Robust Motion Estimation in the Presence of Fixed Pattern Noise

by

Andrew David Copeland

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2003

Author . . . .
Department of Electrical Engineering and Computer Science
May 21, 2003

Certified by. . . . .
Dennis M. Freeman
Associate Professor
Thesis Supervisor

Accepted by . . . .
Arthur C. Smith
Chairman, Department Committee on Graduate Students

Robust Motion Estimation in the Presence of Fixed Pattern Noise

by

Andrew David Copeland

Submitted to the Department of Electrical Engineering and Computer Science
on May 21, 2003, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

Motion estimation algorithms are useful in a variety of applications. These applications include motion characterization of Micro Electrical Mechanical Systems (MEMS) and cochlear mechanisms, video compression, and multi-frame image enhancement methods. One of the factors limiting accuracy in these measurements is a systematic bias towards a zero shift caused by fixed pattern noise (FPN) in the images. FPN is a spatial variation in the input/output relationship of each pixel that can be modeled as signal dependent multiplicative noise. These variations can be attributed to irregularities in pixel sizes and geometries, imperfections or impurities on the sensor surface and in the optical path, and non-uniform illumination. In this study the effect of FPN is examined by analyzing the correlation of the images. Taking the logarithm of the image transforms the noise into additive signal-independent noise that is then removed using conventional linear methods. Standard optical-flow algorithms are then used to measure the motions. These algorithms are performed on a series of test images in simulations. Over a wide range of FPN intensities, the measurements with the pre-processing produce more accurate results than those measurements without, and in some cases the average error is reduced by a factor of six. In addition to the simulations, the pre-filtering is also tested on real-world images of a moving target over a wide range of displacements. In these experiments the new method produces results with errors that were on average 12 dB smaller. The cost of this improvement is a 54% increase in the computational costs.

Thesis Supervisor: Dennis M. Freeman
Title: Associate Professor

Acknowledgments

First of all, I would like to acknowledge my advisor, Professor Denny Freeman. Working with Denny has been great. Whenever I think that I am busy, I just look at how busy he is and find that I might not be that busy after all. Denny has helped me most with his insightful questions, ability to figure out what is important to work on, and, most importantly, by pulling me back when I have gotten too far off track.

Working in the Micromechanics Group has been remarkable. It is truly an eclectic mix of people. My former office-mate and mentor Dr. A.J. Aranyosi has given me invaluable help. He has kept the computers running despite my best efforts, shown me the ropes in getting started on this project and participated in countless brainstorming sessions along the way. Salil Desai showed me how to use the Microvision system. Stan Hong, has shared with me a glimpse of his insight into optics. Jekwan Ryu and Amy Englehart have given me valuable advice about the nature of research.

This document was made possible by valuable input and a great deal of proof-reading by Stan, A.J., and Denny.

I'm especially grateful to my friends, the volleyball team, and the Protestant Student Community. They have provided me a great deal of support, possibly a few too many distractions, and some balance in my life.

Thank you to my parents and to my sister Samantha. Mom you have always kept me on task and helped me keep up with the details, Dad I got here because of a dream you helped me find long ago. Sam, you will always be my big sis.

Lastly, I would like to think the big man upstairs for my life and for the perspective he has given me. Sometimes, I think that life is a little bit like research... "Then you will know the truth, and the truth will set you free." John 8:32

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Image-based motion estimation algorithms are useful in a variety of applications. These applications include motion characterization of Micro Electrical Mechanical Systems (MEMS), motion characterization of biological mechanics, video compression, and multi-frame image enhancement methods. One of the limiting factors in motion estimation is the presence of noise in the images. Of these noises, fixed pattern noise is particularly undesirable because it introduces a systematic bias (Cain & Hayat 2001). The fixed pattern noise introduces an image component that remains perfectly registered from image to image, biasing the answer toward zero. Typically, the FPN has a variation of 1-6% across a CCD (Ilyin 2002). The greater the energy of the fixed pattern noise the greater the bias toward the zero shift (Davis & Freeman 1998). "It is therefore very desirable to have a registration algorithm that is tolerant to fixed-pattern noise. (Cain & Hayat 2001)"

In (Cain & Hayat 2001) it is argued that gradient-based algorithms are sensitive to fixed pattern noise because of their local nature. The method that they propose uses more global information, but can only provide estimates that are integer shifts. An alternate approach is to pre-process images before the application of the gradient algorithm to remove FPN, and thereby improve the accuracy of the measured shifts.

To remove the noise, a homomorphic filter that utilizes the logarithm operator in conjunction with an adaptive Wiener filter is used (Jain 1989). This filter transforms multiplicative noise into additive noise, adaptively filters it, then the filtered result is exponentiated (Campisi, Yan & Hatzinakos 2000, Hadhoud 1999, Lim 1990, Oppenheim, Shafer & Stockham, Jr. 1968). Although the algorithm is designed to combat fixed pattern noise, it also provides low pass filtering, which improves performance even in the presence of little or no fixed pattern noise.

To examine the performance of the motion estimation algorithms in conjunction with homomorphic filtering , both experimental tests and simulations were done. The simulations test a collection of general images over a wide range of noise levels and shifts. To verify the performance in practice, measurements of real motions were made over a range from one nanometer to one micrometer.

Chapter 2

Fixed Pattern Noise (FPN)

2.1 Cause of FPN

FPN is the spatial variation in the input/output relationship of each pixel of an image sensor. Some of the causes of FPN are irregularities in pixel size and geometry, imperfections or impurities on the sensor surface and in the optical path, and nonuniform illumination (Janesick 2001, Healey & Kondepudy 2001, Janesick 2002, Ilyin 2002).

2.2 Correlation Shows Effect of FPN

A good way to view the effect of FPN on motion estimation algorithms is to examine its effect on correlation, as follows. Let $s(x)$ represent the continuous brightness function of a scene. Similarly $s(x-\Delta)$ is that scene shifted by Δ. The corresponding discrete space signals, with the sampling frequency f_s above the Nyquist sampling rate, are $s[n] = s(nX)$ and $s'[n] = s(nX - \Delta)$ where $X = 1/f_s$. The imaging system introduces FPN, which is modeled as multiplicative noise. So the resulting images are $I_1[n] = \alpha[n]s[n]$ and $I_2[n] = \alpha[n]s'[n]$ where $\alpha[n]$ represents the FPN, which is unchanged between images. Let $\alpha[n]$ be a Gaussian random variable with mean μ_α and standard deviation σ_α. Each gain $\alpha[n]$ is uncorrelated with the other gains $\alpha[m]$ for $m \neq n$. Motion estimation can be performed by finding the peak of the cross correlation of the two images $I_1[n]$ and $I_2[n]$,

$$R(I_1[n], I_2[n]) = \sum_{<k>} I_1[k+n] I_2[k]. \tag{2.1}$$

In the absence of FPN this peak occurs near the shift Δ. However FPN can affect the peak. Since the images are stochastic, the expectation of correlation is taken,

$$\mu_\alpha^2 \sum_{<k>} s[k+n]s'[k] + \sigma_\alpha^2 \delta[n] \sum_{<k>} s[k]s'[k]. \tag{2.2}$$

Letting $\mu_\alpha = 1$, yields

$$\sum_{<k>} s[k+n]s'[k] + \sigma_\alpha^2 \delta[n] \cdot \sum_{<k>} s[k]s'[k]. \tag{2.3}$$

The result is the sum of the correlation without FPN, $\sum s[k+n]s'[k]$, and a component due to the noise that contributes only at $n = 0$. This shows that the effect of the random gain and offset of each pixel on the correlation of the images produces an additional peak at $n = 0$. For motions less than $\Delta = 1$, the superposition of this peak with the peak of the correlation without FPN can pull the maximum correlation toward zero for even low noise levels. For large FPN levels the peak at zero may dominate the local maximum due to the motion.

Chapter 3

Methods

The effect of FPN is studied on both simulated and real world images with known displacements. A technique is developed to counter the effects of FPN on motion estimation. This technique and other image based motion estimation algorithms were used on the images and compared to the known motions to characterize the measurement errors.

