

SUBPIXEL ACCURACY ANALYSIS OF PHASE CORRELATION REGISTRATION METHODS APPLIED TO ALIASED IMAGERY

Esteban Vera and Sergio Torres

Department of Electrical Engineering, University of Concepcion
Casilla 160-C, Concepcion, Chile
phone: + (56) 2207476, fax: + (56) 2246669, email: estebanvera@udec.cl
web: www.die.udec.cl

ABSTRACT

Phase Correlation is an efficient technique for subpixel image alignment under translational motion, but its subpixel accuracy normally suffer because of aliasing, which can be found in some undersampled imaging systems. In this work we review, compare, and mix different phase correlation based methods and test them under diverse aliasing conditions, from partial to total. Results show that with some additional features, the methods based on the one or two dimensional Gaussian fit of the spatial phase correlation peak are the less sensitive to severe aliased imagery.

1. INTRODUCTION

Subpixel image registration is a fundamental task for high performance image processing techniques such as image fusion and superresolution, which have been extensively used for applications in remote sensing, medical imaging, surveillance and computer vision. For example, in multiframe superresolution reconstruction, the registration accuracy is directly related to the quality of the final superresolved images obtained, independently of the reconstruction method used.

In the literature [1], frequency based registration methods based on Phase Correlation (PC) have been highlighted because of its accuracy and low complexity for registering motion due to translation, rotation or scale changes between images. The original PC method for image alignment [2] relies on the shift property of the Fourier transform to estimate the translation between two images, and it is extended to estimate rotation and scale changes by using log-polar coordinate changes [3]. Originally limited to discover only integer pixel translations, the algorithm can be naturally extended to provide subpixel accuracy [4], but at a higher computational cost. On the other hand, newer approaches have also allowed to obtain subpixel accuracy with much less complexity [5], and some of its latest variations have reported enhanced accuracy performances [6, 7, 8]. Nevertheless, only a small portion of PC based methods have a explicit treatment for avoiding aliasing effects [9], by working directly on the Fourier domain using the Phase Difference instead of the Phase Correlation, which might be advantageous when dealing with images already obtained in the frequency domain, such as Magnetic Resonance Imaging (MRI).

Anyway, aliasing is recognized as the main responsible on the subpixel accuracy degradation of PC based registration methods [10]. But as aliasing is not avoided nor completely suppressed in some imaging systems, such as in infrared focal-plane arrays (IRFPA), it is still widely found, and sometimes even a key factor in successful multiframe superresolution reconstruction applications [11], for example.

However, no previous work have shown the amount of real degradation to expect on the subpixel accuracy when using the PC method with severely aliased images. If this would be the case, we would be able to know its performance under such conditions, and therefore know its real practical limitations and usage. For that reason, the aim of this paper is to evaluate several versions of the PC subpixel image registration method, specially when dealing with undersampled images of high aliasing content, also verifying if some of the features added that have claimed to enhance its accuracy under normal conditions, have, or have not, a real impact in really improving the performance under severe aliasing conditions.

This paper is organized as follows. In the next section, a description of the phase correlation method for subpixel registration is presented, detailing also the variations to be used in this work. Section 3 describes the way we simulate the subpixel displacements and the aliasing content of the test images. Experimental results and comments are given in section 4. The conclusions and final remarks are summarized in section 5.

2. SUBPIXEL REGISTRATION BY PHASE CORRELATION

Let the image I_2 be a shifted version of the image I_1 by (x_0, y_0) , then

$$I_2(x, y) = I_1(x - x_0, y - y_0) \quad (1)$$

After taking the Fourier Transform (FT) of both images, we have the following relationship due to the shift property of the FT

$$\hat{f}_2(u, v) = \hat{f}_1(u, v)e^{-j(u x_0 + v y_0)} \quad (2)$$

Therefore, a shift in the spatial domain will produce a phase difference in the frequency domain. The normalized cross-power spectrum is finally defined as

$$\frac{\hat{f}_2(u, v)\hat{f}_1^*(u, v)}{|\hat{f}_2(u, v)\hat{f}_1^*(u, v)|} = e^{-j(u x_0 + v y_0)} \quad (3)$$

The Phase Correlation (PC) function is finally obtained by taking the Inverse Fourier Transform (IFT) of the cross-power spectrum, which gives a $\delta(x_0, y_0)$ as a result: a Dirac function centered on the position (x_0, y_0) .

