Colour Maps for Relief Shading

Optimising the presentation of spatial datasets

Peter Kovesi *- The Centre for Exploration Targeting*

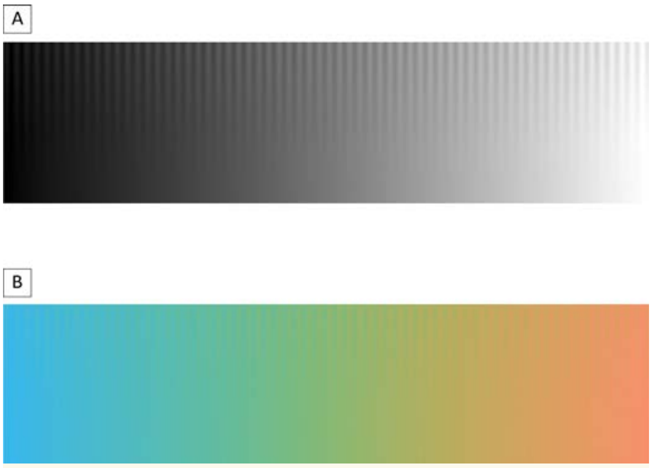

Figure 1: A grey scale colour map (A) and a constant lightness map (B). Both are rendered on a test image consisting of a sine wave superimposed on a ramp function. The sine wave amplitude is progressively reduced to zero towards the bottom of the test image. Notice how the sine wave pattern is almost imperceptible when rendered with the isoluminant colour map.

Relief shading can be a very effective way of presenting spatial datasets. By treating the data as if it is a 3D surface and generating shading corresponding to the surface being 'illuminated' from some direction, we can use the eye's innate ability to interpret shading patterns to invoke a perception of 3D morphology. However, while interpretation of the 'form' of data is enhanced, any sense of the actual data values is diminished because shading only depends on the surface gradient. This can be remedied to some degree by adding colour to convey the value information that is lost in relief shading. Such use of colour in conjunction with relief shading can also enhance the perception of shape induced by the shading. However, if colour is mis-used it is also potentially destructive of the visualisation benefits of relief shading.

An isoluminant colour map – one formed from colours of constant lightness – is very poor at representing structures in the data (Fig. 1). When an image rendered with such an isoluminant map is combined with a shading pattern where the colour gradients are not aligned with the shading gradients, however, then we can produce an amplification of the 3D shading perception. On the other hand, if a shading pattern is rendered with a colour map with significant lightness gradients, those variations in colour intensity can themselves create an implicit shading effect that can interfere with and disrupt the appearance of 3D structure induced by the original shading pattern – leading to poor visualization potential. Looking at the lower-right panel of Fig. 2, notice how the diagonal shading bands are no longer uniform in their darkness, and this confusion in shading pattern is further compounded where the darker blue regions lie alongside the shading bands. Compare this to the image in the top-right panel of the figure and note how the shading has been left untouched by the colouring.

In this example the shading pattern is applied to the colour image multiplicatively. To achieve the perception of a coloured surface being shaded, the luminance of the colours needs to be modulated by the relief shading. This differs from some GIS implementations where it is only possible to combine a shading image with a colour image via a transparency blending of the two images – a weighted sum. This is an inefficient mechanism to use, as it will not result in any amplification of the perception of 3D.

This example is something of a special case because the spatial frequency of the colour lightness variations is close to that of the relief shading variations, thereby maximizing their potential interference. While interesting, this synthetic example is also atypical in that the shading pattern only has a single frequency component. Most data sets derived from natural systems have an amplitude spectrum that decays inversely proportional to the frequency raised to some power. That is, the amplitude spectrum is roughly proportional to $1/f^p$