3.1 Generation of Simulated Data

3.1.1 Test Images

The motion estimation algorithms were tested on simulated motions of four different images (see figure 3-1) with a variety of frequency characteristics and signal energy levels. The images that were chosen are the same as in (Davis & Freeman 1998). The first image was a simulated dark bead on a bright background (Figure 3-1 a). The bead was simulated as a radially symmetric Hanning window. The function $r = \sqrt{(i-c_x)^2 + (j-c_y)^2}$ is the distance from the center (c_x, c_y) of the simulated bead to pixel location $[i, j]$. Two sets of images (Figure 3-1 b, c) were chosen as examples of applications of motion estimation characterizing the motion of micro-electromechanical systems (MEMS) and of the mechanically sensitive bundles of the inner ear. The final image (Figure 3-1 d) was taken from the space shuttle using

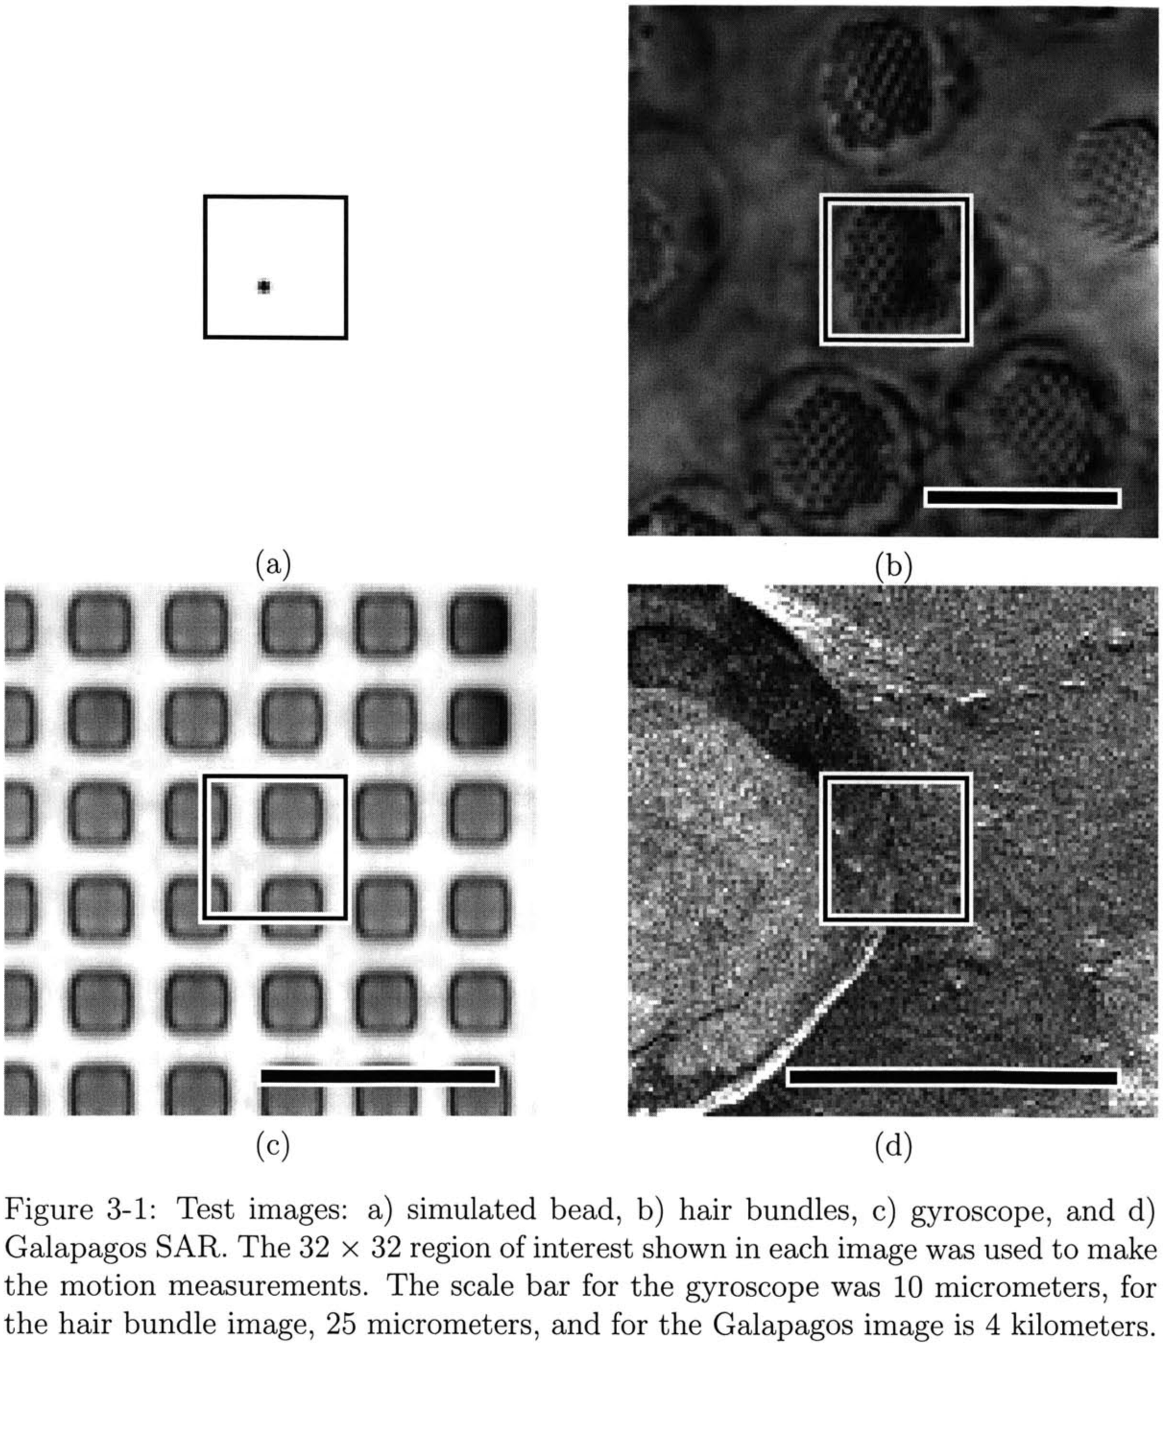

Figure 3-1: Test images: a) simulated bead, b) hair bundles, c) gyroscope, and d) Galapagos SAR. The 32×32 region of interest shown in each image was used to make the motion measurements. The scale bar for the gyroscope was 10 micrometers, for the hair bundle image, 25 micrometers, and for the Galapagos image is 4 kilometers.

Synthetic Aperture Radar (SAR) and is of the Galapagos islands.

3.1.2 Shifting Algorithm

Motions were simulated by generating a sequence of images that were each shifted versions of one original image. The shifted images were produced by (1) taking the Discrete Fourier Transform of the original that is 128×128 pixels; (2) multiplying by a circular shift filter ($e^{j\Delta_x \omega_x + j\Delta_y \omega_y}$); (3) inverse transforming; (4) cropping the edges so that the resultant image is 64×64 pixels. This frequency domain technique can produce arbitrary shifts Δx, Δy with subpixel accuracy.

3.1.3 Simulated Noise

Fixed pattern noise was simulated by multiplying each image by an array of "pixel gains". The gain for each pixel was determined using a pseudo-random Gaussian sequence with a mean of one and standard deviation of σ. The gain for each pixel were chosen independently of the other pixels. The value of σ controls the relative noise level of the fixed-pattern noise. The standard deviation σ was assigned values from 0.0001 to 10 evenly spaced on a log scale. Both the shifted and original images were multiplied by the same array of "pixel gains".

3.2 Acquisition of Real-World Data

3.2.1 Apparatus

A system as detailed in (Davis 1997) and further developed in (Desai 2002, Aranyosi 2002) was used to capture images of a high contrast target moving with displacements ranging from one nanometer to one micrometer. The images were acquired with a Pulnix TM1010 CCD camera attached to a Zeiss Axioplan II microscope with a Zeiss 20x Epiplan LD 0.4 NA microscope objective. A green led was strobed at 8 different phases of a base frequency of one kilohertz. The motions were generated by applying a one kilohertz sinusoidal stimulus to a piezo device. The amplitude of the motion

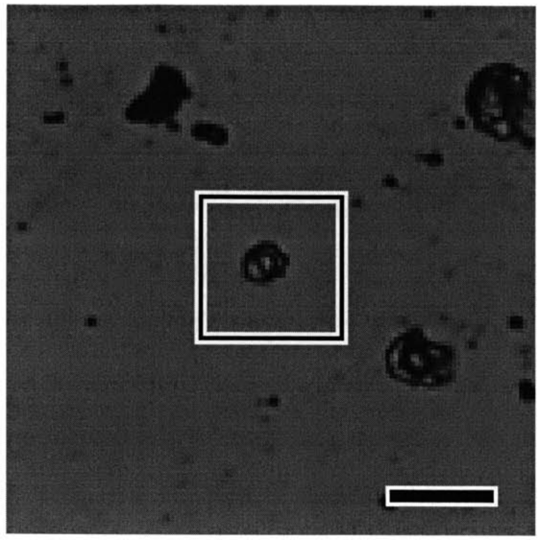

Figure 3-2: Experimental Image: piece of silicon wafer with impurities on its surface. Images of this target were taken of 8 different phases for motions ranging from one nanometer to one micrometer. The 32×32 region of interest shown in the image was used to make the motion measurements. The scale bar is 50 micrometers.

was controlled by varying the amplitude of the sinusoidal stimulus. A mirror was attached to the end of the piezo device so that a laser Doppler vibrometer (Polytec models OFV 3001 and OFV 511) could be used as a motion standard to compare to the image based methods. The laser Doppler is accurate to about one picometer.

