DICTA2002: Digital Image Computing Techniques and Applications, 21--22 January 2002, Melbourne, Australia.

Application of Optical Flow Techniques in the Restoration of Non-Uniformly Warped Images

David Clyde[1], Ian Scott-Fleming[2], Donald Fraser[1], Andrew Lambert[1]
[1]*School of Electrical Engineering, The University of New South Wales, ADFA, Northcott Drive, Canberra ACT 2600, Australia*
d-clyde@defence.gov.au, d-fraser@adfa.edu.au, a-lambert@adfa.edu.au
[2]*Department of Electrical and Computer Engineering, The University of Arizona, Tucson, AZ 85721, USA*
icsf@ece.arizona.edu

Abstract

When viewing a scene through a turbulent atmosphere, received images will be non-uniformly distorted. A major component of this distortion is a random x and y shift at each point in the received image. A motion blurred but geometrically accurate prototype is obtained by averaging a large number of individual frames. Previous efforts to remove the motion blur utilised cross-correlation techniques to estimate a warping function for each individual frame, allowing the random shifts to be removed. We are currently investigating the implementation of an alternative method where gradient techniques are used to estimate the warping function. Results obtained from both simulated and real data are presented. Current activities are refining the implementation of the gradient techniques and evaluating the performance of the gradient based approach against the cross-correlation approach.

1. Introduction

There exist a number of situations where a sequence of images will be subjected to non-uniform, random distortions, for example, in the fields of terrestrial and wide-area astronomical surveillance. When a scene is viewed through the atmosphere, it is in reality being viewed through a boiling mass of different "atmospheric cells", resulting in a non-uniform distortion of the received image. A major component of this distortion is a random x and y shift at each point in the received image. The point-spread function for each image approximates a position-dependent, randomly displaced delta function. A motion blurred but geometrically accurate prototype can be obtained by averaging a large number of individual frames. Previous efforts have attempted to estimate a warping function for each individual frame via an iterative cross-correlation technique between points in the individual frame and the prototype [1-5]. The individual frames are restored using the estimated warping functions and, by averaging the restored images, an updated prototype is generated. The updated prototype provides a geometrically accurate representation of the scene, with a reduction in the motion blur. The residual blurring due to instantaneous speckle or instrument blur has been ensemble averaged in the final, motion-blur-restored, prototype [2], and can be removed using standard deconvolution techniques. The results of this work, previously published in [2], show a significant improvement in the quality of the restored image over the original.

The process of restoring the individual frames can be viewed as an attempt to estimate the optical flow in an image sequence where the prototype is the first image in the series and the individual frame is the subsequent image. In an effort to reduce the computational effort and time required to estimate the warping function, and improve the accuracy, we are investigating an alternative method where gradient techniques are used to estimate the optical flow rather than the cross-correlation techniques previously implemented. In this approach, the prototype and the individual frames are decomposed into a multi-resolution stack of images. The shifts are estimated for each point in the lowest resolution level of the stack producing a dense vector field estimate of the shifts. This is considered to be the initial estimate of the optical flow. This estimate is used to warp the next level in the multi-resolution stack and the "restored" image is used to calculate a "refinement" vector field that is added to the current optical flow estimate to create a refined estimate.

The process repeats until all the layers except the highest resolution in the multi-resolution stack have been processed and a completed estimate of the warping function obtained. The motion blur present in the prototype at the highest resolution introduces a significant amount of noise into the calculation and degrades the result if the highest resolution is included in the process.

This implementation of the gradient techniques is similar to implementations used in other applications involving these same techniques. The primary difference in this instance arises from the noise that is introduced by the motion blur in the prototype. Ideally, in attempting to map the shifts between the geometrically correct prototype and an individual frame, the prototype would contain no noise or blurring. Unfortunately, as a motion blur removed prototype is the result we are hoping to obtain, such a prototype is unavailable for use in the calculation. Instead we use the motion blurred average of a large number of frames that have been individually warped. The motion blur introduces additional noise into the process and, due to the instability of gradient techniques in the presence of noise, this causes problems in the accurate estimation of the warping functions.

To date, the gradient techniques have been evaluated primarily with simulated data, although initial results obtained from the de-warping of real turbulence-degraded data are also presented. Current activities are refining the implementation of the gradient techniques, particularly in regards to handling the noise introduced by the motion blur and reducing the time required for the estimation of the warping functions, and evaluating the performance of the gradient based approach against the cross-correlation approach.