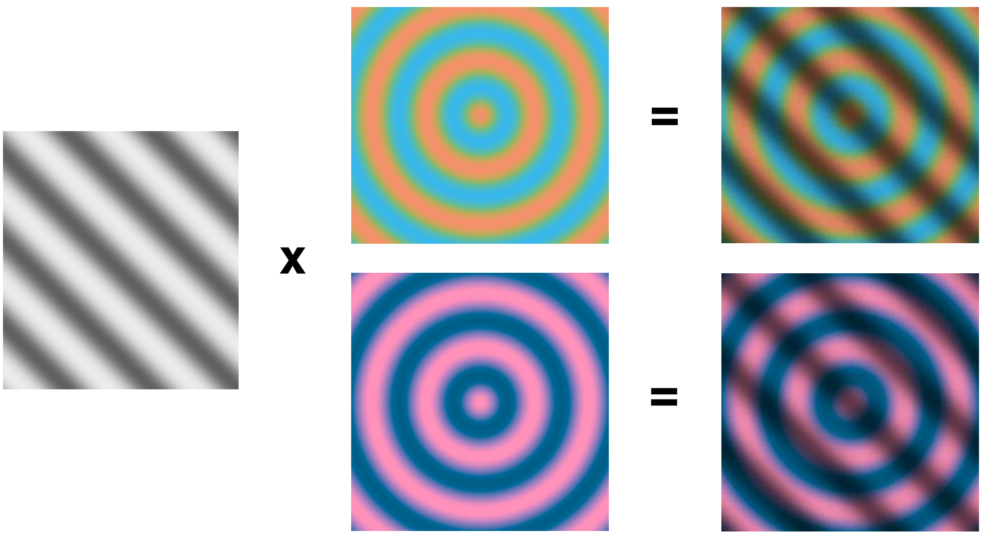

Figure 2: When a shading pattern is combined with an isoluminant image with colour gradients that are not aligned with the shading gradients an amplified 3D perception of the structure is obtained (top right). Combining a shading pattern with an image having strong lightness variations can disrupt the perception of 3D structure (bottom right).

where p typically ranges between 1 and 2. If the relief shading pattern has a more distributed frequency spectrum of this form, the colour-shading interaction effects seen in this simple sinusoidal shading pattern are not generally so strongly reproduced.

Empirically then, our work shows that as long as the image is not closely matched to the frequency spectrum of the relief shading with which it is combined, there is generally no need to employ an isoluminant colour map – especially if the image is predominantly of lower frequencies than the relief shading. The results presented here would indicate that the common practice of draping a colour image over the relief shading of the dataset from which it is derived is largely valid, and unlikely to create any perceptual problems. Forming a relief image from a data set is akin to taking a derivative of the surface – in effect amplifying the spectral content of the image as a function of frequency. Thus the original data will typically have a stronger low frequency content than the relief shading image. In the example shown in Fig. 3, the DEM amplitude spectrum falls away at approximately $1/f^{1.7}$ whereas the amplitude spectrum of the raw relief shading falls away at approximately $1/f^{1.2}$. This difference appears to be more than sufficient to avoid adverse interaction between the two. Another reason for expecting little interaction is that the image gradient values (and hence shading values) are generally independent of the image data values themselves.

Fig. 4 presents a comparison of relief shading overlain by the original DEM data values rendered in an isoluminant colour map, a low contrast colour map, and a grey scale map. What is quite surprising is that even for the grey scale case there is little disruption of the relief shading pattern and, in addition, the highest and lowest regions in the surface can still be identified. At the same time, however, none of the potential 3D perception amplification benefits of using colour – as seen for the simple sine wave shading example above – are apparent in the image. It would appear then that once the relief shading pattern is, in some sense, rich enough, the addition of colour makes little difference.

As mentioned earlier an important advantage of combining relief shading with an image that has been derived from the data is that it enhances communication of the overall data properties – metric information – in addition to the form and structure provided by the shading. A raw grey scale relief image conveys only a sense of the local surface normal information, offering no insight into absolute data values. If the data range is very large this can be useful as the relief shading acts as a form of dynamic range reduction allowing small scale features to be enhanced within an arbitrarily large range of data values. However, in other cases the loss of any sense of absolute data value can be a disadvantage. Overlaying an image derived from the data values overcomes this problem – allowing the best of both worlds.

