

A Multiresolution Spline With Application to Image Mosaics

PETER J. BURT and EDWARD H. ADELSON
RCA David Sarnoff Research Center

We define a multiresolution spline technique for combining two or more images into a larger image mosaic. In this procedure, the images to be splined are first decomposed into a set of band-pass filtered component images. Next, the component images in each spatial frequency band are assembled into a corresponding bandpass mosaic. In this step, component images are joined using a weighted average within a transition zone which is proportional in size to the wave lengths represented in the band. Finally, these band-pass mosaic images are summed to obtain the desired image mosaic. In this way, the spline is matched to the scale of features within the images themselves. When coarse features occur near borders, these are blended gradually over a relatively large distance without blurring or otherwise degrading finer image details in the neighborhood of the border.

Categories and Subject Descriptors: I.3.3 [Computer Graphics]: Picture/Image Generation; I.4.3 [Image Processing]: Enhancement

General Terms: Algorithms

Additional Key Words and Phrases: Image mosaics, photomosaics, splines, pyramid algorithms, multiresolution analysis, frequency analysis, fast algorithms

1. INTRODUCTION

The need to combine two or more images into a larger mosaic has arisen in a number of contexts. Panoramic views of Jupiter and Saturn have been assembled for multiple images returned to Earth from the two Voyager spacecraft. In a similar way, Landsat photographs are routinely assembled into panoramic views of Earth. Detailed images of galaxies and nebulae have been assembled from mul-

The work reported in this paper was supported by NSF grant ECS-8206321. A shorter description of this work was published in the *Proceedings of SPIE, vol. 432, Applications of Digital Image Processing VI*, The International Society for Optical Engineering, Bellingham, Washington.

Authors' address: RCA David Sarnoff Research Center, Princeton, NJ 08540.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1983 ACM 0730-0301/83/1000-0217 \$00.75

ACM Transactions on Graphics, Vol. 2, No. 4, October 1983, Pages 217-236.




Fig. 1. A pair of images may be represented as a pair of surfaces above the (x, y) plane. The problem of image splining is to join these surfaces with a smooth seam, with as little distortion of each surface as possible.

tiple telescope photographs. In each of these cases, the mosaic technique is used to construct an image with a far larger field of view or level of detail than could be obtained with a single photograph. In advertising or computer graphics, the technique can be used to create synthetic images from possibly unrelated components.

A technical problem common to all applications of photomosaics is joining two images so that the edge between them is not visible. Even slight differences in image gray level across an extended boundary can make that boundary quite noticeable. Unfortunately, such gray level differences are frequently unavoidable; they may be due to such factors as differences in camera position or in image processing prior to assembly. Thus, a technique is required which will modify image gray levels in the vicinity of a boundary to obtain a smooth transition between images. The two images to be joined may be considered as two surfaces, where the image intensity $I(x, y)$ corresponds to the elevation above the x, y plane. The problem, as illustrated in Figure 1, may be stated as follows: How can the two surfaces be gently distorted so that they can be joined together with a smooth seam? We will use the term *image spline* to refer to digital techniques for making these adjustments. A good image spline will make the seam perfectly smooth, yet will preserve as much of the original image information as possible.

It is probably safe to say that no fully satisfactory splining technique has yet been found. Most image mosaics are still produced without any attempt at remov-




Fig. 2. The weighted average method may be used to avoid seams when mosaics are constructed from overlapped images. Each image is multiplied by a weighting function which decreases monotonically across its border; the resulting images are then summed to form the mosaic. Example weighting functions are shown here in one dimension. The width of the transition zone T is a critical parameter for this method.

ing visible boundaries (e.g., [4]). The magnitude of the gray level difference across a mosaic boundary can be reduced to some extent by a judicious choice of boundary location when splining overlapped images. The match may be improved by adding a linear ramp to pixel values on either side of the boundary to obtain equal values at the boundary itself [6, 7]. A still smoother transition can be obtained using a technique recently proposed by Peleg [9]: The "smoothest possible" correction function is constructed which can be added to each image of a mosaic to eliminate edge differences. However, this technique may not be practical for large images, since the correction functions must be computed using an iterative relaxation algorithm.

