

Interactive Multi-Image Blending for Data Visualisation and Interpretation

Peter Kovesi

Centre for Exploration Targeting
 The University of Western Australia
 35 Stirling Highway, Crawley WA. 6009
 peter.kovesi@uwa.edu.au

Eun-Jung Holden

Centre for Exploration Targeting
 The University of Western Australia
 35 Stirling Highway, Crawley WA. 6009
 eun-jung.holden@uwa.edu.au

Jason C. Wong

Centre for Exploration Targeting
 The University of Western Australia
 35 Stirling Highway, Crawley WA. 6009
 jason.wong.uwa.edu.au

SUMMARY

The ability to integrate data from a range of different images is often a crucial requirement for successful interpretation. Interactive multi-image blending is presented as a tool for facilitating the interpretation of complex information from multiple data sources. Traditionally, image blending has only been considered for cross-dissolving effects between two images. However, it is common for there to be more than just two images of interest in an interpretation task. We have developed a family of different multi-image blending tools to fill this need. These have been designed to support a number of different interpretation tasks and image types. For image blending to be a useful tool for multiple image interpretation it is important that the association between features and individual input images remain identifiable and distinct within the blend. We argue that interactivity of the blend is an important component for achieving this. Blending can also be usefully employed to interactively explore parameter variations for enhancement techniques. Often the best parameter values to use cannot be known beforehand, and it is common for different regions of an image to require different parameter values for best enhancement. By preparing a set of images processed over a sequence of scales and parameter values, and then interactively blending between these images, the interpretation of a data set can be greatly facilitated.

Key words: image blending, visualisation, interpretation.

INTRODUCTION

Multi-modal imagery is increasingly important for the interpretation of complex spatial structures. The spatial data available might include gravity, magnetic, topography, geochemistry, and hyperspectral imagery. In addition to data from different modalities some of the different images to be interpreted may be a result of different treatments of these datasets using varying colour maps, enhancement techniques, spatial filtering and so on.

The traditional interface for switching between multiple images is a list of the images with checkboxes against them that operate in a mutually exclusive manner. Only one image can be displayed at a time and switching images requires the user to move their attention away from the image display and attend to a list of checkboxes. This limits one's ability to compare and assimilate information across multiple images.

If we initially limit ourselves to interpreting data from just two images an alternative approach is to employ cross dissolving of the images. Cross dissolving, or blending, between two images, I_1 and I_2 , to obtain a composite image I_b is typically achieved by a linear interpolation between the two images (Porter and Duff 1984, Smith 1995, Brinkmann 1999).

$$I_b = w_1 I_1 + w_2 I_2$$

where: $w_1 + w_2 = 1$ and $0 \leq w_1 \leq 1$. An interactive blend can be readily facilitated by providing the user with a slider that controls the value of the weights w_1 and w_2 .




Figure 1. Interactive blending between two images via a horizontal slider.

Image blending has traditionally been used for cross dissolving effects in movies and typically the cross dissolve only occurs between two images. This paper develops and extends this approach to allow blending operations across multiple images.

SWITCHING, BLENDING AND MORPHING

For the purposes of this paper it is useful to clarify our use of the terms image switching, blending and morphing (or at least the *aims* of switching, blending and morphing, as the implementations of all might be quite similar).

Image switching is just that, switching. Here we are seeking to conveniently switch between several images, perhaps toggling back and forth between two images at a time to identify regions of contrast and similarity. While we are not interested in obtaining intermediate blended images there may be some value in performing the switch in the form of a very short cross-dissolve in order to soften the transition and reduce image transition after-effects resulting from adaptation of the eye.

With morphing the aim is to continually create the impression of a single image. At all stages during a morph there should only be a perception of a single image, the input images should not be identifiable or distinct. When applied for movie special effects the input images may be of objects of different shape so morphing requires geometric distortion of the objects

as well as cross-dissolving of the images in order to maintain the perception of a single image.

When blending we are seeking to display two, or more, images simultaneously in a way that keeps all images distinct and allows the perception of all images simultaneously, though with varying emphasis.

Blending is appropriate when combining images of different modalities, for example topography, gravity and magnetic data. On the other hand the aim of morphing, which is to always maintain the perception of a single image, is appropriate when one is attempting to visualise a set of images that are of the same modality but differ only in some processing parameter, say scale. Given that all the input images are of the same modality and hence will share the same shape characteristics morphing can typically be implemented just with simple cross-dissolving. So, ultimately, for our purposes the implementations of switching, blending and morphing might all be achieved in the same way, via some form of linear blending. However, it is worth remembering that each has quite different aims.