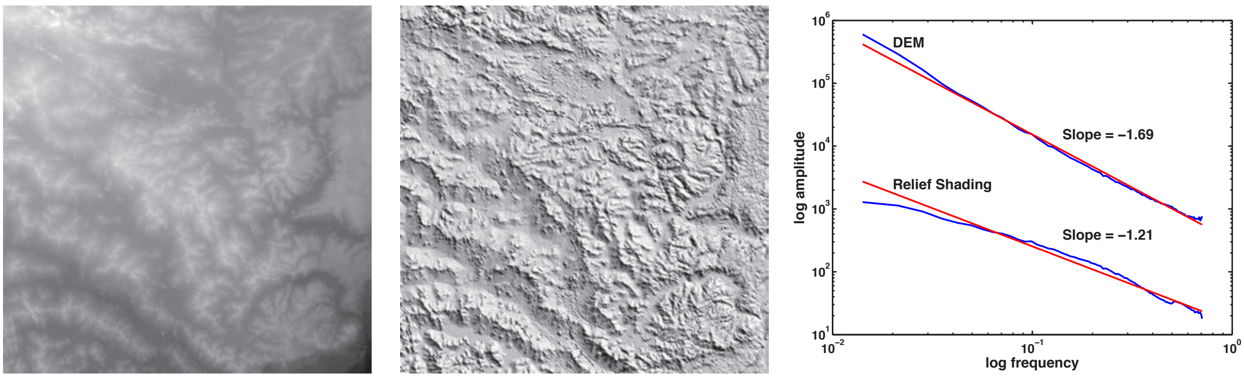

Figure 3: A DEM, a relief shading of the DEM, and a log-log plot of the amplitude spectra of the DEM and its relief shading image. The lines of best fit, from which the slopes are derived, are plotted in red.

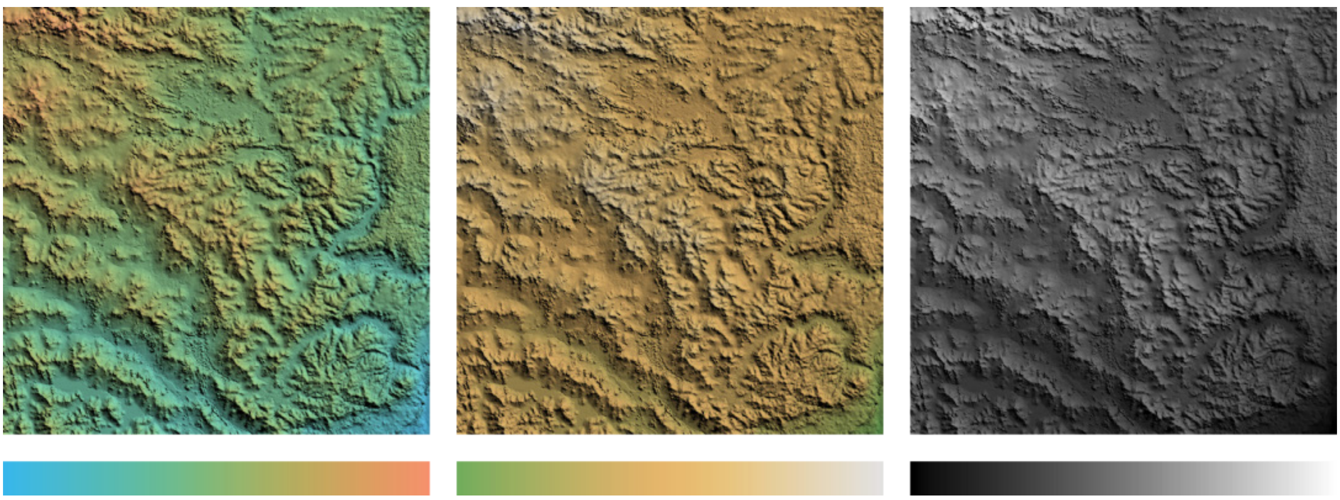

Figure 4: Relief shading combined with images of the original DEM data rendered with an isoluminant colour map, a low contrast green-brown-white colour map and a grey scale colour map.

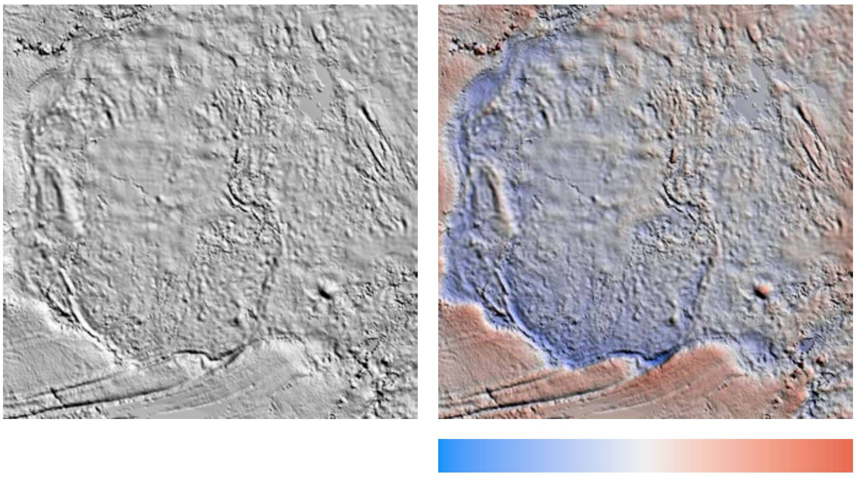

Figure 5: Relief shading of residual gravity data of West Africa and the relief shading combined with a diverging colour map image. Note the use of a light colour map to compensate for the darkening induced by the shading.

