

SEAMLESS STITCHING OF IMAGES BASED ON A HAAR WAVELET 2D INTEGRATION METHOD

Ioana S. Sevcenco, Peter J. Hampton and Pan Agathoklis

Department of Electrical and Computer Engineering
University of Victoria
Victoria, BC, V8W 3P6, Canada

ABSTRACT

A method for seamless stitching of images with photometric inconsistencies in the overlapping region is presented. The method is based on generating a set of stitched gradients from the gradients of the input images and then reconstructing the mosaic image using the Haar wavelet integration technique of [11], [12]. This $O(N)$ reconstruction technique is based on obtaining the Haar wavelet decomposition of the mosaic image directly from the stitched gradients and then using Haar synthesis to obtain the mosaic image. The Haar synthesis step includes a Poisson smoother at each resolution level leading to results without visual artifacts despite the non conservative nature of the stitched gradient field. Experimental results illustrate the method and show that it leads to seamless mosaic images despite intensity differences in the overlap region of the input images.

Index Terms— Image mosaicking, image reconstruction from gradient measurements, Poisson equation.

1. INTRODUCTION

There are many applications where large view field images are of great importance. In areas ranging from medical imaging to computer graphics and satellite imagery, a computationally efficient and easy to implement method to produce high resolution wide angle images will continue to draw research interest. One of the most common methods is image mosaicking, which consists of combining multiple overlapping representations of a scene into a single, wide angle image. Depending on the scene content, on the camera position, and on the intended application, two major stages in mosaicking algorithms are identified and consecutively addressed. The first is the registration of the images to be combined and the second deals with the inconsistencies of luminance and chrominance within the overlapping regions.

In image registration the source images are spatially aligned. This is typically [10] done by choosing one of the images as the reference and then finding the geometric

transformations which map the other images onto the coordinate system of the reference image. Upon completion of this step an initial mosaic is created, which represents an expanded view consisting of the union of all component images. Various intensity and feature based registration algorithms have been developed and are presented in [5] and in [22, 4] criteria for quality evaluation of image mosaics are reviewed.

This initial mosaic may contain visual artifacts because of the different lighting conditions associated with each source image. The second stage is finding a blending function to combine the overlapping regions of the source images, in such a way that in the final mosaic the transition from one image to the other becomes imperceptible. Good blending techniques should produce seamless mosaics by compensating for existing exposure differences in component images, without introducing blurring or ghosting artifacts. Previously proposed image blending algorithms operate on image intensities or image gradient, at full resolution of the image or at multiple resolution scales. A well known intensity based technique is *feathering* [18], where the mosaic is generated by computing a weighted average of the source images. In the composite mosaic image, pixels are assigned weights proportional to their distance from the centre of the image they come from, resulting in a smoother transition from one image to the other in the final mosaic. As it is discussed in [19], this method may generate blurring and ghosting artifacts when the images are not registered properly. The first intensity based multiscale method is presented in [6] and relies on *pyramid decomposition*. The source images are decomposed into bandpass components and the blending is done at each scale, in a transition zone inversely proportional to the spatial frequency content in the band. A similar approach in [13] replaces pyramid decomposition with wavelet decomposition.

Image gradient blending techniques prove effective because the human visual system is known to be more sensitive to local contrast changes than to global intensity variations. Therefore, the gradient domain provides an excellent framework for image processing applications such

as image editing [17, 2, 8], high dynamic range imaging and compression [7, 21] and image mosaicking [16, 19, 20, 3]. In all these methods the gradients of the source images are modified and the modified gradients are used to obtain the final image. In general, these gradient modifications lead to a gradient which typically is not a conservative vector field, and image reconstruction from this gradient no longer has an exact solution, but can be formulated as an optimization problem. A general framework for image reconstruction from non-integrable gradients is discussed in [1].

Two main approaches to gradient domain image mosaicking are the Gradient-domain Image Stitching algorithm [16], and the optimal seam techniques [15, 3]. Gradient-domain Image Stitching, as described in [16], blends the gradients of the source images in the overlap region and the image is then reconstructed from this gradient data set by an FFT based technique. For two given registered images, optimal seam techniques search the overlap regions for a curve along which the difference between the images is minimum. Once such a curve is found, the mosaic is constructed by pasting the image parts on the corresponding sides of the curve. The method proposed in [3] is based on an optimal seam approach in the intensity and gradient domains. The image is reconstructed using a conjugate gradients method.