2. Optical Flow

Optical flow is the estimation of the apparent motion in a series of time varying images. This apparent motion may be identical to the actual motion that is occurring, but not necessarily. One common example where the apparent motion and the actual motion are not identical is a rotating barber pole. In this case the actual motion is horizontal (i.e. rotating around the axis) and the apparent motion is vertical (i.e. up the pole).

Numerous techniques have been developed over the years for estimating optical flow but most can be grouped into one of four general types: correlation, gradient, spatiotemporal filtering, and Fourier phase or energy techniques. Correlation is the most common technique, but gradient implementations are often more efficient and can provide greater accuracy. The gradient method is capable of providing sub-pixel displacement estimates without first enlarging the images being operated on. However, the gradient techniques are only useful in situations where small displacements are expected [6].

All four of the general types rely on a basic assumption that there is a constant moving brightness pattern. In other words, the change in brightness at a particular point in an image is due solely to a shift in the pattern at that point. Other aspects such as the lighting and reflectance of the object are assumed to remain constant.

2.1. Estimation of Optical Flow via a Gradient Technique

Calculation of optical flow via a gradient technique captures the basic assumption of a constant moving brightness pattern in a common constraint known as Horn s Eq. [7]:

$$\frac{dI}{dt} = 0 \quad \Rightarrow \quad \frac{\partial I}{\partial x}\frac{dx}{dt} + \frac{\partial I}{\partial y}\frac{dy}{dt} + \frac{\partial I}{\partial t} = 0 \tag{1}$$

where I is the image intensity. Re-arranging Eq. (1) gives us what is known as the Image Brightness Constancy Eq. [8]:

$$(\nabla I) \cdot \mathbf{v} + I_t = 0 \tag{2}$$

where the subscript t indicates a partial derivative with respect to time. The values for the spatial and temporal derivatives required by Eq. (2) can be obtained by direct observation of a point, p, in an image sequence. However, due to the two unknowns in $\mathbf{v}$ (the velocity in the x direction and the velocity in the y direction), we are not able to find a solution for $\mathbf{v}$ based upon a single point in an image sequence. An additional constraint is required. By assuming that neighbouring points in an image will have experienced similar displacements and evaluating Eq. 2 over a patch , P, of N by N points centred on the point of interest, we are able to find a solution for $\mathbf{v}$ by minimising the total error. The total error over a patch is given by:

$$E(\mathbf{v}) = \sum_{p(x_i, y_i) \in P} \left[\left(\nabla I(x_i, y_i) \right) \cdot \mathbf{v} + I_t(x_i, y_i) \right]^2 \tag{3}$$

and we minimise the error using a least squares minimisation algorithm:

$$\mathbf{v} = \left(\mathbf{A}^T \mathbf{A} \right)^{-1} \mathbf{A}^T \mathbf{b} \tag{4}$$

where A and b are given by:

$$\mathbf{A}=\begin{bmatrix}\nabla I(x_1,y_1)\\ \nabla I(x_2,y_2)\\ \cdot\\ \cdot\\ \nabla I\left(x_{N^2},y_{N^2}\right)\end{bmatrix}\mathbf{b}=\begin{bmatrix}\frac{\partial I(x_1,y_1)}{\partial t}\\ \frac{\partial I(x_2,y_2)}{\partial t}\\ \cdot\\ \cdot\\ \frac{\partial I\left(x_{N^2},y_{N^2}\right)}{\partial t}\end{bmatrix}$$

2.2. Limiting Factors in the Gradient Techniques

Despite the advantages of the gradient techniques in calculating the optical flow, particularly in regards to accuracy, there are a number of inherent limitations. One of these is instability in the presence of noise, where noise is any factor that causes a divergence from the basic assumption of a constant brightness pattern. If we re-arrange Eq. (2) to obtain an expression for v we get:

$$\mathbf{v}=\frac{-I_t}{\nabla I} \tag{5}$$

Upon the introduction of noise, Eq. (5) becomes:

$$\mathbf{v}+\Delta\mathbf{v}=\frac{-I_t+n}{\nabla I} \tag{6}$$

where $\Delta\mathbf{v}$ is the error in the calculated vector, and n is the noise in the image sequence. Using equations (5) and (6), we can show that the error in the calculated vector is given by:

$$\Delta\mathbf{v}=\frac{n}{\nabla I} \tag{7}$$

and as the spatial gradient of an image goes to zero, the error in the calculated vector approaches infinity. It is clear that the best results will be obtained in areas of high spatial gradients.