3.2.2 Test Structures

Figure 3-2 shows an image of the high contrast target that was moved and later tracked with the motion estimation algorithms. The image is a magnified view of the surface of a piece of silicon with impurities on its surface.

3.2.3 Image Acquisition

One hundred images were acquired at each of 8 phases using stroboscopic illumination and the CCD camera. A hundred images of each phase were averaged to reduce the effect of the variability in the number of photon arrivals over the exposure time, also known as shot noise.

3.3 Motion Measurement from Image Data

3.3.1 Homomorphic Filter

Each motion measurement starts with a sequence of two images. Each image was homomorphically filtered independently of the other. The filtering was done by (1) taking the natural log of each pixel value; (2) performing Matlab's 2-D adaptive Wiener filter of size four by four on the resultant log of each image; (3) exponentiating the output of the filter. For large noise levels the simulations resulted in negative pixel values. Images from CCD cameras are non-negative so the absolute value was taken of each pixel value. A value of 2.2204×10^{-16}, the smallest number in floating point precision, was added to each pixel value to prevent taking the log of zero. The 2-D adaptive Wiener filter estimates the spectrum of the FPN and the signal in a region and then adjusts the values of the filter to remove the noise without removing signal. The adaptive Wiener filter generates filter coefficients using estimates of the local mean and variance along with a global estimate of the noise. There are several different algorithms for this, one is described in (Lim 1990). A filter size of 4 was chosen because it gave the best power of the signal to that of the error in the image values over the noise levels tested.

3.3.2 Lowpass Filter

A lowpass equiripple filter with a transition band from $\pi/50$ to $\pi/10$ of length 35 was used to filter the images. The rows were filtered first using this filter, the resulting columns were then filtered. The cutoff was chosen, to preserve the spectral information in the bead image. The bead image contained all of its energy below this cutoff frequency. For low FPN levels this would act similarly to the Wiener filter on the image of the bead.

3.3.3 Gradient Algorithm

The motion estimation algorithm in this study is based on optical flow and utilizes local spatial and temporal gradients as developed in (Horn 1986, Horn & Schunck 1981, Horn & Schunck 1993, Horn & Weldon, Jr. 1988). The algorithm uses linear bias correction (LBC) to decrease biases in the gradient algorithms; this is further described in (Davis & Freeman 1998).

3.3.4 Correlation

The correlation plots of the two images were generated by (1) upsampling and interpolating each image by a factor of ten; (2) subtracting the mean from each of the images; (3) flipping the shifted upsampled image along each axis; (4) and taking the FFT of both the flipped image and upsampled original; (5) multiplying the transformed results; (6) taking the magnitude of the IFFT of that result.

3.3.5 Two Point Correction

Two point correction is a method for the attenuation of FPN (Healey & Kondepudy 2001, Agard, Hiraoka, Shaw & Sedat 1989). It is performed by first generating a set of two images, a bright and a dark image, then using those images to adjust the pixel values of the image to be corrected. The bright image is generated by taking the average of 100 images of the out of focus target or of the illumination source in the absence of a target. The dark image is generated by taking the averages of 100 images with the illumination off. The corrected image $C[i,j] = \frac{M[i,j]-D[i,j]}{B[i,j]-D[i,j]}$, where $M[i,j]$ is the measured pixel value, $B[i,j]$ the bright pixel value, and $D[i,j]$ the dark pixel value all at location $[i,j]$. The idea behind two-point correction is that $B[i,j]$ should result from constant illumination intensity at each pixel, so variations in $B[i,j]$ are proportional to FPN.

Chapter 4

Results

The performance of the sub-pixel motion estimation used in conjunction with homomorphic filtering was tested between pairs of computer-generated images with simulated motions and fixed pattern noise. Each pair of images consisted of an original and a computer-generated shift of that image. In addition to the simulations, the algorithms were tested on images of real targets captured experimentally.

4.1 Simulation Results

Figure 4-1 shows the correlation between the original gyro image and the shifted gyro image, both containing FPN. It contains a peak near the zero shift which dominates the local maximum near the location of the simulated shift. Figure 4-2 shows the resulting correlation following the application of the homomorphic filter to both images. The peak located at the zero shift, seen in Figure 4-1, is greatly attenuated and no longer visible, while the peak near the imposed shift now contains the maximum correlation.

Each simulated shift in x and y is represented by a pair of displacements (d_x, d_y). Displacements were simulated for 21×21 shifts of (d_x, d_y) from $(-1, -1)$ to $(1, 1)$ on each of the images. Figure 4-3 shows the estimates of the shift in x for each imposed d_x. The estimates of d_x were averaged over the estimates from the 21 different d_y. The results show that over each d_x the measurements on the images that were filtered with

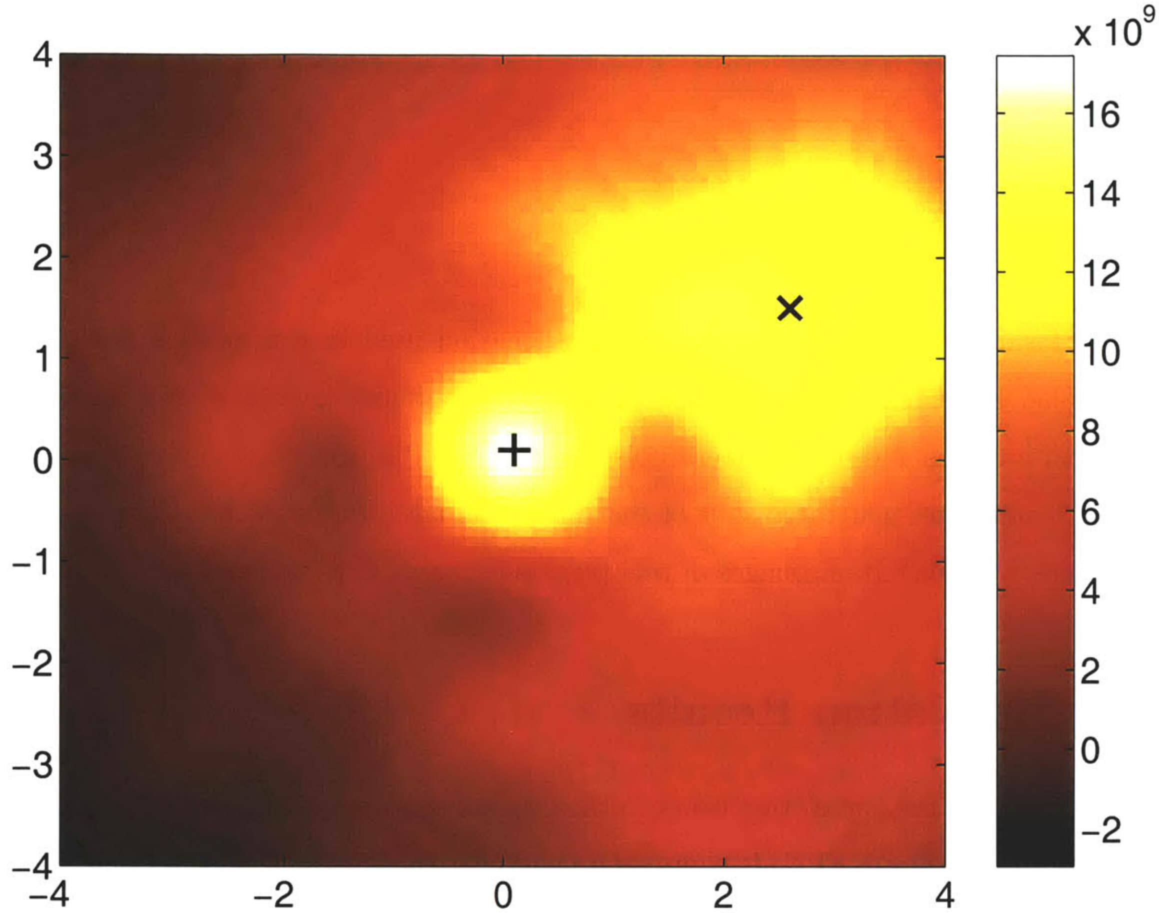

Figure 4-1: The values of the cross-correlation for a pair of the gyro images in the presence of FPN with a standard deviation of $2 \cdot 10^{-1}$ as a function of the shift in x and in y. The $+$ indicates the location of the maximum correlation and the $\times$ indicates the position of the simulated shift which was $(2.6, 1.5)$.