MULTI-IMAGE BLENDING

Traditionally, cross-dissolving or blending has only been applied across two images at a time. We are seeking to use blending across multiple images as a means of facilitating the interpretation and integration of all the information in the images. The linear blend described earlier can be extended to an arbitrary number of images

$$I_b = w_1 I_1 + w_2 I_2 + \dots + w_n I_n$$

where $\sum w_i = 1$. However, interactively manipulating such a blend is awkward. One approach might be to provide the user with an array of sliders that can be adjusted to set the individual weights. Given that the image weights have to sum to one, $n-1$ sliders will be required. Manipulating all the sliders in order to maximise one's ability to interpret information from the images is likely to be a difficult task.

Blending three images

In the case of three images we have a system with two degrees of freedom, weights w_1 and w_2 . The third weight, w_3 , is fixed at $1 - w_1 - w_2$. The position of a mouse on a screen provides two degrees of freedom of input. This suggests a user interface as follows: A 2D 'slider' is formed using three nodes arranged in a triangle, one node for each of the three images to be blended. The user positions the mouse within the triangle and the homogeneous barycentric coordinates of the cursor position within the triangle are then used as the image weights to form the blended image. Barycentric coordinates (Coxeter, 1969) form a natural interpolation scheme across a triangle. A physical interpretation of a point's barycentric coordinates are the masses that should be applied at each vertex in order for the triangle to balance at that point. To obtain homogeneous coordinates we impose the additional condition that the masses/image weights sum to one. Positioning the cursor at the centre results in a blend consisting of the average of the three images. Positioning the cursor at a node results in just that node's image being displayed. Some example weightings that arise from various cursor positions are shown in Figure 2.

An obvious application for such a blender is an interactive ternary image for rendering radiometric data. The three input

images are potassium, thorium, and uranium images each assigned to the red, green and blue colour channels respectively. Moving the cursor within the triangle adjusts the weights for each colour channel. Positioning the cursor at the centre results in the traditional radiometric ternary image, moving the cursor away from the centre allows the user to interactively explore the relative abundance of individual components, or pairs of components.

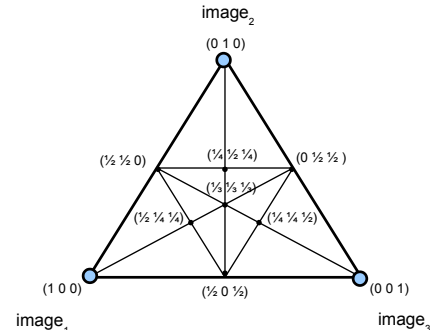


Figure 2. Using barycentric coordinates to generate three image blending weights from a 2D mouse position.

BLENDING MORE THAN THREE IMAGES

Trying to set $n-1$ sliders to blend n images is likely to be an onerous task for a user. An alternative approach is to consider ways of projecting a 2D cursor position into a, hopefully useful, subspace of the $n-1$ dimensional space describing all possible image blends.

The Image Wheel

This approach is inspired by a model of mixing different frequencies of light from the spectrum and is analogous to adjusting the Hue and Saturation axes of the HSV colour space. The interface consists of a circle with equi-spaced nodes placed around it, one node for each image to be blended. An additional node is placed at the centre of the circle. The average of all the images is assigned to this node.

The blending is performed as follows: Given the angular position of the cursor a linear blend is formed from the two images that the cursor is between. This linear blend is then 'desaturated' by performing a linear blend between it and the average image according to the radial position of the cursor. See Figure 3.




Figure 3. Blending multiple images using the Image Wheel. Here we are blending between images 1 and 5, and the average image. First a blend is formed between images 1 and 5. This is then 'desaturated' by blending with the average image according to the radial position of the cursor.

There are some limitations of this blending scheme. The ordering of the images nodes around the wheel dictate the image blends that can be obtained. In the example shown in

Figure 3 it is only possible to blend image 1 with images 5 or 2. This could be partially addressed by allowing the user to interactively alter the ordering of the vertices. A second issue is that the centre node represents a point of singularity for angular movements of the cursor. Movements of the cursor close to the centre node results in very rapid changes between the images being considered for blending. Also, if the images being blended are of different sensing modalities the use of an average image at the centre may not make much sense. In this situation the interpretation of up to perhaps four blended images is probably the limit of human ability.

