Compositing Colour Images

Binh Pham
School of Computing and Information Technology
Griffith University, Nathan, Brisbane, Australia

Abstract

This paper presents a method based on regression analysis for compositing colour images. The compositing process is carried out in a hierarchical fashion using two sets of images which form a Gaussian and a Laplacian pyramid. The method provides a gradual transition between composited images with minimum changes to the original data.

Introduction

There are two common situations in which image compositing is required. The first situation occurs when a set of images taken at the same time (or almost at the same time) are mosaicked together to form a bigger image, e.g. mosaicking Landsat images, or joining a spatial sequence of images to form a panoramic scene. Another situation is when two or more images of different types or from different sources are combined together to create another image for some special purposes, e.g. compositing images for graphical design or animation. In the first case, the images are of the same type and captured by the same device using the same technique. Adjacent images also have an overlap area. Although some colour differences do exists between these images and are noticeable by human eyes when the images are placed side-by-side, the amount of difference is not substantial. This is generally not true in the second case, where the colour in images to be composited could be significantly different.

Most current methods perform mosaicking by finding the seam which minimizes the sum of the absolute difference in intensity or minimizes the cost of the seam path[(e.g. [5,7,9])]. Once the seam is found, smoothing is carried out to remove any discontinuities still present in the neighbourhood of the seams. In [8], we presented a color correction method which minimize the differences between adjacent images of an image sequence. The method involved the construction of a mapping from the overlap region in one image to that of its adjacent image using a multiple regression scheme. Once the mapping is found, it will then be applied to the rest of the first image. Various schemes were introduced to take care of colours which are not present in the overlap section of first image, but exist in its non-overlap section. In many cases, this method produced images which are very well matched that the task of finding an optimal seam for joining the images is not necessary.

This paper extends the above method to deal with the second situation where two images to be joint are very different. The aim is to form a combined image in such a way that there is a smooth transition in colour from the first image to the second at the local neighbourhood of the joint, while the rest of both images remain unchanged. The method makes use of a hierarchical sets of images which are formed using the concept of Gaussian and Laplacian pyramids[1]. The next section briefly describes these pyramids and discusses their relevant properties. The method for compositing images is presented in the following section.

Gaussian and Laplacian Pyramids

Gaussian and Laplacian Pyramids were first introduced by Burt[1] as a means to calculate efficiently the correlations of an image with kernels of the same shape, but of different sizes. It was shown that the correlation of an image with a large kernel can be computed repeatedly as a weighted sum of the correlations with smaller and smaller kernels. Furthermore, this method was much more efficient than that using direct correlation or FFT. Other image processing tasks which have been carried out by a few researchers using these pyramids, includes image coding[2], image interpolation[6] and image mosaicking[3, 6].

Briefly speaking, the Gaussian pyramid is created by first applying a Gaussian-like kernel on an image to obtain an image of reduced size, whose pixels are the weighted average of those in the original image. This process is repeated so that the size of the image is reduced by half at each level and the resultant images form successive levels of a pyramid. Each step of reduction may be seen as equivalent to carrying out a linear low-pass filter and a down-sampling. Hence, the Gaussian pyramid images $G_0, G_1, \ldots G_K$ may be formed by

$$G_k(i,j) = \sum_{m,n=-2}^{2} w(m,n) G_{k-1}(2i+m, 2j+n)$$

where $k = 0, 1, \ldots, K$ and $w(m, n)$ is the weight function. Details of how to construct this weight function was described in [1,3]. The width dimension of G_{k-1}, and G_k are $2^{p-1} + 1$ and $2^p + 1$ respectively, where p is a positive integer.

At each level k, a band-pass image called a Laplacian image L_k can be created by taking the difference between the low-pass images at the previous level and this level.

As the image size of the two levels are different, the image at the current level needs to be expanded before the difference is carried out. The expanding process is equivalent to performing an interpolation of the image using the same weight function. Hence,

$$L_k = G_{k\text{-}1} EXPAND[G_k]$$

The set of images L_k forms a Laplacian pyramid. One special property of this pyramid is that the original image can be reconstructed by summing all the Laplacian images together with the lowest Gaussian image. Thus,

$$G_0 = L_1 + L_2 + \ldots + L_K + G_K$$

This sum is carried out one step at a time, from the lowest level. As the images are of different sizes, the resultant image at each step is expanded before adding to the image in the next level.

Hierarchical Compositing

There have been two attempts to employ the Laplacian pyramid for mosaicking. Burt and Adelson's method[3] first constructed two Laplacian pyramids for the two images to be joint, then formed each third Laplacian image by copying the left half of the first Laplacian image to the left of the third Laplacian image, and the right half of the second image to the right of the third Laplacian image. As all images are of odd dimension, the centre line of the third Laplacian image remains to be constructed as the average of the centre line of the left and right Laplacian image. In addition, at the lowest level, the third Gaussian image is formed from the first two Gaussian images in the same manner. When the final image is constructed by summing all of its images together with the corresponding lowest Gaussian image, a gradual transition from the left image to the right one has been achieved. This is due to the fact that the averaging of the centre line was done at different levels of resolution, hence the pixels along the centre line is most affected by the averaging process, while pixels towards the boundary are less affected and thus retain more characteristics of the original image.

Another method was introduced by Pei and Tsai[6], who constructed morphological Gaussian and Laplacian pyramids by using morphological filters instead of linear low-pass filters. Mosaicking is then performed in the same way as in the previous method, using these morphological pyramids.

Both of the above-mentioned methods are simple because they only involve taking the averages of pixels, once the pyramids are built. However, these methods do not care to what extent each image has to be adjusted to obtain the smooth transition. Our method aims to find a smooth transition between two images such that the overall amount of adjustment is minimized. This can be achieved by using multiple regression which is a common method for finding the relationship between two given sets of data[4].