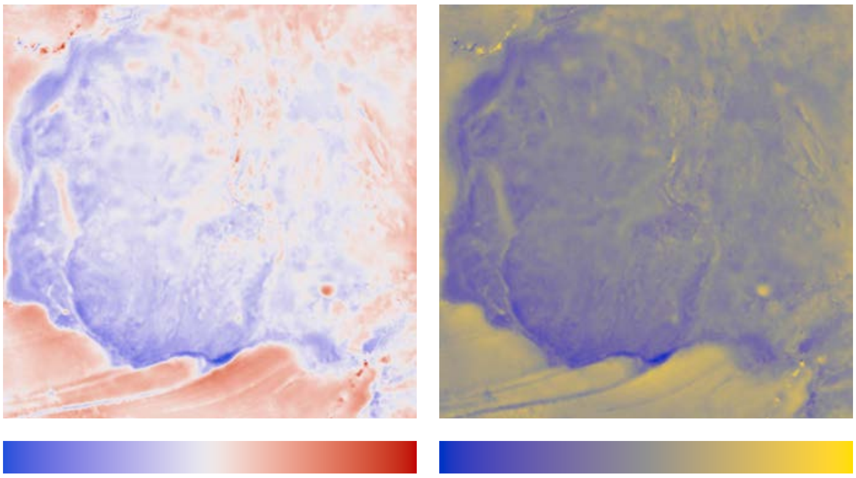

Figure 6: Residual gravity data rendered with a diverging colour map (left) and with a linear-diverging map (right). Note that a standard diverging colour map has a perceptual dead spot in the middle that makes resolving features at the mid point difficult. This is why structures in the white regions of the left image are hard to resolve. This problem is avoided using a linear-diverging map because the lightness values of the colours in the map increase in a linear fashion.

Using a residual gravity image of West Africa as an example (Fig. 6) we can see that a relief shaded image allows small scale structures to be identified readily – but it is hard to get a sense of the magnitude of the deviation of features above and below zero. Combining the relief shading with an image of the data rendered via a diverging colour map allows the fine structures to be seen in conjunction with the magnitude and polarity of the data values.

Summary

Relief shading combined with a coloured image, even a grey scale image, can be a very useful way to present data. If the frequency content of the coloured image is significantly different from that of the relief shading then no particular precautions are needed with the colour map, other than ensuring it does not have any particularly dark colours that could mask the shading altogether. As a general guide it is probably wise to use a low contrast colour map to minimise any such potential disruption to the shading pattern.

However, if the data from which the relief shading is being generated from is smooth and/or has a limited spatial frequency content then it becomes important that an isoluminant, or very low contrast, colour map is used. If the image being combined with the relief shading is derived from another source – perhaps when combining magnetic and gravity data, for example – then one should also be mindful that the frequency content of the two images may interfere. Should this be the case then, again, an isoluminant colour map should be used.

For more information about the CET Perceptually Uniform Colour Maps project visit: **http://www.cet.edu.au/research-projects/geophysics-and-image-analysis/projects/colour-maps-with-uniform-perceptual-contrast**

AUTHOR INFORMATION

Professor Kovesi is a member of the Geophysics and Image Analysis Group within CET. His research interests are focused in low-level computer vision, with much of his recent work concentrated on using local phase information in images for feature detection, symmetry/asymmetry analysis, texture analysis, and feature matching. Some of these techniques have been commercialised through the CET Grid Analysis and the CET Porphyry Detection extensions for Geosoft Oasis Montaj.

Peter also has research interests in photogrammetry and 3D reconstruction, forensic image processing and biometrics.

Further Reading

Bak, P., Tang, C., and Wiesenfeld, K., 1987, Self-organized criticality: An explanation of the 1/f noise. *Physical Review Letters*, 59: 381-384.

Field, D., 1987, Relations between the statistics of natural images and the response profiles of cortical cells. *Journal of the Optical Society of America A*, 4: 2379-2394.

Kingdom, F. A., 2003, Color brings relief to human vision. *Nature Neuroscience*, 6: 641-644.