In this paper a new method for seamless mosaicking of images which are assumed to be registered is presented. The proposed method is based on generating a set of gradients for the mosaic image by blending the gradients of the source images in the overlap region and pasting the gradients from the source images in the rest. The image mosaic will be reconstructed from this gradient data set. The image reconstruction from the gradient will be done here based on the approach presented in [11] and [12] which relies on the Haar wavelet decomposition. The basic idea is that the Haar wavelet decomposition of the resulting image can be directly computed from the gradient data and the image can then be recovered from this decomposition using Haar synthesis. The computational complexity of this method is $O(N)$, where N is the number of pixels and, as it was discussed in [11], this is faster than FFT based methods. Further, this reconstruction method allows for denoising by processing the HH part of the Haar wavelet decomposition as discussed in [11]. In [12] this method was further extended to include a Poisson smoother at each level of resolution during the Haar synthesis step. The addition of the Poisson smoother has the effect that the final reconstruction satisfies the Poisson equation at full resolution and leads to a result with no visual artifacts in the case of a non conservative gradient field. This improvement in the quality of the resulting image comes at a very low computational cost since there is no need for more than one or two iterations of the Poisson smoother at each resolution level. Since the gradient generated by the blending of the gradients of the input images is expected to

be a non conservative field, this approach will be used here for obtaining the mosaic image from the stitched gradients.

The remainder of this paper is structured as follows. In Section 2 the proposed method is presented. Its performance and effectiveness are shown through experimental results in Section 3, and in Section 4 the key concepts and main results are reviewed.

2. PROPOSED METHOD

The method for seamless image stitching proposed here is based on combining the gradients of the two input images to generate a set of gradients for the mosaic image and then reconstructing the mosaic image from these gradients.

The proposed method is sketched in Fig. 1. Consider two registered input images Φ_1 and Φ_2 with a common overlapping region as indicated in Fig. 1. The gradients

$$\tilde{\Phi}_1 = \left[\frac{\partial \Phi_1}{\partial x}, \frac{\partial \Phi_1}{\partial y} \right]^T \quad \text{and} \quad \tilde{\Phi}_2 = \left[\frac{\partial \Phi_2}{\partial x}, \frac{\partial \Phi_2}{\partial y} \right]^T$$

of the input images are computed and stitched together to generate the gradient from which the mosaic image will be reconstructed. In the overlapping region, the gradients are blended using a weighting function, while in the rest the gradients are simply stitched together.

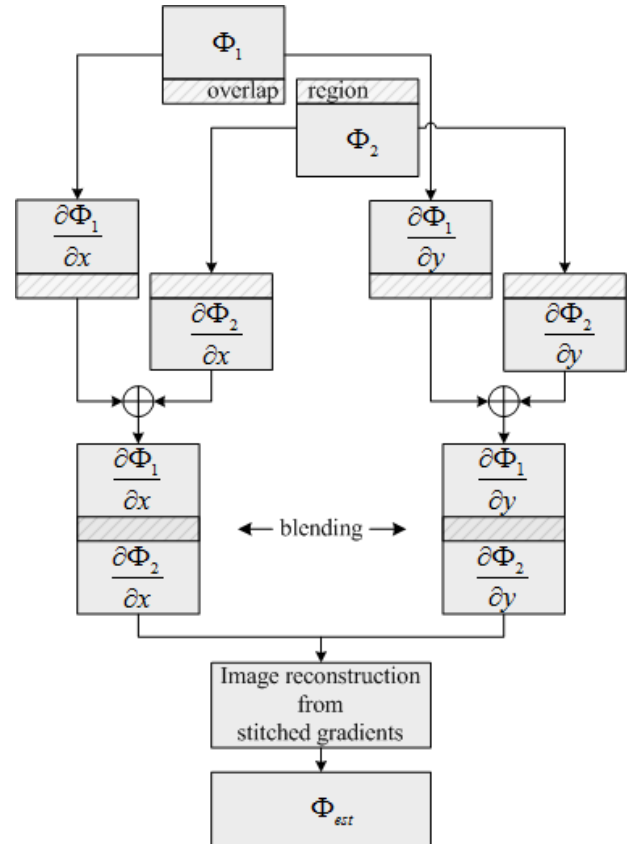


Fig. 1. Overview of proposed method for seamless stitching of images.