In addition to the errors introduced by noise, it can also be shown that gradient techniques will fail in areas where the image motion is parallel to the image gradient over the entire patch. This failure is commonly referred to as the aperture problem [8,9], and arises as a result of the property of the techniques that only motion normal to the spatial gradient can be estimated. To correctly estimate the shift for a given point therefore, a multifaceted edge must be present within its neighbourhood.

In order to ensure that a multifaceted edge is present in the neighbourhood of a point, the neighbourhood can be enlarged, but this increases exposure to another difficulty. One of the assumptions made in solving Horn s equation over a patch is that all the points in the patch will have experienced a similar shift. The larger the patch, the greater the possibility that this assumption will be invalid. Any deviation from the assumption introduces noise into the process.

We can partially compensate for the variation in the shifts that occur over a patch by applying a weight to the error that is obtained when evaluating Eq. (2) at each point in the patch. These weights are used to emphasize errors in the centre of the patch, i.e. close to the point of interest, over those towards the edges of the patch. When weights are applied, a modified form of Eq. (3) gives the total error:

$$E_w(\mathbf{v})=\sum_{p(x_i,y_i)\in P} w_i\left[\left(\nabla I(x_i,y_i)\right)\cdot\mathbf{v}+I_t(x_i,y_i)\right]^2 \tag{8}$$

where w_i is the weight for a given location in the patch. The least squares minimisation of the error now becomes:

$$\mathbf{v}=\left(\mathbf{A}^T\mathbf{W}^2\mathbf{A}\right)^{-1}\mathbf{A}^T\mathbf{W}^2\mathbf{b} \tag{9}$$

where **W** is the weight matrix.

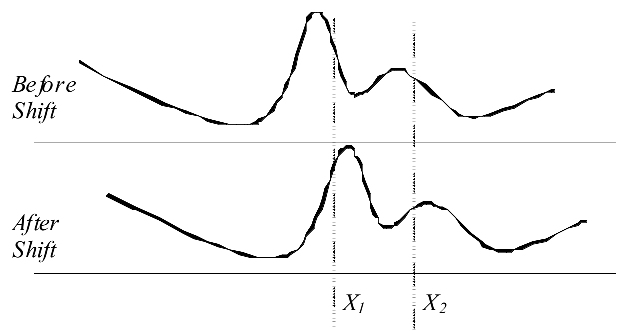

Fig. 1. These curves show potential problems that occur when using gradient techniques in the presence of relatively large scale shifts.

Another major limitation of the gradient techniques is their failure due to aliasing problems in the presence of large shifts. An example of the potential problems caused by large shifts is clearly represented in one dimension in Fig. 1. Here the second curve has been shifted to the right of the first. At X_1 the value of the curve has increased. The gradient of the curve at this point before the shift indicates that a small shift to the right had occurred. The gradient of the curve at the same location after the shift indicates that a small shift to the left had occurred. We might avoid the problem of which gradient to use by taking the average of the two. However, in this case, the average of the gradients is approximately zero, indicating a near infinite shift.

Alternatively, at X_2 in Fig. 1, the value of the curve at this point after the shift is identical to what it was before the shift. In this circumstance, the gradient techniques described will indicate that no shift has occurred.

The difficulties associated with large shifts are avoided if the change in slope over the extent of the shift is negligible.

The current implementation of the algorithm attempts to avoid the problems caused by large scale shifts by taking an iterative approach to computing the optical flow. The images involved in the computation are first decomposed into a multi-resolution stack and the optical flow is calculated on the coarsest image. The resultant vector field is used to warp the image in the next layer of the stack and the optical flow is calculated using the warped image. The resultant vector field is added to all previous vector fields as a refinement to the optical flow estimate. The process iterates until the optical flow has been computed at the second highest resolution. At the highest resolution, the noise introduced by the motion blurring in the prototype is too significant to allow any significant improvement on the vector field generated from the lower resolutions. In fact, including the highest resolution in the process is likely to be detrimental to the final result.

4. Results using Simulated Data

Initial investigations into the viability of the gradient techniques have been carried out using simulated data. The first stage of these investigations was to take a known image and distort it in a manner so as to simulate the effects of a turbulent atmosphere. Then, using the known image as the prototype, the warping function used to distort the known image was estimated using the gradient techniques. Example results from these initial investigations are shown in Fig. 2. Cross-sections of the original and recovered warping functions are shown in Fig. 3.

These results show that the gradient techniques have recovered a large proportion of the detail in the x and y shift maps, however it is obvious there is considerable room for improvement.

The second stage of the testing of the algorithm involved the generation of fifty "received" images by warping a source image according to randomly generated x and y shift maps. These received images were averaged to provide a prototype and then estimations of the warping functions were calculated between each image and the prototype. The received images were restored according to the calculated optical flow and the restored images were averaged to provide an updated prototype. Examples of the results obtained are shown in Fig. 4.