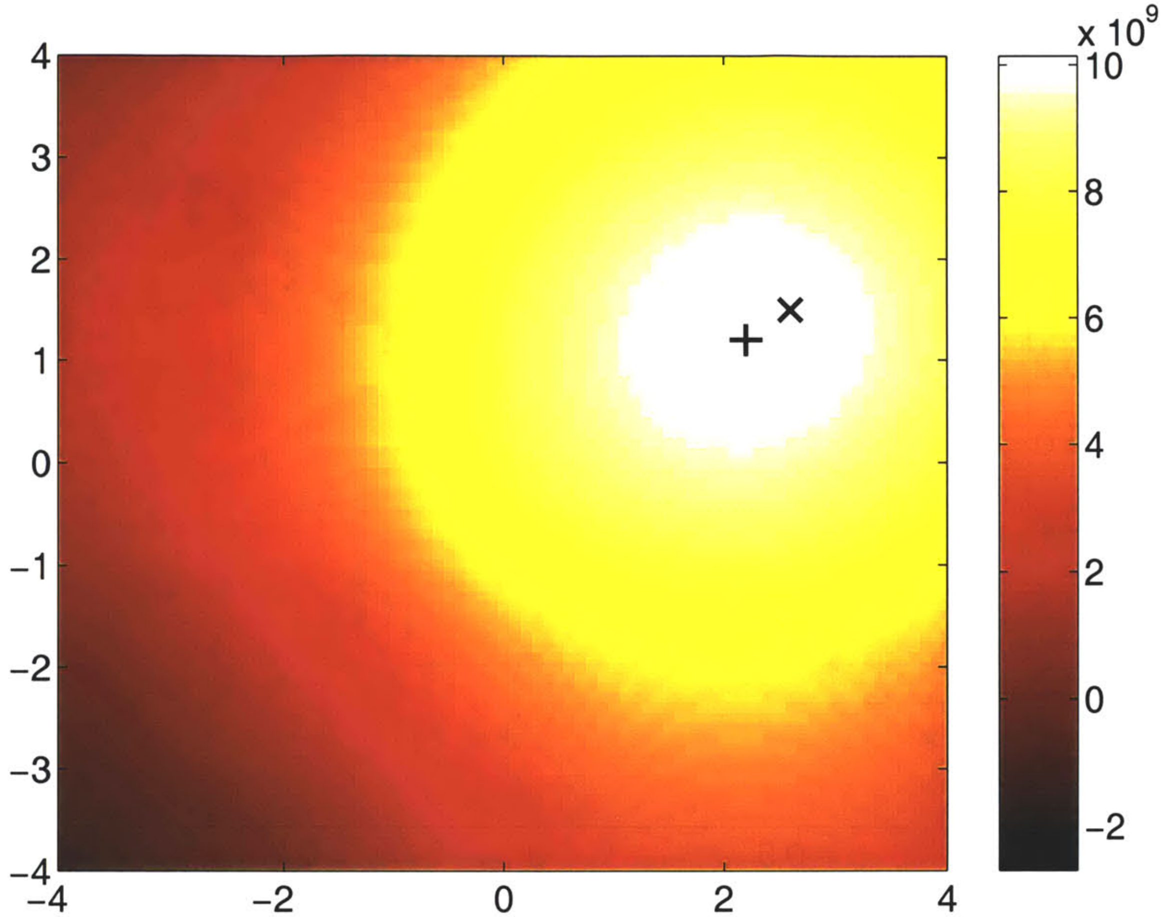

Figure 4-2: Values of the cross-correlation for a pair of images in the presence of fixed pattern noise following the application of the homomorphic Weiner filter. The + indicates the location of the maximum correlation and the × indicates the position of the simulated shift. Observe that the peak corresponding to the (0,0) shift in figure 4-1 is no longer visible and the peak remaining is near the imposed displacement of $(2.6, 1.5)$.

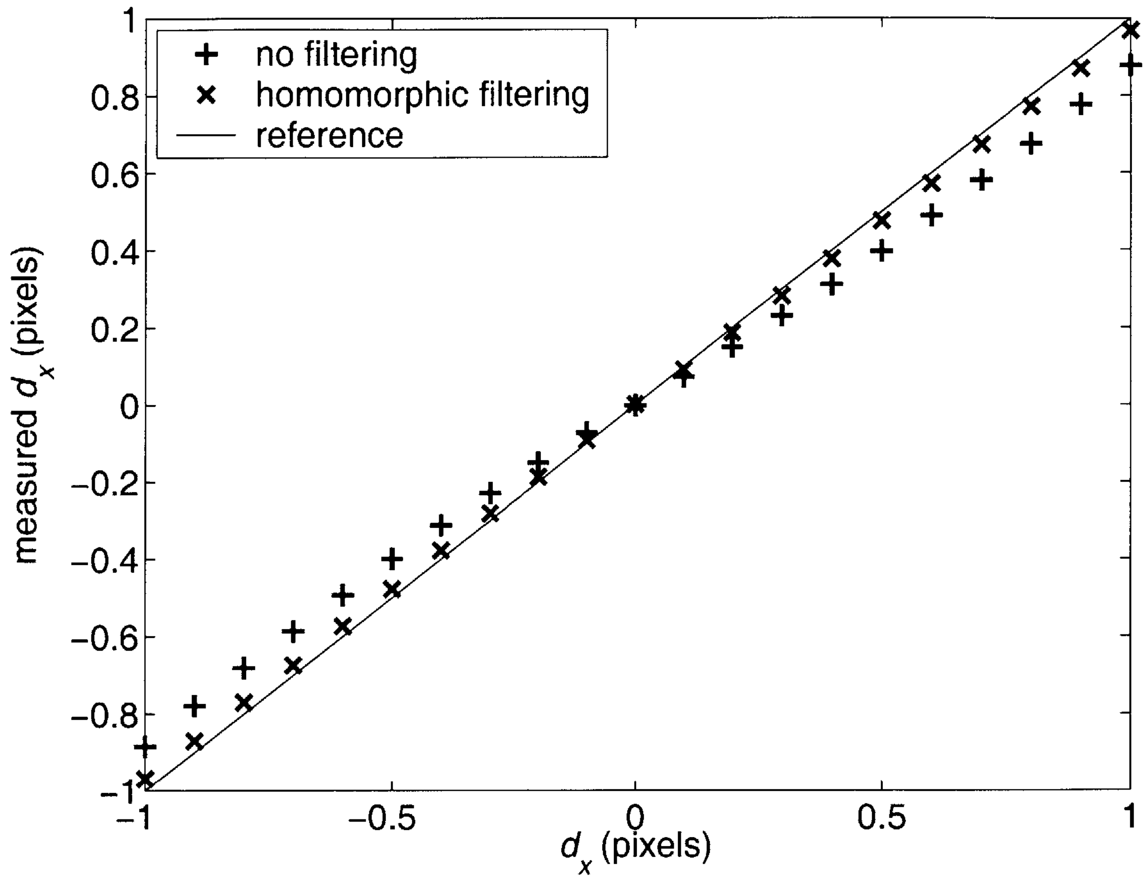

Figure 4-3: The measured shift in x averaged over each y shift as a function of d_x for both algorithms. Displacements were simulated for 21 $\times$ 21 values of (d_x, d_y) from $(-1, -1)$ to $(1, 1)$ on the image of the hair bundles. The standard deviation of the FPN was $3.2 \cdot 10^{-2}$.