There have been many algorithms proposed in literature for the blending in the overlap region such as feathering [18], pyramid blending [6], etc. In our approach the blending of gradients in the overlap region is carried out using:

$$\tilde{\omega} = w(x, y) \cdot \tilde{\omega}_1 + [1 - w(x, y)] \cdot \tilde{\omega}_2 \quad (1)$$

where $\tilde{\omega}_1$ and $\tilde{\omega}_2$ denote the gradients of Φ_1 and Φ_2 in the overlap region respectively and $w: \mathbf{R}^2 \rightarrow [0, 1]$ is a linearly or exponentially decreasing function. This simple approach results in the gradient from the upper image dominating the upper part of the overlapping region while the lower gradient dominates the lower part. The gradient of the mosaic image outside the overlapping region is obtained by pasting the gradients from the non-overlapping parts of the input images.

The mosaic image will now be generated from the stitched gradients. The two dimensional (2D) integration method for reconstructing the mosaic image is based on the one presented in [11] and [12]. In this method, the Haar wavelet decomposition of the image is obtained directly from the image gradient and then the image is reconstructed using Haar synthesis. This technique has computational complexity $O(N)$, where N is the number of pixels and, as shown in [11] leads to fast and accurate results when the gradient field is a conservative field. However, for stitching of images using their gradients, the stitched gradient is not expected to be a conservative field in general and this technique may lead to a mosaic image with artifacts. It is well known that image reconstruction from its gradient can be formulated as an ℓ_2 optimization problem whose solution can be obtained by solving the Poisson equation:

$$\Delta \Phi = \frac{\partial}{\partial x} \tilde{\Phi}_x + \frac{\partial}{\partial y} \tilde{\Phi}_y \quad (2)$$

where $\tilde{\Phi}_x$ and $\tilde{\Phi}_y$ are the given input gradients (stitched gradients), $\Phi: \mathbf{R}^2 \rightarrow \mathbf{R}$ is the unknown mosaic image, and

$$\Delta \Phi = \frac{\partial^2 \Phi}{\partial x^2} + \frac{\partial^2 \Phi}{\partial y^2}$$

is the Laplacian operator.

Equation (2) can be discretized using an approximation of the continuous gradients. Two of the commonly used 2D approximations are based on the Hudgin [14] and the Fried [9] discretization models. In the case of Hudgin geometry the vertical and horizontal derivatives are approximated between two vertical and two horizontal points respectively. In Fried geometry, a gradient value is approximated in the center of 4 points. In image processing and vision applications the Hudgin geometry is commonly used, while in adaptive optics the Fried geometry is common because this geometry is consistent with the model of wavefront sensors. The transformation from Hudgin to Fried model is straightforward and Fried geometry is being used here.

Discretizing the Poisson equation (2) using the Fried geometry leads to:

$$\begin{bmatrix} 1 & 0 & 1 \\ 0 & -4 & 0 \\ 1 & 0 & 1 \end{bmatrix} \otimes \Phi = \begin{bmatrix} 1 & -1 \\ 1 & -1 \end{bmatrix} \otimes \tilde{\Phi}_x + \begin{bmatrix} 1 & 1 \\ -1 & -1 \end{bmatrix} \otimes \tilde{\Phi}_y \quad (3)$$

where \otimes denotes filter convolution. An iterative solution to equation (3) can be obtained using the Jacobi method [23] leading to the following Poisson smoother:

$$\Phi_{k+1} = \Phi_k - 0.25 \left(\begin{bmatrix} -1 & 0 & -1 \\ 0 & 4 & 0 \\ -1 & 0 & -1 \end{bmatrix} \otimes \Phi_k + \begin{bmatrix} 1 & -1 \\ 1 & -1 \end{bmatrix} \otimes \tilde{\Phi}_x + \begin{bmatrix} 1 & 1 \\ -1 & -1 \end{bmatrix} \otimes \tilde{\Phi}_y \right) \quad (4)$$

In [12] the Haar wavelet reconstruction technique was further developed to obtain the solution to the Poisson equation (2) and thus minimize the Frobenius norm of the difference between the Laplacian and the divergence of the gradient of the reconstructed image. This was done by modifying the Haar wavelet synthesis step to include a Poisson smoother, eq. (4), at each resolution level. By using the Poisson smoother at each resolution, only one or two iterations are required at each level, which is much less than the number of iterations necessary when the Poisson smoother is applied on the full resolution image.