The results in Fig. 4 show that there has been considerable improvement in the quality of the image between the initial prototype and the updated prototype. Details in the hair, the feathers on the hat, and around the eyes in particular highlight the reduction of the motion blur in the updated prototype.

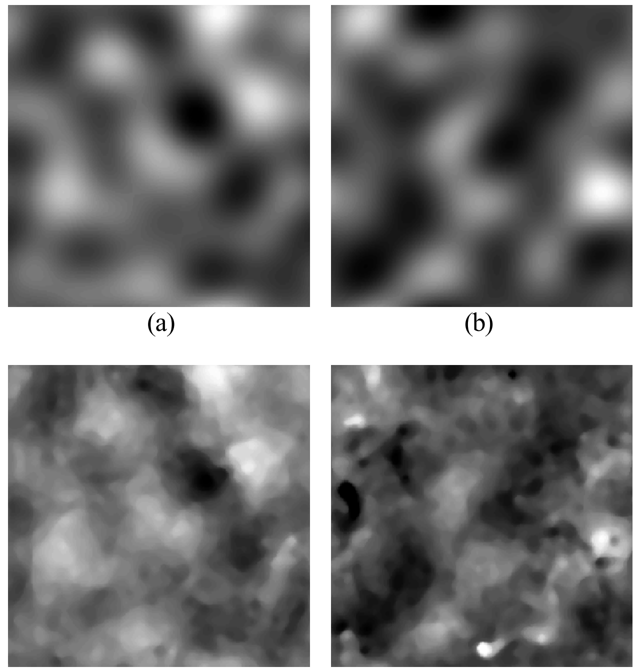

(a) (b) (c) (d)

Fig. 2. (a) The original *x* shift map (b) The original *y* shift map (c) Estimated *x* shift map (d) Estimated *y* shift map

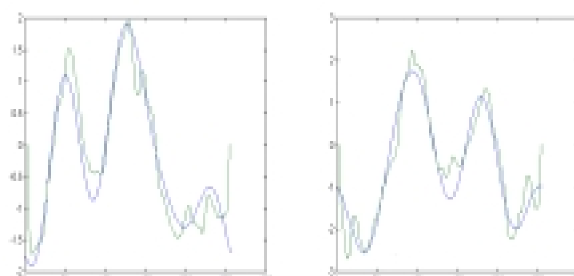

(a) (b)

Fig. 3. (a) Cross-section of the original *x* shift map (dark curve) and the estimated *x* shift map (light curve) (b) Cross-section of the original *y* shift map (dark curve) and the estimated *y* shift map (light curve)

(a)

(b)

Fig. 4. (a) Average of 50 independently warped images prior to processing (b) Average of 50 images after the motion blur has been removed.

5. Results with Real Data

Following the testing of the gradient approach with simulated data, the method was applied to sequences of real images captured with an interline-transfer CCD camera with a resolution of 1K by 1K. Example results are shown in Fig. 5.

To obtain these results, the source data was passed through the motion blur removal algorithm three times, with the results from the first iteration providing the prototype for the second iteration, and so on. The output of this process was then passed through a blind deconvolution process to remove the residual blurring due to instantaneous speckle effects [10]. The increased detail in the restored images is clearly visible.

(a)

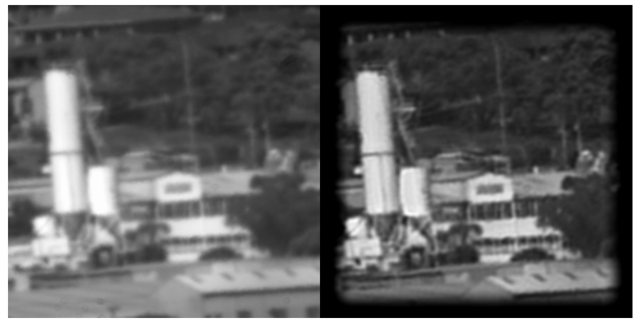

(b)

Fig. 5. (a) Average of the 'truck' image sequence. The left hand image is prior to processing and the right hand image is he final result. (b) Average of the 'industrial' image sequence. The left hand image is prior to processing and the right hand image is the final result.

One other result from the application of the gradient method to real data is shown in Fig. 6. This figure shows the estimate of the x and y shift maps in the area around the cabin of the truck in for one frame in the truck image sequence. These maps are similar to those obtained using the cross-correlation method, as shown in [2].