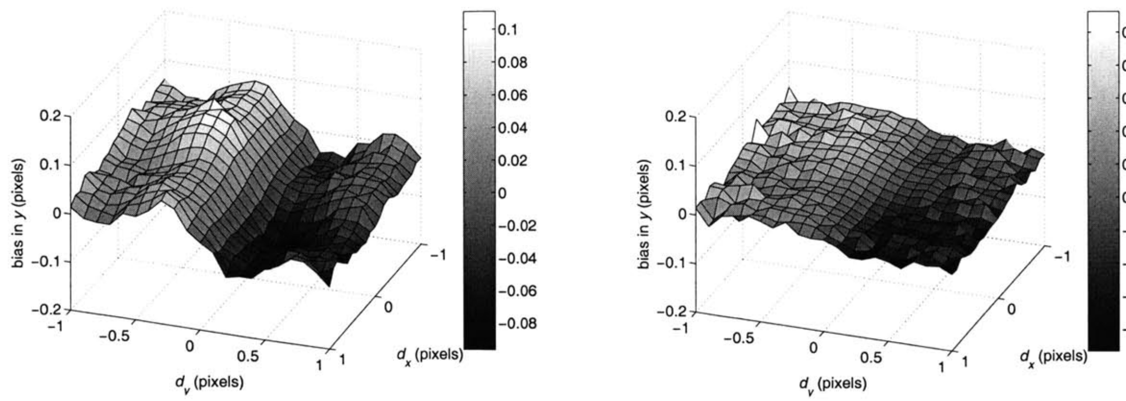

(a) (b)

Figure 4-4: Bias in estimates of y displacements without (a) and with (b) homomorphic pre-filtering for the SAR image of Galapagos Islands. The measured bias is the difference between the measured shift in y and the imposed displacement d_y. Displacements were simulated for 21×21 values of (d_x, d_y) from $(-1, -1)$ to $(1, 1)$ of the Galapagos image. The standard deviation of the fixed pattern noise was $3.2 \cdot 10^{-2}$.

homomorphic filtering were closer to the actual value than were those measurements without the filtering. The magnitudes of the estimates of both algorithms are smaller than the imposed displacements d_x indicating a systematic bias towards zero shift.

Figures 4-4 a and b show the measured bias of the measurement of y as a function of (d_x, d_y). For constant y values and any x value the magnitude of the bias does not vary significantly compared to the variation as d_y is varied for a fixed d_x, i.e, the bias in y was mostly independent of d_x. The bias in the filtered images is on average 2.4 times smaller than that of the unfiltered images.

The vector fields in figure 4-5 provide a way to see both the direction and magnitude of errors of each imposed shift simultaneously. Each arrow points from the applied shift to the measured shift. The length of the arrow is indicative of the magnitude of the bias.

The arrows for the errors of motion estimation using both the unfiltered and the homomorphic filtered images point toward the center, indicating that both algorithms are biased toward zero. Although the vectors in both graphs are pointing in similar

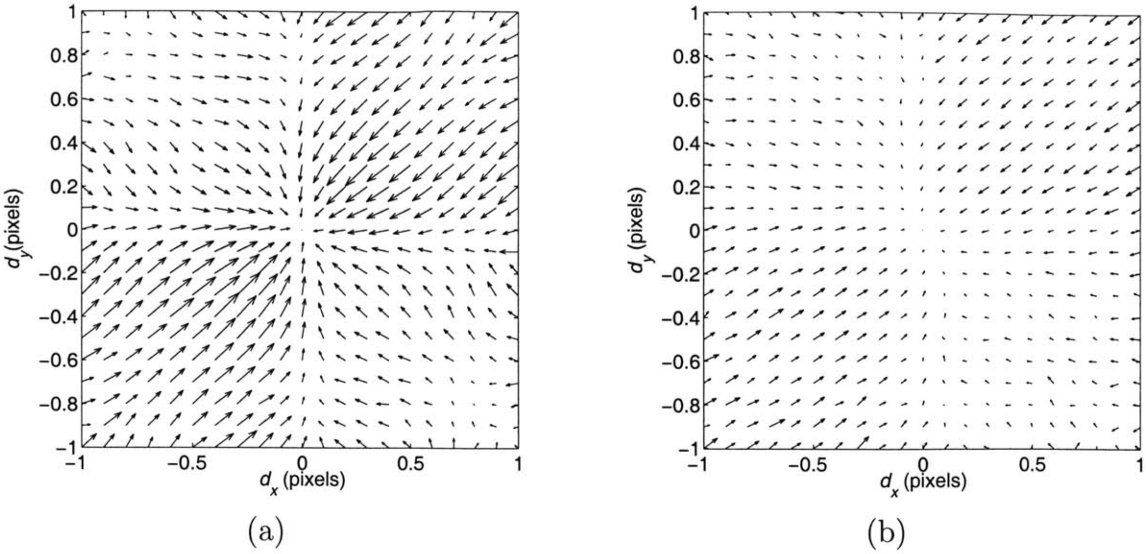

Figure 4-5: Vector field representation of the errors without (a) and with (b) homomorphic filtering for each 21 $\times$ 21 shifts (d_x, d_y) . The tail of each arrow lies on the location of the simulated shift (d_x, d_y). The head of that arrow points to the measured displacement $(\hat{d_x}, \hat{d_y})$.

directions the length of each vector (amount of bias) of each measurement is smaller in figure 4-5 b than in figure 4-5 a. This shows that the application of the homomorphic filter reduces bias.

Figure 4-6 shows the average of the magnitude of the error over all of the simulated displacements for each of the test images. In each of the plots the measurements of the homomorphic filtered images perform either better than or the same as those of the unfiltered images. Each of these plots has the same characteristic shape with three different regions: a low noise region in which the behavior of both algorithms does not vary as a function of the FPN, an intermediate range in which the error of both algorithms increases with the level of the FPN, and a high noise level in which the errors in both algorithms perform do not vary as a function of the FPN. The ratios of the errors of the two algorithms in both the low and middle noise ranges are shown in Table 4.1. In the high noise range, the measurements on the images with homomorphic filtering perform a small amount better (less than a dB) than without.

In addition to the results in Figure 4-6, Figure 4-7 contains the results of measurements using a simple lowpass filter with a transition band from $\pi/50$ to $\pi/10$.

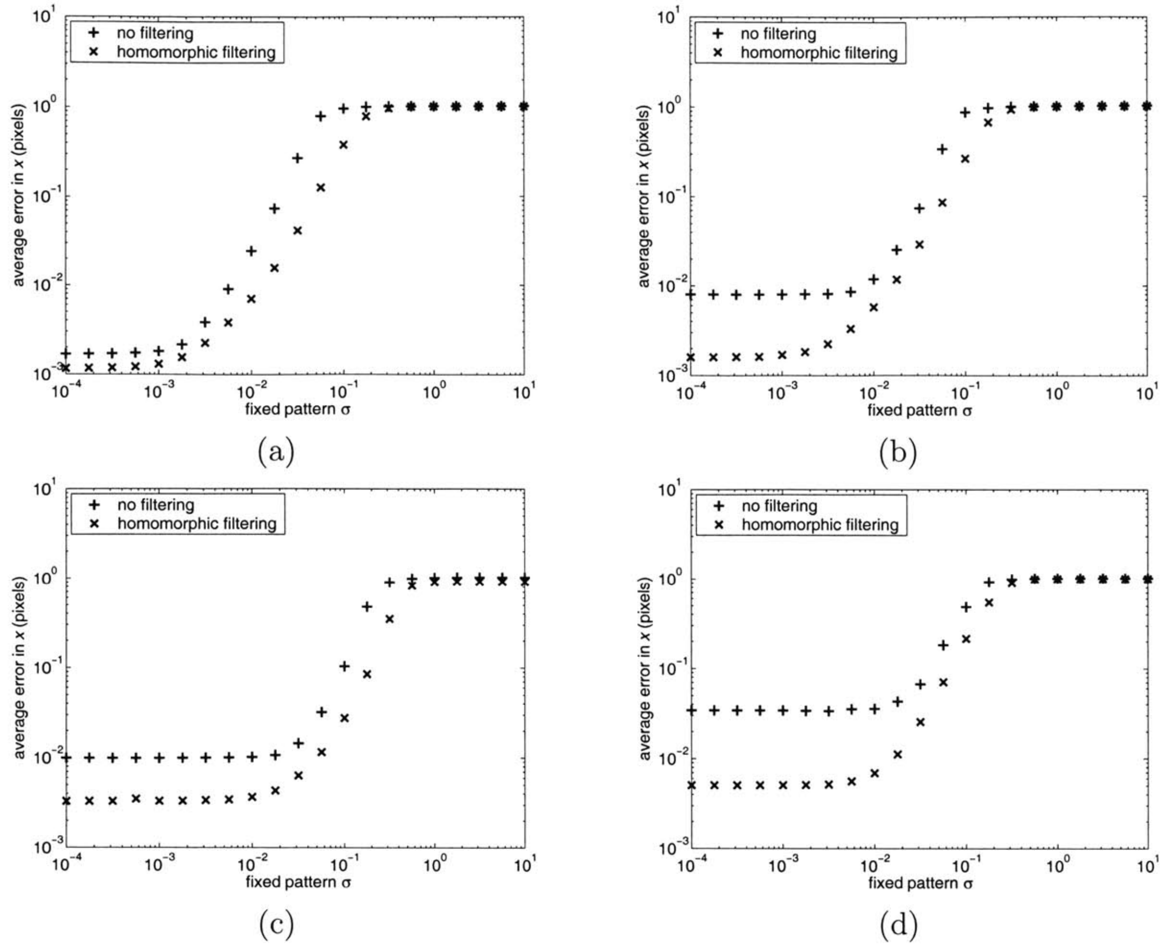

Figure 4-6: Average of the absolute value of errors in the measured values of x for each of the test images — a) simulated bead, b) Hair Bundles, c) Gyro, and d) Galapagos SAR — as a function of the standard deviation of the fixed pattern noise. Errors were determined for the algorithms with and without homomorphic filtering. Each point shows the average error of each of the 41×41 imposed displacements from $(-2,-2)$ to $(2,2)$ of (d_x, d_y).