The method proposed here can thus be summarized as follows:

1. Compute the gradients of the input images:

$$\tilde{\Phi}_1 = \left[\frac{\partial \Phi_1}{\partial x}, \frac{\partial \Phi_1}{\partial y} \right]^T \quad \text{and} \quad \tilde{\Phi}_2 = \left[\frac{\partial \Phi_2}{\partial x}, \frac{\partial \Phi_2}{\partial y} \right]^T$$

2. Create the stitched gradients using the blending function given by eq. (1) in the overlap region and ‘‘cut and paste’’ in the rest to obtain:

$$\tilde{\Phi} = \left[\tilde{\Phi}_x, \tilde{\Phi}_y \right]^T$$

3. Reconstruct the mosaic image from $\tilde{\Phi}$ using the Haar wavelet reconstruction technique of [12] with Poisson smoother at each resolution level.

The performance of the proposed method will be illustrated and discussed in the next section.

3. EXPERIMENTAL RESULTS

In this section the performance of the proposed method is illustrated with experiments. Given two partially overlapping, previously registered images with photometric inconsistencies (i.e., with differences in light intensity within the overlap region) the objective is to stitch them and produce a larger mosaic image which looks smooth and continuous without any noticeable artifacts such as seams or blurring.



Fig. 2. Left to right: (a) Input images; (b) Direct pasting of intensities; (c) Result obtained using the proposed method. Mosaic image: 512×512 pixels, overlap region: 30 pixels.

The input images are constructed from standard test images, by cropping them and modifying the light intensity values to generate photometric inconsistency. Fig. 2a shows an example of the input images used. The two images have an overlapping region indicated by the two horizontal lines. The intensity levels in the two images have been modified and are clearly different: the average intensity of the top image becomes lighter from bottom to top, while the intensity in the lower image becomes darker from left to right. This has the effect that in a “direct paste” mosaic image the seam is less visible where the two image intensities are close (left hand-side of the image shown in Fig. 2b, along the central horizontal line) and more visible where the difference in intensity is larger (right hand-side of Fig. 2b). The result of stitching the two images using the method proposed in the previous section based on [12] can be seen in Fig. 2c. Clearly there is no seam and the image has no visual artifacts.

To illustrate the effect of Poisson smoothing at each resolution level during the Haar synthesis step of the image reconstruction from the stitched gradient, consider Fig. 3. Reconstructing the image from the stitched gradient using the Haar wavelet approach [11] leads to blocking artifacts as

shown in Fig. 3a. These blocking artifacts are due to the fact that the stitched gradient is not a conservative field and therefore there is a difference between the gradient of the reconstructed image and the stitched gradient. This difference can be reduced, and the visual appearance of the image improved, by using the reconstructed image in Fig. 3a as the input to an iterative Poisson smoother at full resolution given by equation (4). This leads, after several iterations, to Fig. 3b with reduced, but still visible, blocking artifacts. Using the method of [12] which employs a Poisson smoother at each resolution level during the Haar synthesis step, as proposed in the previous section, produces the image in Fig. 3c (same as Fig. 2c). It can be seen that there are no blocking artifacts in this image. To evaluate the performance of the proposed method for various levels of photometric inconsistency, the amount of intensity modification in the two input images was varied using a parameter k . For high values of k the intensity of the top image is becoming lighter while at the same time the bottom image is becoming darker. The difference between the stitched gradient and the gradient of the mosaic image reconstructed from the stitched gradient was evaluated for various values of k . The ℓ_2 norm of the difference is plotted



Fig. 3. Left to right: (a) Reconstruction from the stitched gradient; (b) Poisson smoothing at full resolution after reconstruction from the stitched gradient; (c) Poisson smoothing at each resolution during reconstruction.

in Fig. 4. It can be seen that the use of the Poisson solver at each resolution level leads to the smallest difference among the three options. In our simulations various sizes of the overlap region led to comparable results. For a 512×512 mosaic image, the smallest overlap region we considered was 6 pixels wide. The results are still acceptable and an example is shown in Fig. 5.