Nonetheless, as pointed out in [5], when dealing with discrete images and using the Discrete Fourier Transform (DFT)

to generate the PC, the Dirac is turned into a Dirichlet kernel, whose maximum peak is found at the closest integer displacement, so finding the PC peak is equivalent to finding the translation at a pixel resolution. In order to obtain subpixel resolution and keep using the same technique of finding the peak position of the PC function, interpolation by zero padding the cross-power spectrum is suggested in [4], but accuracy is limited by the interpolation factor used which is also limited by the size of the IDFT that can be computed. Lately, this approach has been improved with a more efficient implementation proposed in [12]. This is the first method we will use for the tests, and we will name it UDFT.

Using a different approach, in [5] an extension of the original PC method is presented where using not only the PC information of the main peak, but also its surrounding pixels, leads to an estimation of the amount of subpixel displacement as well. Let $C(0,0)$ be the main peak, and $C(1,0)$ and $C(0,1)$ be the neighbors with the largest value in both horizontal and vertical direction respectively. The subpixel displacement is then calculated as:

$$\Delta x = \frac{C(1,0)}{C(1,0) \pm C(0,0)} \quad \Delta y = \frac{C(0,1)}{C(0,1) \pm C(0,0)} \quad (4)$$

This will be the second method to be used in our tests, and we will denomine it FSPE.

By using the same principle of using the information of the main peak and its surroundings, in [8] some separable fitting of different functions into a given neighborhood is proposed. The 1D fitting of a Gaussian function was reported to be superior than the FSPE method, so it will be used as our third method, or 1DGF.

In a similar fashion, in [6] a 2D Gaussian fit to a $N \times N$ neighborhood of the main correlation peak has shown subpixel accuracies of 1/100 of a pixel. This time a Gaussian smoother to the cross-power spectrum should be applied in order to improve the fit quality. This method, named 2DGF, will be used as our fourth method under test.

Other subpixel PC methods work directly in the Fourier domain, considering that the phase difference (PD) of two translated images should remain constant, so the displacements can be obtained by estimating the slope in both directions of the plane that crosses the origin. To account for some limited aliasing, in [9] the slope is calculated by a least square fitting of the phase data bounded to be within a given low frequency radius, and that also satisfies a robustness condition given by a threshold. In [11] a simplified version that only limits the frequency range of the phase data to a low frequency square to perform the least square fitting has shown good performance, even for some small amounts of aliasing. This will be the last method to be tested, named LSPD.

In addition to the methods already described, we would like to mention two pre-processing steps that have shown to improve subpixel accuracy of PC based methods. The first one, proposed in [9], is the use of a window to soften the image borders in order to reduce undesired artifacts in the Fourier domain, and in particular a Blackman window was indicated as adequate. The second one, explored in [7], is the use of the image gradient in both directions instead of the image itself, so each input image G should be calculated as follows

$$G = G^h + jG^v \quad (5)$$

where $G^h(x,y)$ and $G^v(x,y)$ are the image gradients in the horizontal and vertical direction respectively, and they can be obtained as follows

$$G^h(x,y) = I(x+1,y) - I(x-1,y) \quad (6)$$

$$G^v(x,y) = I(x,y+1) - I(x,y-1) \quad (7)$$

3. EXPERIMENTAL SETUP

3.1 Test Images

In order to simulate the subpixel displacements and also have control of the aliasing, we used a couple of large bandlimited images (2200×2200) to generate the final undersampled/aliased imagery sets that will be used to test the algorithms here reviewed.

The original images, displayed in figure 1 a) and c), were passed through a lowpass antialiasing filter in order to control the amount of aliasing content of the further downsampled versions of each one. After choosing the cutoff frequency, the antialiasing filtering is performed using a Gaussian kernel designed in the frequency domain using the desired cutoff. Then, we choose a downsampling factor D and started to generate several shifted and downsampled versions of the original image cropped to a size of $2048 + D + 1 \times 2048 + D + 1$ up to the point we finally have $(D+1)^2$ images of size $2048/D \times 2048/D$, thus covering all the combinations of possible downsampling shifts from 1 to $D+1$ in each direction. Hence, after selecting a reference image, all the others are a complete set of translated versions of the same image within 0 to 1 pixel in both directions, with exact subpixel displacements in portions of $1/D$.