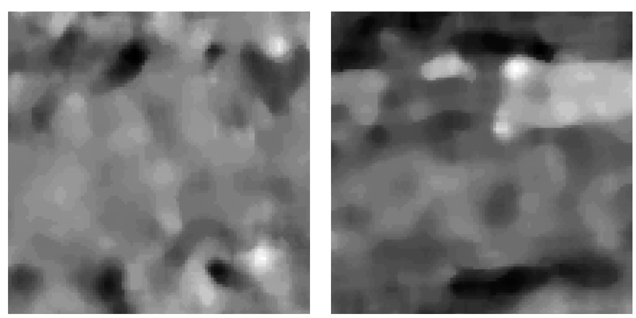

(a) (b)

Fig. 6. (a) Estimate of the *x* shift map for one frame in the truck image sequence. (b) Estimate of the *y* shift map for one frame in the truck image sequence.

6. Conclusions and Future Direction

The use of gradient techniques to estimate the warping functions required to restore images degraded by atmospheric turbulence has been shown to provide reasonable results for simulated data in cases where the warping is not overly severe. Problems arise when the magnitude of the shifts increases, in the presence of excessive noise, and when the nature of the spatial gradients in a neighbourhood does not allow an accurate estimation of the shift to be made. Measures exist in the current implementation of the gradient techniques to overcome some of the inherent limitations. Other measures that may enhance the restoration process are being trialed, and the outcomes of these trials will feed into development of the gradient method. In particular, two options are being investigated. Firstly, we are investigating the selection of points in the image that are likely to provide a good estimate of the shift at that location and interpolating between the selected points. Secondly, we are testing the performance of the method when an additional weighting matrix is applied in the error minimisation process to emphasise the errors at points of higher spatial gradient.

Results have been obtained using real data and the intention is to compare these results to those obtained using cross-correlation techniques. Initial impressions are that the results obtained using the current gradient method compare favourably with the results that were obtained using cross-correlation techniques.

7. References

[1] SPIE: D Fraser, and A.J. Lambert, " Wide area image restoration using a new iterative registration method", in Proc. *Conference 4123 Image Reconstruction from Incomplete Data*, SPIE 45th Annual Meeting, San Diego, July 2000.

[2] D. Fraser, G Thorpe, and A.J. Lambert, " Visualization of Turbulence and motion-blur removal in wide-area imaging through the atmosphere" *J. Opt. Soc. Am. A* 16, 1751-1758 (1999).

[3] OSA: D Fraser, G Thorpe, and A.J. Lambert, " Visualization of Turbulence and motion-blur removal in wide-area imaging through the atmosphere", in Proc., *Optical Soc. America, Summer Topical Meeting on Signal Recovery and Synthesis*, Kailua-Kona, Hawaii, June 9-11, 1998 (Optical Society of America Digest Series, Washington, D.C., 1998), pp. 16-19.

[4] SPIE: G. Thorpe, and D. Fraser, "Wide area imaging through the atmosphere", in Proc., *EUROPTO SERIES, Conference on Optics in Atmospheric Propagation*, Adaptive Systems, Adam D. Devir, Anton Kohnle, Christian Werner, eds., SPIE Proceedings Series, Vol. 2956 (1996), pp. 188-197.

[5] IAU: G. Thorpe, and D Fraser, "Restoration over fields of view wider than the isoplanatic patch", in Proc., *Very high angular resolution imaging*, J.G. Robertson, and W.J. Tango, eds., International Astronomical Union Symposium, Sydney, Australia, 11-15 January, 1993 (Kluwer Academic Publishers, Dordrecht, 1994), pp. 221-223.

[6] E.P. Simoncelli, " Bayesian Multi-scale Differential Optical Flow", in *Handbook of Computer Vision and Applications*, eds B Jahne, H Haussecker, and P Geissler, Volume II, chapter 14, pages 397-422. Academic Press, April 1999.

[7] B.K.P. Horn, and B. Schunck, "Determining optical flow", *Artif. Intell.* 17, 185-203 (1981).

[8] Trucco and Verri, *Introductory Techniques for 3-D Computer Vision*, (Prentice Hall, 1998), Chap. 8.

[9] E.P. Simoncelli, and E.H. Adelson, "Computing Optical Flow Distributions Using Spatio-temporal Filters", *MIT Media Laboratory Vision and Modeling Technical Report* #165, March 1991.

[10] N.J. Donaldson, "Image restoration through blind deconvolution", *Final Year Thesis*, School of Electrical Engineering, University College, UNSW, ADFA, 2001.