery with Other Remotely Sensed Data: *Journal of Photogrammetric Engineering and Remote Sensing*, Vol. 56, No. 12, p. 1631-1641.
- Grant, G., Fraser P., and Viljoen, D., 2006, How to SVM: accessed at <ftp://nrd:imanrd2@ftp.gis.nrcan.gc.ca/viljoen/downloads/satvalmod/HowToSVM.pdf>.
- Harris, J.R., Bowie, C., Rencz, A.N., and Graham, D., 1994, Computer Enhancement Techniques for the Integration of Remotely Sensed, Geophysical and Thematic Data for the Geosciences: *Canadian Journal of Remote Sensing*, Vol. 20, No. 3, p. 210-221.
- Viljoen, D., 2006, Saturation-Value-Modulation FTP site: accessed at <ftp://nrd:imanrd2@ftp.gis.nrcan.gc.ca/viljoen/downloads/satvalmod/>.