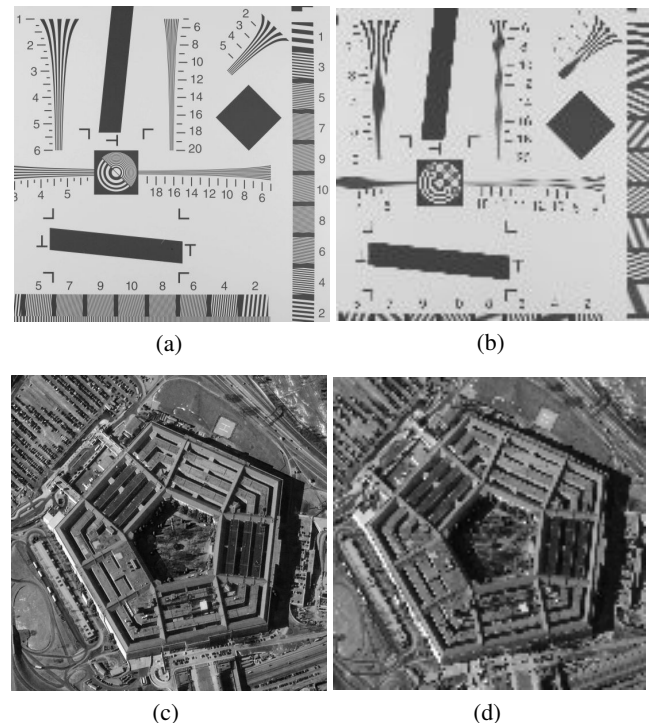


Figure 1: Test Images: a)Resolution Chart and c)Aerial View; b)downsampled and aliased version of a); d)downsampled and aliased version of c).

Thus, we can know exactly how much the real shifts are, so we can also calculate in an objective way the accuracy of a given subpixel estimation technique over a wide range of possible displacements.

For the results presented here, we choose a downsampling factor of $D = 16$, generating a set of 289 images of 128×128 for a particular aliasing setup. The aliasing cutoff frequencies were set in order to allow from 0%, 50%, 100%, 200%, and up to 400% of possible aliasing for each one of the original images, assuming they are both bandlimited. The 400% aliasing case is simply a direct downsampling of the original images with no antialias filtering. Perhaps there is no imaging system that allow for such amount of aliasing, but it can be seen as a worst case scenario. Some samples of highly aliased versions of the original images are shown in figure 1 b) and d).

3.2 Phase Correlation Methods Used

We will use five Phase Correlation based methods that were previously described in section 2, they are: UDFT [4], FSPE [5], 1DGF [8], 2DGF [6] and LSPD [11].

The UDFT was implemented by the enhanced version proposed in [12] using a 1/100 pixel resolution setup. The FSPE method was implemented as it was first proposed. The 2DGF was implemented following the best results obtained in the original paper, with a standard deviation of the Gaussian window of 0.71 and a fitting window of 5×5 pixels. The 1DGF was implemented adding the same Gaussian frequency smoother function as it is proposed for the 2DGF, and also using 5 pixels in each direction because it allows for better results and help to see if the separable fitting is worthwhile. The LSPD was implemented using the original code given in [11], and the frequency cutoff was set to half the frequency content.

Finally, and as described in the previous section, we will try two of the features that are claimed in the literature to improve the subpixel accuracy performance. Therefore, each of the five methods are tested in four different ways: 1. with no preprocessing of the input images; 2. using a Blackman window; 3. using the image gradient; 4. using the image gradient and a Blackman window.

4. RESULTS

For each one of the test images, five sets of 289 images with different and known subpixel shifts were generated, one per each level of aliasing allowed.

As mentioned in section 3, each one of the five algorithms used here were tested in four different modes. Then, after selecting the reference image, for example the first one in a given set of images, the algorithms are applied to the same set of images in the same order, giving as a result two vectors of size 289 that contain the subpixel shift values found for the horizontal (X) and vertical (Y) directions for every image in the set. This two vectors can be compared to the one that contains the exact values for the real translation in order to obtain a figure of merit, which, as an example, can be done visually in a plot like the one shown in figure 2, where the black circles are centered in the real shift values whereas the different points represent the shift estimates obtained from different algorithms. Error estimates from the comparison are usually made separately for the X and Y coordinates, calculating the differences in each axis, but it would be better to




Figure 2: Three PC methods compared to the real shifts for the Aerial View at 100% of aliasing : Real Shifts (black circles); UDFT (green squares); 2DGF (red triangles); LSPD (blue dots).

include the overall error for a given estimation, in a similar way as it is presented in [13].

In this way, we calculated the estimation error as the Euclidean distance between each estimation and the real shift. For a set of images, the mean and standard deviation of the distance is calculated for all the available shifts, and they are finally averaged. For example, the results for all the methods after incorporating the two preprocessing steps of windowing and gradient are displayed in table 1 for the Resolution Chart image of figure 1a). The first thing we can mention is that when no aliasing is present all methods have their best performance, and between all of them the LSPD is the most accurate. All of them, except the FSPE, present results that look like an uniform grid, with no visible differences with the real shifts, and that lie within 1/100 of a pixel and even less for the LSPD. As soon as the aliasing starts to appear, the performance is overall degraded, remaining within 1/50 of a pixel in all methods except the FSPE. When the aliasing is total (100%), all methods except again the FSPE show more similar results, but only the Gaussian fit based methods 1DGF and 2DGF have still an acceptable performance of 0.05 pixel accuracy, that is expanded up to 0.1 and 0.15 pixels for the 200% and 400% aliasing case respectively.