However, for image sets that form a cyclic sequence, such as a set of images that have been filtered with a sequence of oriented filters, or had light shading applied in a sequence of directions, this is an ideal blending interface.

Clique Blending

This approach abandons any attempt to simultaneously blend more than two images. Instead it is designed to allow convenient comparison between any two images from a large collection.

Nodes representing all the images to be considered are arranged in a circle and every node is connected to every other by an edge, forming a clique, as shown in Figure 4. The location of the cursor on any edge is then used to perform a simple blend between the two images connected by the edge. Mouse movements are projected onto the closest edge to obtain the appropriate image weights. A degree of hysteresis is incorporated in the edge selection to prevent unwanted switching of edges as the cursor crosses other edges while traversing between the two selected image nodes.

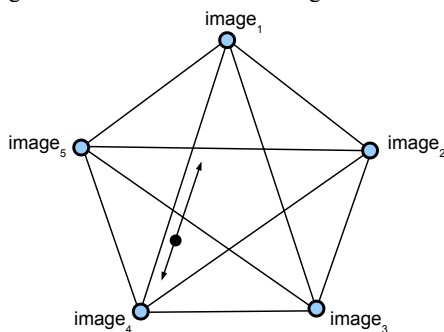


Figure 4. Blending pairs of images within a clique.

This results in a simple, easy to understand, interface which scales reasonably well as the number of images are increased. However, with larger numbers of images the variation in lengths of the edges between the nodes in the clique increases. This can make interactions less natural for the user. Also, as the clique size increases the number of edges meeting at each node increases and the user has to attend rather carefully to the interface to ensure the desired edge is being used. To overcome this we find it useful to allow the user to lock their selection of an edge. Having selected an edge the user can then drive the blending interface without having to visually attend to it, this allows all attention to be directed at the blended images. We believe the ability to drive the interface with a minimal need to visually attend to it is important. We have found this interface useful for exploring multichannel datasets such as ASTER images.

Linear Blending

Another application for blending is to provide rapid interactive viewing of an image that has been filtered in some way over a range of different parameter values. Often the best parameter

values to use cannot be known beforehand. Also it is not uncommon for different regions of an image to require different parameter values for best enhancement.

The ideal solution would be to compute a sequence of images over a finely graduated range of parameter values and then scan through them looking for the images that provide the best enhancement. However, it is quite possible that this approach will be computationally expensive and encounter memory constraints. An alternative is to precompute a fewer number of filtering results over a more coarsely quantised range of parameter values. These can then be displayed with an interactive blender to approximate the filtering results that would be obtained using a more finely graduated range of parameter values. In this context the aim of image morphing, to maintain the perception of a single image, is appropriate here.

Given that the images to be blended have a clear ordering and that we only wish to allow blending between an image and its ordered neighbours the appropriate interface is simply a linear array of nodes, one for each image, along which the cursor is positioned in order to select the desired blend. Hence the name Linear Blending, in essence this is simply a manually controlled movie player. Given that the interface is so simple it is not necessary to explicitly display it. We have used the image itself as the interface. Conceptually one can imagine the nodes being distributed across the width, or height, of the image and the input to the blender being obtained from the column, or row, position of the cursor. This provides a truly interactive image. All of the user's attention is directed at the image, there is no interface that has to be attended to.




Figure 5. Linear blending of an ordered sequence of images processed with different parameter values. The horizontal cursor position specifies a desired parameter value which is then approximated by the blending process. The diagram also shows the linear weighting functions associated with each image.

Using this interactive blending approach one can readily find the filtering result/parameter value that provides the best enhancement for any particular region of the image. We have used this interface to good effect in exploring different scales of analysis associated with the CET phase preserving dynamic range compression filter (Kovesi, 2012).

Bilinear Blending

The linear blending idea can be extended to 2D for filtering procedures that are a function of two different parameter values. In these circumstances the search through parameter space for the best values to use can be uncertain and quite laborious due to the combinatorial complexity. In this case the interpolation nodes form a rectangular grid and bilinear interpolation is used to obtain the image blend given a cursor position and its surrounding nodes. Again, as with the linear blending interface, if one imagines the interpolation nodes being distributed across the image display panel then one can use the position of the cursor within the image as the input to the blending process to approximate the desired parameter combination.