In the above experiments the assumption was made that the images are registered. An interesting question is how possible registration errors between input images affect the quality of the mosaic image. Such errors may cause the appearance of ghosting artifacts in some methods. Fig. 6 shows the stitching results when there is a misregistration between the input images of 10 pixels in both directions. The registration error can be observed in Fig. 6b. The stitched image using the proposed method, Fig. 6c, indicates that the proposed method can yield satisfactory results even in the presence of registration errors.

The above experiments illustrate that the proposed method can lead to mosaic images without any visual artifacts despite differences in the intensities of the input images in the overlapping region. The algorithm was implemented using Matlab and was running on a x86 32bit PC Architecture, Intel Core Duo T2080, 1.73GHz. The av-

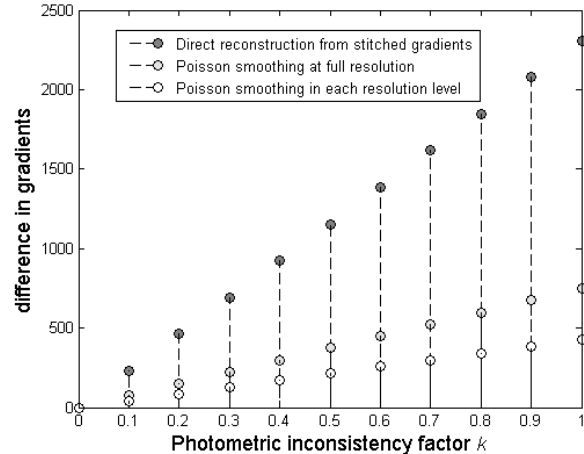


Fig. 4. ℓ_2 norm of the difference between the stitched gradient and the gradient of the mosaic image. Mosaic size: 512×512 ; overlap region height: 100 pixels.

erage computation time for a 512×512 mosaic was 0.5 sec, while for 1024×1024 mosaic images the average time was 2.05 sec. These average computation times are consistent with the $O(N)$ complexity of the Haar wavelet reconstruction method as discussed in [11] and [12].



Fig. 5. From left to right: (a) Input images; (b) Direct pasting of intensities; (c) Result obtained using the proposed method. Mosaic image: 512×512 pixels; overlap region height: 6 pixels.




Fig. 6. From left to right: (a) Input images superimposed on the overlap region; (b) Direct pasting of intensities, showing registration error; (c) Result obtained using the proposed method.



Fig. 7. Performance of proposed method on other types of illumination inconsistencies. (a) Direct pasting intensities; (b) Result of the proposed method.

Fig. 7 illustrates the performance of the proposed method with other types of illumination inconsistencies. In Fig. 7, the intensity of the upper image was obtained by gradually making the image darker towards its borders.

4. CONCLUSIONS

A method for seamless stitching of images with photometric inconsistencies in the overlapping region is presented. The method is based on generating a set of stitched gradients from the gradients of the input images and then reconstructing the mosaic image using the Haar wavelet based reconstruction technique of [11], [12]. This $O(N)$ technique includes a Poisson solver at each resolution level leading to mosaic images without visual artifacts despite the non conservative nature of the stitched gradient field. Experimental results illustrate the method and show that it can lead to seamless mosaic images despite intensity differences in the overlap region of the input images.