To simplify the explanation, we now describe our method for compositing two images so that the overlap regions along the vertical direction vary smoothly from one image to another. The method first constructs a Gaussian pyramid from each of the overlap images. A middle strip is then extracted from each of these Gaussian images, and used for regression. The horizontal dimension of the strip which is the same for each level, depends on the number of levels in the Gaussian pyramid. It is chosen to be equal to the horizontal dimension of the image in the lowest level. In other words, in this level, the strip is identical to the whole overlap image. Figure 1 shows two Gaussian pyramids which correspond to the two overlap images. The shaded areas display the strips taken out of the image at each level of the pyramid. The regression is carried out between each pair of strips to match one to another. A detailed description of how the regression mapping is constructed to match two colour images was described in [8]. Briefly speaking, colour components of two images are used to form two sets of data. The coefficients for the mapping are found by minimizing the sum of the squares of errors which occurred when one set of data is expressed as a polynomial combination of the other set of data. Once the mapping is found, it will be used to transform the colour components of one image so that it matches the other. The regression can be carried out in various ways: to match the left image to the right one, or vice versa, or to match the left and right image to another image which is the average of the two original images.

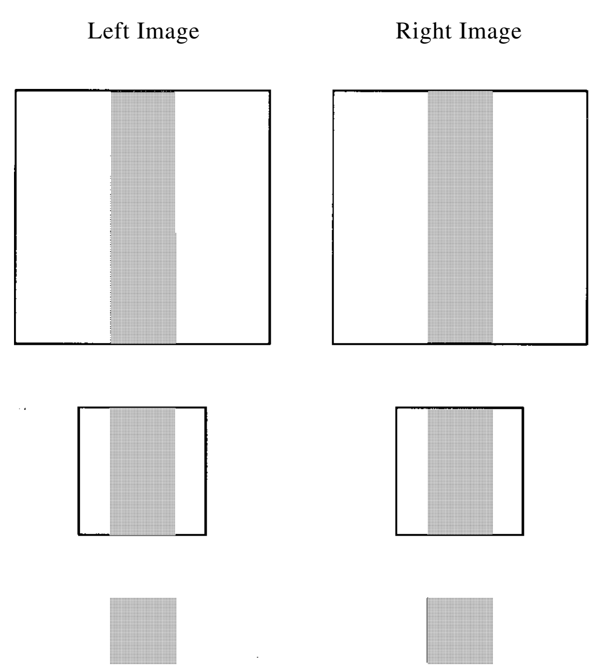

Figure 1. 3-level Gaussian Pyramids

The regression results in a set of regressed images which forms a new Gaussian pyramid. A Laplacian pyramid is created by taking the difference of successive regressed Gaussian images. The final image is then obtained by summing all of these Laplacian images together with the regressed Gaussian image at the lowest level. Since the horizontal dimension of the strip is the same in all pyramid levels, the strip at a lower level covers more contents of the overlap image than its counterpart at a higher level. Hence, the regression affects a larger area away from the centre line of the overlap image at a lower pyramid level. Thus, in the final image, the transition from one image to another occurs gradually.

One advantage of this method is that one can choose to bias the matching to favour one of the two images by performing the regression so that the left image match the right one or vice versa. The choice of the weight function also affects the outcome of the result. For example, as the weight in the centre of the Gaussian mask decreases, the Gaussian curve is broader and the final result is less sharp. The method has been described for joining images along the vertical direction, but it can be readily modified for other types of compositing. The only requirement is that an area of overlap between two images can be identified, on which regression analysis will be applied.

As the test results can only be demonstrated effectively in colour, they will be presented in the talk at the conference and are not included in this paper.

Conclusion

A method for compositing colour images in a hierarchical way has been presented, which provides a gradual transition from one image to another with minimum changes to the original data. This method will be analysed further to identify ways in which more user control can be exerted in order to bias results to satisfy some specific requirement on the colour characteristics of the final image. For example, it may be required to composite two images so that the transition from one to another is as smooth as possible, yet retaining certain hues.

References

1. Burt, P. J., Fast Filter Transforms for Image Processing, *Comput. Graph. Image Proces.*, **16**, 1981, 20-51.
2. Burt, P. J., Adelson, E. H., The Laplacian Pyramid as a Compact Image Code, *IEEE Trans. on Commun.,* **31**, 1983, 532-540.
3. Burt, P. J., Adelson, E. H., A Multiresolution Spline with Application to Image Mosaics, *ACM Trans. on Graph.,* **2**, 1983, 217-236.
4. Mendenhall, W., Scheaffer, R. L., Wackerley, D. D., *Mathematical Statistics with Applications,* Duxbury Pres., Boston, 3rd ed., 1986.
5. Milgram, D. L., Adaptive Techniques for Photomosaicking, *Trans. IEEE on Comput.*, **C-26**, 1977, 1175-1180.
6. Pei, S. C., Tsai, H. B., Image Mosaic and Interpolation by Multiresolution Morphological Pyramids, *Proc. of SPIE Indust. Soc. for Optic. Eng.,* No. **1360**, 1990, 284-292.
7. Peleg, S., Elimination of Seams for Photomosaics, *IEEE Conf. on Patt. Recog. Image Proces.*, 1981, 436-429.
8. Pringle, G., Pham, B., Colour Correction for an Image Sequence, *Proc. of Color Imaging Conference: Transforms & Transportability of Color*, November 1993, Arizona, 238-242.
9. Yang, S., Li, L., Gao P., Two-dimensional Seam-Point Searching in Digital Image Processing, *Photogram. Eng. and Remote Sens.*, **55** (1), 1989, 49-53.