	low		intermediate	
Image	range in σ	ratio (dB)	range in σ	ratio (dB)
simulated bead	10^{-4}-10^{-3}	3.2	$2 \cdot 10^{-3}$-10^{-1}	11.9
gyroscope	10^{-4}-$2 \cdot 10^{-3}$	9.6	$3 \cdot 10^{-3}$-10^{-1}	10.2
hair bundle	10^{-4}-10^{-2}	14.0	$2 \cdot 10^{-2}$-$2 \cdot 10^{-1}$	8.5
SAR Galapagos	10^{-4}-$6 \cdot 10^{-3}$	16.6	10^{-2}-$3 \cdot 10^{-1}$	7.8

Table 4.1: Average ratio between the errors in motion estimates of unprocessed images and processed images. Low and intermediate noise ranges for each test image are shown. The measurement errors come from Figure 4-6.

The results show that low pass filtering of the images yields more accurate motion estimates than no filtering at all over the entire range of FPN levels. The homomorphic filtering produces more accurate estimates in all intermediate and high noise levels than do the low pass filtered and un-filtered images. In the low noise ranges the homomorphic filter yields the best estimates on all but the Galapagos image, where the motion estimates of the low pass filtered images have 1.3 dB smaller error than do the homomorphic filtered images. In the low noise ranges the homomorphic filtered images bested the low pass filtered images by 3.2 dB, 2.7 dB, and 7.3 dB for the bead, gyro,and hair bundle images, respectively. The improvements over the standard algorithm for the low pass filtered images over the unfiltered images in that range were 3.2 dB, 6.9 dB, 6.3 dB and 17.9 dB, respectively.

Figure 4-8 shows the results for the adaptive Wiener filter without using the log operator along with the homomorphic filter and the unfiltered case. The results show that the performance on the simulated bead and hair bundles is the same, while the homomorphic filtering yields more accurate measurements on the gyro image and the Galapagos image. The greatest differences are in the low to mid FPN ranges.

4.2 Experimental Results

In addition to the simulations, the motion estimation algorithms were tested on images obtained experimentally. The measured displacements of the image based algorithms are plotted against measurements of motion of the same target made using the laser Doppler in figure 4-9. For each displacement ranging from about one nanometer to one micrometer, the measurement made by using the homomorphic filtering is closer to the reference measurement indicated by the straight line.

Figure 4-10 shows the magnitude of the errors for the experiment as a function of the displacements measured with the laser Doppler. The error for the measurements on the pre-processed images are smaller than those for the unprocessed images for the two point corrected images for all displacements. The two point corrected images

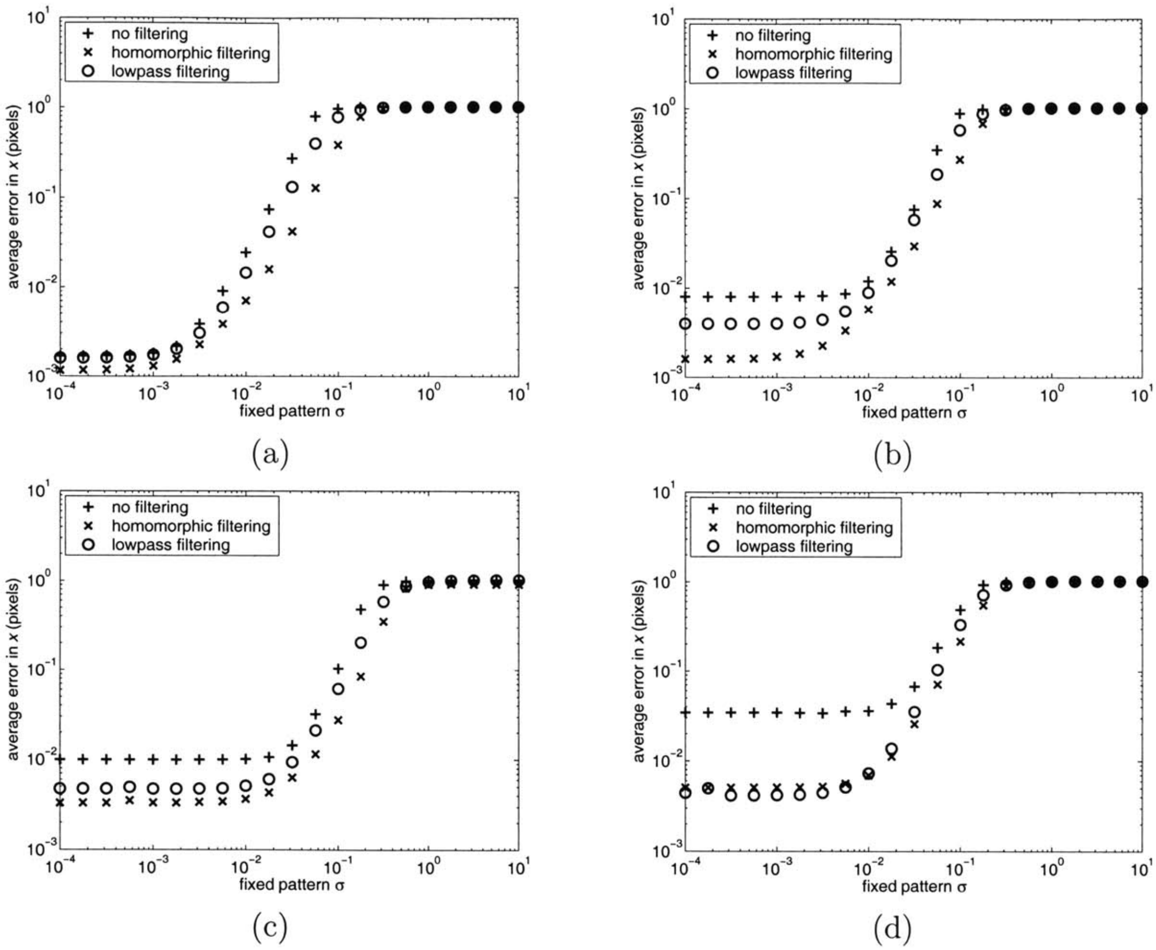

Figure 4-7: The same plot as in figure 4-6, except that the results from the algorithm using a simple lowpass filter with a transition band from $\pi/50$ to $\pi/10$ are included. The images used are a) simulated bead b) hair bundles c) gyro and d) Galapagos SAR.