REFERENCES

- [1] A. Agrawal, R. Raskar, and R. Chellappa, "What is the range of surface reconstructions from a gradient field?," in A. Leonardis, H. Bischof, and A. Pinz, editors, *Computer Vision a ECCV 2006*, vol. 3951 of Lecture Notes in Computer Science, pp 578–591, Springer Berlin, Heidelberg, 2006.
- [2] A. Agrawal, R. Raskar, S. Nayar, and Y. Li, "Removing Photography Artifacts using Gradient Projection and Flash-Exposure Sampling," *ACM Trans. Graphics* vol. 24, no. 3, pp. 828–835, Jul. 2005.
- [3] A. Agarwala, M. Dontcheva, M. Agrawala, S. Drucker, A. Colburn, and B. Curless, "Interactive Digital Photomontage," *ACM Trans. Graphics*, vol. 23, no. 3, pp. 294–302, Aug. 2004.
- [4] J. Boutellier, O. Silvén, M. Tico, and L. Korhonen, "Objective evaluation of image mosaics," book chapter in *Computer Vision and Computer Graphics. Theory and Applications*, J. Braz et al. (Eds.), Springer Berlin, Heidelberg, pp. 107–117, 2009.
- [5] L. Brown, "A survey of image registration techniques," *ACM Computing Surveys (CSUR)*, vol. 24, no. 4, pp. 325–376, Dec. 1992.
- [6] P. J. Burt, and E.H. Adelson, "A multiresolution spline with application to image mosaics," *ACM Trans. Graphics*, vol. 2, pp. 217–236, Oct. 1983.
- [7] R. Fattal, D. Lischinski, and M. Werman, "Gradient domain high dynamic range compression," *Proceedings of SIGGRAPH*, pp. 249–256, 2002.
- [8] G. D. Finlayson, S. D. Hordley, and M. S. Drew, "Removing shadows from images," *Proc. European Conf. Computer Vision*, pp. 823–836, 2002.
- [9] D. L. Fried, "Least square fitting a wavefront distortion estimate to an array of phase difference measurements," *Journal of the Optical Society of America*, vol. 67, pp. 370–374, 1977.
- [10] A., Goshtasby, *2-D and 3-D image registration: for medical, remote sensing, and industrial applications*, John Wiley & Sons, Inc., Hoboken, New Jersey, 2005.
- [11] P. Hampton, P. Agathoklis, and C. Bradley, "A new wavefront reconstruction method for adaptive optics system using wavelets," *IEEE Journal of Selected Topics in Signal Processing* 2, pp. 781–792, Oct. 2008.
- [12] P. Hampton, and P. Agathoklis, "Comparison of Haar wavelet-based and Poisson-based numerical integration techniques," *Proceedings of 2010 IEEE International Symposium Circuits and Systems (ISCAS)*, pp. 1623–1626, 2010.
- [13] C. T. Hsu, and J. L. Wu, "Multiresolution mosaic," *IEEE Trans. on Consumer Electronics*, vol. 42, pp. 981–990, Aug. 1996.
- [14] R. Hudgin, "Wave-front reconstruction for compensated imaging," *Journal of the Optical Society of America*, vol. 67, pp. 375–378, 1977.
- [15] J. Jia, and C.-K. Tang, "Eliminating structure and intensity misalignment in image stitching," *ICCV 2005. Tenth IEEE International Conference on Computer Vision*, vol.2, pp.1651–1658, Oct. 2005.
- [16] A. Levin, A. Zomet, S. Peleg, and Y. Weiss, "Seamless image stitching in the gradient domain," *Eighth European Conference on Computer Vision*, vol. 4, pp. 377–389, May 2004.
- [17] P. Perez, M. Gangnet, and A. Blake, "Poisson image editing," *Proceedings of SIGGRAPH*, pp. 313–318, Jul. 2003.
- [18] R. Szeliski, "Video mosaics for virtual environments," *IEEE Computer Graphics and Applications*, pp. 22–30, Mar. 1996.
- [19] R. Szeliski, "Image alignment and stitching: a tutorial," Technical Report MSR-TR-2004-92, Microsoft Research, Dec. 2004.
- [20] M.-S. Su, W.-L. Hwang, and K.-Y. Cheng, "Analysis on multiresolution mosaic images," *IEEE Trans. on Image Processing*, vol. 13, no. 7, pp.952–959, Jul. 2004.
- [21] J. Tumblin, A. Agarwal, and A. Raskar, "Why I want a gradient camera," *Proc. IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR 2005)*, vol. 1, pp. 103–110, 2005.
- [22] Z. Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Trans. on Image Processing*, vol.13, no.4, pp. 600–612, Apr. 2004.
- [23] D. S., Watkins, *Fundamentals of Matrix Computations*, John Wiley & Sons, Inc., New York, USA, 2002.