Figure 7. Bilinear blending: The image nodes form a grid distributed over the display area. The cursor position within the grid specifies the image blend to be generated.

This bilinear blending approach is also very effective in the situation where one has two images of different modalities, say gravity and magnetic data, which you want to view under some varying parameter, say the scale of phase preserving dynamic range compression (Kovesi 2012). In this case the node grid will consist of two columns. The left column being the gravity data processed over a range of scales and the right column the magnetic data processed over the same range of scales. Left-right movements of the cursor allow you to see where gravity and magnetic features coincide or differ, and up-down movements allow you to vary the scale at which features in both data sets are compared.

USING THE BLENDERS

Much of the value of an image blender comes through its interactivity. Indeed, this is crucial. A static image that displays the result of a blend is hard to interpret. You cannot readily tell which features in the resulting blend are associated with any of the input images. However, in an interactive blending situation the user is able to establish which images contribute to a particular feature by making regular exploratory perturbations to the cursor position. Thus interactive cursor movements support the aims of *image blending* in that they help us maintain the perception of the individual input images within the overall blend. This emphasises the fact that the act of interactive image blending is very much a personal process, it is hard for an onlooker to get the same perceptual benefit that the operator is gaining.

How many images can one sensibly blend? As the number of images being blended simultaneously increases, the ability to use exploratory cursor moments to identify which input image contributes to a particular feature becomes increasingly difficult. At some point the association of features with input images becomes too hard and the blend loses its usefulness. In addition, as the number of images in the blend increases, the contrast problem highlighted by Grundland et al. (2006) becomes increasingly significant. If we are blending three images the mid-blend image will be constructed from three images, each having one-third contrast. With four images at the mid point we will be combining images having only one-quarter contrast. This increasing contrast loss compounds the difficulties one has in associating features in the blend with individual input images. In practice we find the number of images that can be simultaneously blended in a useful manner is limited to three images, perhaps four, but certainly no more than that.

In many applications the value of a blender may not be in its ability to blend images. Rather its value may lie in its ability to provide a fast and convenient way of switching between different images, say, through movement of the cursor around the perimeter of the image wheel or clique blender. Images can be changed without the user having to take their eyes away from the image display panel and the ease of switching between images encourages experimentation. In this case the blending function can be changed from a linear ramp to a sigmoidal, or step, function so that transitions between images are made distinct.

Linear and bilinear blending of images that vary in one or two parameters is a simple and highly effective approach to the problem of searching through parameter space to maximise image interpretation. Interactive cursor movements play a useful, but different, role here. They allow you to explore the range of parameter values over which a feature maintains its visibility allowing the user to determine the stability and prominence of particular features.

CONCLUSIONS

We have presented a variety of novel interactive blending tools for visualising and integrating visual information from multiple image datasets. It is useful to recognise the different aims of image switching, morphing and blending and to identify which of these is most appropriate for the data sets being analysed. Key attributes for useful image blending are interactivity and an interface that minimises the need of a user to attend to a separate blending control tool. Interactivity permits the user to make associations between observed features and the input images through exploratory cursor movements. Minimising the need to attend to a separate blending control allows the user to maximise their attention on the image display. The ease of switching between images encourages experimentation and facilitates interpretation.

REFERENCES

- Brinkmann R., 1999, *The Art and Science of Digital Compositing*, Morgan Kaufmann.
- Coxeter, H. S. M., 1969, "Barycentric Coordinates." §13.7 in *Introduction to Geometry*, 2nd ed. Wiley, pp. 216-221.
- Grundland, M., Vohra, R., Williams, G.P. and Dodgson, N.A., 2006, *Cross Dissolve Without Cross Fade: Preserving Contrast, Color and Salience in Image Compositing. EUROGRAPHICS 2006*, Vol 25, No 3.
- Kovesi, P., 2012, Phase Preserving Tone Mapping of Non-Photographic High Dynamic Range Images. *Australian Pattern Recognition Society Conference on Digital Image Computing Techniques and Applications (DICTA 2012)*.
- Porter T. and Duff T., 1984, Compositing digital images. *SIGGRAPH*, 253-259.
- Smith A. R., 1995, Alpha and the history of digital compositing. Microsoft Technical Memo 7.