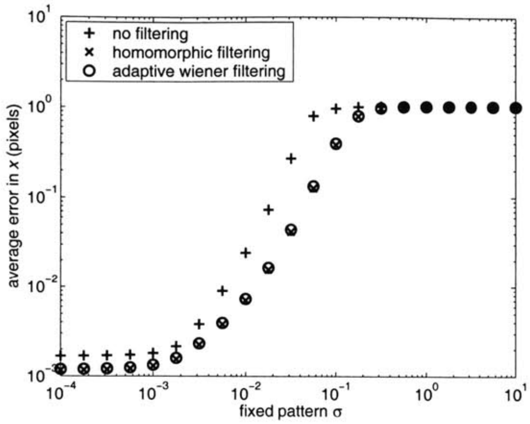

(a)

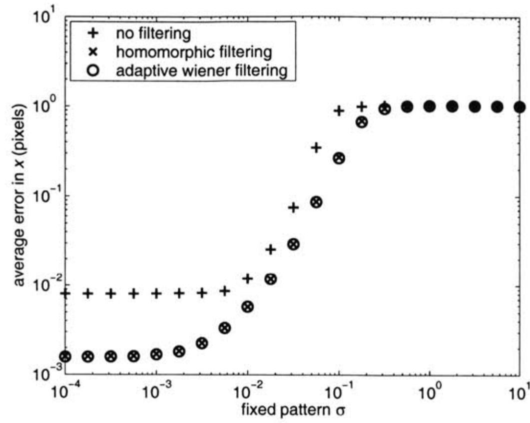

(b)

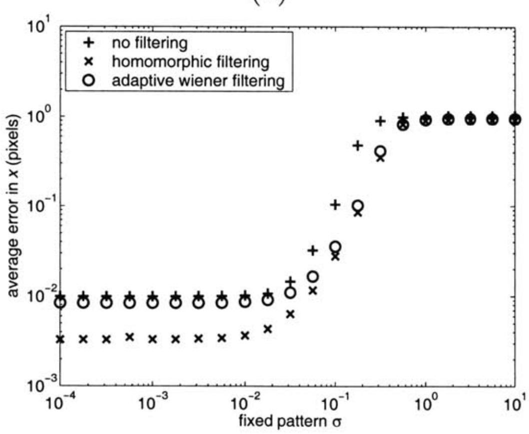

(c)

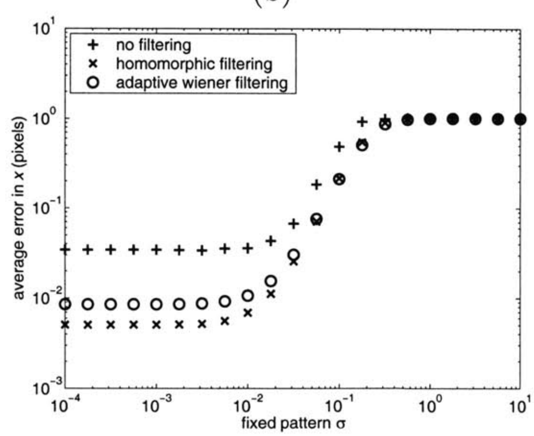

(d)

Figure 4-8: The same plot as before in figure 4-6, except that the results from the algorithm using an adaptive wiener filter (of size 4) are included. The images used are a) simulated bead b) hair bundles c) gyro and d) Galapagos SAR.

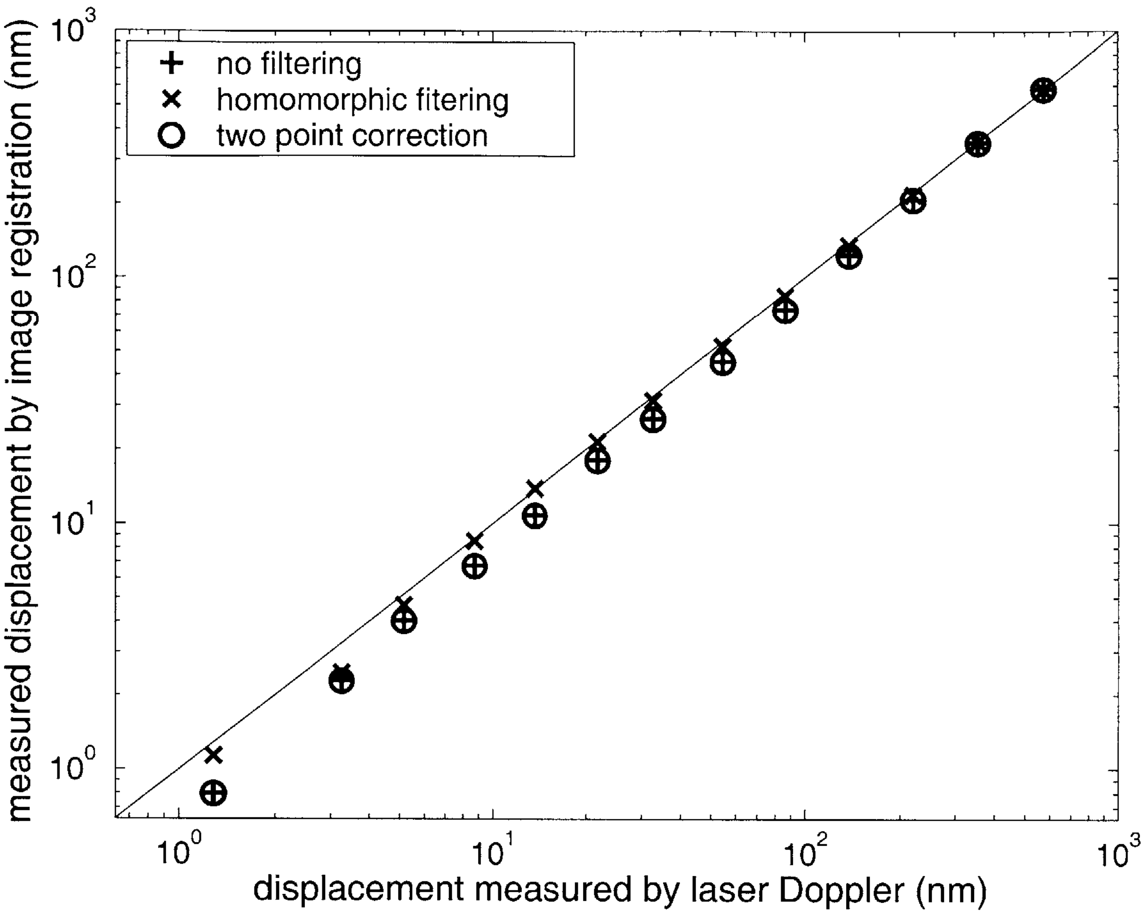

Figure 4-9: The measured displacements from the image based algorithms compared to the motion measurements from the laser Doppler. A straight line is shown as a reference. Each measurement is made from the average of 100 images to reduce the effect of shot noise.

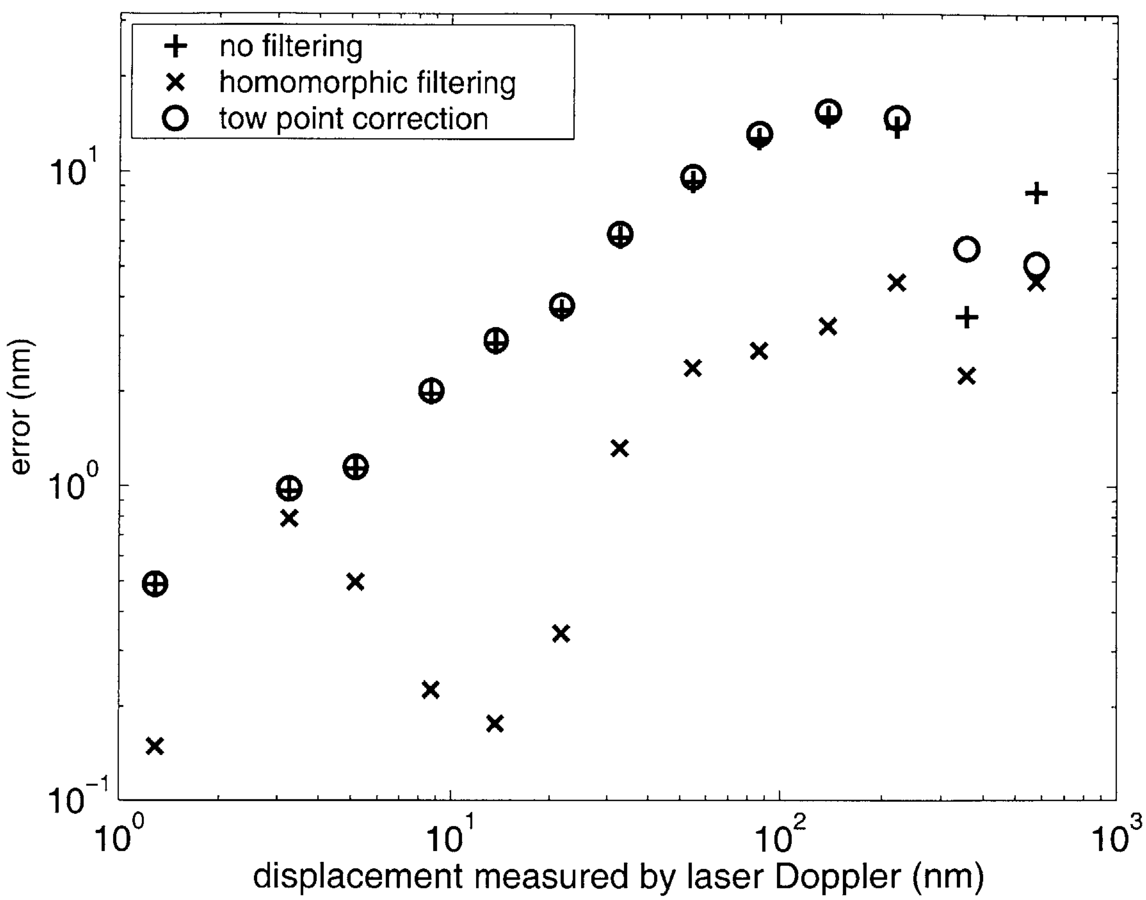

Figure 4-10: Magnitude error in the measurement of the laser Doppler and the image based algorithms plotted against displacement as measured by the laser Doppler. The error for the method using homomorphic filtering is the smallest over all imposed displacements.

give approximately the same result as for the uncorrected images except at the largest two displacements. The average differences in the size of the measured error is 12dB. The differences range from the worst case of 1.75 dB to the best case of 24.2 dB.