We are concerned with a weighted average splining technique. To begin, it is assumed that the images to be joined overlap so that it is possible to compute the gray level value of points within a transition zone as a weighted average of the corresponding points in each image. Suppose that one image, $Fl(i)$, is on the left and the other, $Fr(i)$, is on the right, and that the images are to be splined at a point \hat{i} (expressed in one dimension to simplify notation). Let $Hl(i)$ be a weighting function which decreases monotonically from left to right and let $Hr(i) = 1 - Hl(i)$ (see Figure 2). Then, the splined image F is given by

$$F(i) = Hl(i - \hat{i}) Fl(i) + Hr(i - \hat{i}) Fr(i).$$

It is clear that with an appropriate choice of H , the weighted average technique will result in a transition which is smooth. However, this alone does not ensure that the location of the boundary will be invisible. Let T be the width of a transition zone over which Hl changes from 1 to 0. If T is small compared to image features, then the boundary may still appear as a step in image gray level, albeit a somewhat blurred step. If, on the other hand, T is large compared to image features, features from both images may appear superimposed within the transition zone, as in a photographic double exposure.

ACM Transactions on Graphics, Vol. 2, No. 4, October 1983.

These extremes are illustrated in Figure 3 with several attempts to spline two synthetic images of stars. The original images, Figures 3a and 3b (257 x 257 pixels) are identical except for a slight shift in vertical position and a slight shift in mean gray level. The first of these differences can arise from optical distortions or misalignments of actual photographs, while the second can be due to differences in atmospheric conditions or in photographic development.

In this example, photomosaics are obtained by joining the left half of Figure 3a with the right half of Figure 3b. If this is done without any attempt to smooth the image transition ($T = 0$), the boundary will appear as a sharp edge (Figure 3c). If instead the images are combined by the method of weighted average within a relatively narrow transition zone ($T = 8$), the edge appears blurred but remains visible (Figure 3d). When the images are splined with a broad transition ($T = 64$), the edge is no longer visible but stars have a "double exposed" look within the transition zone (Figure 3e).

Clearly, the size of the transition zone, relative to the size of image features, plays a critical role in image splining. To eliminate a visible edge the transition width should be at least comparable in size to the largest prominent features in the image. On the other hand, to avoid a double exposure effect, the zone should not be much larger than the smallest prominent image features. There is no choice of T which satisfies both requirements in the star images of Figure 3 because these contain both a diffuse background and small bright stars.

These constraints can be stated more precisely in terms of the image spatial frequency content. In particular, a suitable T can only be selected if the images to be splined occupy a relatively narrow spatial frequency band. As a rough requirement, we may stipulate that T should be comparable in size to the wavelength of the lowest prominent frequency in the image. If T is smaller than this the spline will introduce a noticeable edge. On the other hand, to avoid a double exposure effect, T should not be much larger than two wave lengths of the highest prominent frequency component in the images. This ensures that there will not be room for multiple features within the transition zone. While it is likely that these limits can be exceeded somewhat without noticeable degradation, the general conclusion—that the band width of images to be splined should be roughly one octave—is an important one.

How can images which occupy more than an octave be splined? The approach proposed here is that such images should first be decomposed into a set of band-pass component images. A separate spline with an appropriately selected T can then be performed in each band. Finally, the splined band-pass components are recombined into the desired mosaic image. We call this approach the multi-resolution spline. It was used to obtain the image shown in Figure 3f.

In decomposing the image into frequency bands, it is important that the range of frequencies in the original be covered uniformly, although the bands themselves may overlap. As a practical matter, a set of low-pass filters are applied to generate a sequence of images in which the band limit is reduced from image to image in one-octave steps. Band-pass images can then be obtained simply by subtracting each low-pass image from the previous image in the sequence. This not only ensures complete coverage of spatial frequencies but means that the final mosaic can be obtained simply by summing the band-pass component images.

ACM Transactions on Graphics, Vol. 2, No. 4, October 1983.




Fig. 3. Common artifacts of the weighted average techniques are demonstrated in these attempts to spline two synthetic images of stars (Figure 3a and 3b). These differ only in mean gray level and a slight vertical shift. A seam is clearly visible when the left half of figure 3a is joined with the right half of Figure 3b without any adjustment in gray level, as shown in Figure 3c. The seam is still visible when the weighted average technique is used with a narrow transition zone (Figure 3d). However, if the transition zone is wide, features within the zone appear double (Figure 3e). The first of these artifacts is due to a gray level mismatch at low spatial frequencies, while the second is due to a position mismatch at high frequencies. Both are avoided in the multiresolution method (Figure 3f).

ACM Transactions on Graphics, Vol. 2, No. 4, October 1983.

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.