4.2.1 Computational Costs

The relative computational costs of the optical flow algorithms and the homomorphic filter were estimated by timing a thousand runs of each the 64×64 original and shifted images of the Galapagos islands. The tests were run on a Dell Dimension 8250 computer with a 2.8 GHz Pentium 4 Processor and one gigabyte of RAM. It took 2.3 seconds to read a thousand pairs of images from disk. The iterations of the optical flow took 3.55 seconds and the iterations of the filter took 7.68 seconds. The LBC algorithm makes a minimum of four optical flow measurements which took 14.19 seconds to perform the thousand iterations. Both images were loaded into memory before the algorithms were tested. Performing the homomorphic filtering in addition to the LBC algorithm takes 54% more computation time.

Chapter 5

Discussion

5.1 The Effect of FPN on Motion Estimation

The effect of FPN on the accuracy of the gradient algorithms is consistent with what is shown in the correlation plots. As FPN increases the measurements are biased more and more towards zero. By applying the pre-filtering to the images the peak located at the zero-shift is attenuated in the correlation plots and the error in the measurements is diminished. This shows that as the energy of the correlation due to the FPN is reduced, more accurate motion estimations are possible.

5.2 Frequency Content of the Images Matters in Motion Estimation

Table 4.1 is arranged by the portion of energy in the images at frequencies less than $\pi/10$. The simulated bead, which contains only low frequency energy, appears first in the table, while the Galapagos image containing significant energy at all frequencies appears last. The table shows two major trends: (1) the improvement in accuracy gained by homomorphic filtering in the low FPN region increases with the amount energy at higher frequencies in the images and (2) this improvement decreases in the intermediate FPN region.

In the low FPN region the magnitude of the errors is independent of the FPN. This suggests that the reduction in the error is not governed by the removal of fixed pattern noise. For these noise levels the homomorphic filter acts as a homomorphic low pass filter. In (Davis & Freeman 1998) it is shown that as the frequency content of an image increases so does the bias. This result suggests that removing the high frequency content from an image will decrease the bias in the measured displacement. Since the bead image consists of only low frequency energy not much improvement in motion estimation is gained by filtering, whereas a large improvement is seen with the Galapagos image.

Since the errors in the intermediate range depend on FPN, it is clear that by removing FPN, more accurate measurements of motion can be made. Because the bead image contains only low frequency information, the FPN — which is white in the Cepstral domain — can be isolated and removed by the adaptive Wiener filter. In the Galapagos image it is more difficult for the adaptive Wiener filter to remove the noise because the signal contains energy over a larger portion of the spectrum.

This analysis is further supported by the measurements made on images using low pass filtered images. In the low FPN range, for the image with the least amount of information over $\pi/10$, the low pass filtering had the smallest performance gains over the unfiltered images. The smallest gains were obtained on the bead image, which consisted of only low frequency information. By reducing information at higher frequencies, the errors in the measurements were reduced. The deviation between the two different filtering schemes in this low FPN region is due to their different low pass filter character.

The images of the gyro and of the Galapagos have a larger portion of their energy over $\pi/2$ than do the other two images. This makes it more difficult for the Wiener filter to separate signal from the noise. As seen in Figure 4-8, by taking the log of the image values the adaptive Wiener filter can then better separate the signal from the noise.

5.3 Homomorphic Pre-filtering Improves Motion Estimation

Using a low pass filter, improvements were gained over the unfiltered images for all FPN noise ranges. In regions where FPN affects the accuracy of the measurements, the homomorphic filtered images yield better motion estimates than the low pass filtered images in all cases. In the low noise ranges, the low pass filter has good performance, but the homomorphic filter still works best in three of the four test images. On the image that it performs worse on, the difference is only 1.3 dB, while on the other three images it performed better by 3.2 dB, 2.7 dB, and 7.3 dB.

A simple adaptive Wiener filter provides improvements in motion estimation over the un-filtered case. However, by taking the log of the image values before Wiener filtering, better accuracy in motion estimation is achieved.

The amount of variation in gain across the image sensor can be seen as a floor on the level of the FPN since there is typically additional noise in the optical setup. This places the FPN level of a typical camera setup, usually between 1-6% across a CCD (Ilyin 2002), in a range where the homomorphic filtering is the best choice for removing the FPN and producing the most accurate measurements.

In both the simulations and the experiments, the proposed pre-filtering produced more accurate results than did the algorithms without the filtering. The algorithms produced improvements in accuracy of 7.8 dB to 11.9 dB in the simulations, and an average of 11 dB in the experiments. This improvement results from two factors: reducing the FPN in the images and in low-pass filtering of the images. Under no condition was the accuracy of the motion estimation algorithm reduced by the filtering. This technique has a relatively small computation cost and therefore should be applied whenever gradient-based motion estimation is used.

Bibliography

Agard, D. A., Hiraoka, Y., Shaw, P. & Sedat, J. W. (1989). Fluorescence microscopy in three dimensions, *Methods in Cell Biology* **30**: 353.

Aranyosi, A. J. (2002). *Measuring Sound-Induced Motions of the Alligator Lizard Cochlea*, PhD thesis, Massachusetts Institute of Technology, Cambridge, MA.

Cain, S. C. & Hayat, M. M. (2001). Projection-based image registration in the presence of fixed-pattern noise, *IEEE Transactions on Image Processing* **10**: 1860–1872.

Campisi, P., Yan, J. C. K. & Hatzinakos, D. (2000). Signal-dependent film grain noise generation using homomorphic adaptive filtering, *IEE Proc.-Vis. Transactions on Image Signal Process* **147**: 283–287.

Davis, C. Q. (1997). *Measuring nanometer, three-dimensional motions with light microscopy*, PhD thesis, Massachusetts Institute of Technology, Cambridge, MA.

Davis, C. Q. & Freeman, D. M. (1998). Statistics of subpixel registration algorithms based on spatio-temporal gradients or block matching, *Optical Engineering* pp. 1290–1298.

Desai, S. P. (2002). *Mirau interferometric computer microvision*, Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Hadhoud, M. M. (1999). Image contrast enhancement using homomorphic processing and adaptive filters, *16th National Radio Science Conference*, Cairo, Egypt.

Healey, G. E. & Kondepudy, R. (2001). Radiometric CCD camera calibration and noise estimation, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **10**: 1860–1872.

Horn, B. K. P. (1986). *Robot Vision*, MIT Press, Cambridge, MA.

Horn, B. K. P. & Schunck, B. G. (1981). Determining optical flow, *Artificial Intelligence* **17**: 185–203.

Horn, B. K. P. & Schunck, B. G. (1993). *"Determining Optical Flow"*: a retrospective, *Artificial Intelligence* **59**: 81–87.

Horn, B. K. P. & Weldon, Jr., E. J. (1988). Direct methods for recovering motion, *Internatl J of Computer Vision* **2**: 51–76.

Ilyin, I. (2002). Optimization of CCD performance.
URL: *http://cc.oulu.fi/ ilyin/sofin/sofin/node3.html*

Jain, A. K. (1989). *Fundamentals of Digital Signal Processing*, Prentice Hall Inc., Upper Saddle River, New Jersey.

Janesick, J. R. (2001). *Scientific Charge Coupled Devices*, SPIE - The International Society for Optical Engineering, Bellingham, Washington.

Janesick, J. R. (2002). Lux transfer: Complementary metal oxide semiconductors versus charge-coupled devices, *Optical Engineering* **41**: 1203–1215.

Lim, J. S. (1990). *Two-Dimensional Signal and Image Processing*, Prentice Hall Inc., Upper Saddle River, New Jersey.

Oppenheim, A. V., Shafer, R. W. & Stockham, Jr., T. G. (1968). Nonlinear filtering of multiplied and convolved signals, *Proc. IEEE* **56**: 1264–1291.