A similar tendency is also seen in the results of table 2, where the same methods are applied to the other test image (Aerial View), again including the two preprocessing steps. The performance with no aliasing is even better now, within 1/200 of a pixel, which can be explained because of the data dependency of the PC algorithm due to a different frequency content of the original image used. The performance is degraded with partial aliasing, but remains within 1/100 of a pixel for all algorithms except the FSPE. With total aliasing the mean accuracy is about 0.03 pixels for all methods ex-

Table 1: Mean and Standard Deviation of the Shift Estimation Error for the Resolution Chart Image.

Aliasing	0%	50%	100%	200%	400%
UDFT	0.0043 ± 0.0022	0.0137 ± 0.0061	0.0529 ± 0.0198	0.1242 ± 0.0453	0.1574 ± 0.0663
FSPE	0.0412 ± 0.1119	0.0502 ± 0.1322	0.0683 ± 0.1594	0.1297 ± 0.1628	0.4373 ± 0.4996
1DGF	0.0087 ± 0.0041	0.0133 ± 0.0059	0.0334 ± 0.0144	0.0784 ± 0.0336	0.0926 ± 0.0470
2DGF	0.0062 ± 0.0030	0.0138 ± 0.0063	0.0377 ± 0.0155	0.0887 ± 0.0352	0.0940 ± 0.0525
LSPD	0.0021 ± 0.0016	0.0088 ± 0.0051	0.0430 ± 0.0240	0.1951 ± 0.1031	0.3680 ± 0.1541

Table 2: Mean and Standard Deviation of the Shift Estimation Error for the Aerial View Image.

Aliasing	0%	50%	100%	200%	400%
UDFT	0.0038 ± 0.0018	0.0054 ± 0.0028	0.0195 ± 0.0082	0.0540 ± 0.0207	0.0691 ± 0.0276
FSPE	0.0266 ± 0.0427	0.0291 ± 0.0374	0.0404 ± 0.0566	0.0533 ± 0.0419	0.0698 ± 0.0474
1DGF	0.0076 ± 0.0036	0.0111 ± 0.0042	0.0202 ± 0.0084	0.0398 ± 0.0194	0.0448 ± 0.0211
2DGF	0.0034 ± 0.0015	0.0064 ± 0.0027	0.0145 ± 0.0055	0.0288 ± 0.0133	0.0319 ± 0.0160
LSPD	0.0016 ± 0.0017	0.0026 ± 0.0015	0.0076 ± 0.0043	0.0299 ± 0.0173	0.0722 ± 0.0421

cept the FSPE, but the LSPD is still the most accurate, which can be confirmed by inspecting the graphical comparison in figure 2, where the LSPD related dots never leave the black circle. As the aliasing started to rise, again the Gaussian fit methods outperform the others, being the 2DGF the best in the worst scenario with an error below 0.05 pixels.

The reason we only display the results that were obtained by using the two preprocessing steps is because they allow all the algorithms to perform better. The windowing of the images always improved the performance, but the gradient only improved the performance when the aliasing was larger. However, when both preprocessing were used together, all methods improved their performance more than using windowing alone for example. Anyway, in order to summarize the results, we present in table 3 the averaged results obtained with every method using all the different preprocessing configurations and for both images used. In this way we can verify what algorithm is the most reliable under diverse aliasing conditions, with or without any preprocessing, or using any type of image. From no aliasing to partial aliasing conditions, the LSPD algorithm is the most accurate, which is maybe related to the fact that it was tuned to handle up to 50% of aliasing. From total aliasing on so forth, the Gaussian fit methods are more stable in terms of the degradation of its accuracy, being the 2DGF the most accurate overall. This might be explained by the fact that there is less degree of freedom, and thus more robustness when doing the least square fit by using both axis instead of separately.

The results obtained can be better perceived in the scatter plots of figure 3, where for the 200% aliasing Resolution Chart image, UDFT and 2DGF still keep the grid shape, but less uniform, however for the LSPD the grid is changing its overall shape. Moreover, when the aliasing is at 400%, the 2DGF presents a more uniform distribution of the estimated shifts than UDFT, while the LSPD shape is completely deformed.

A similar comparison can be performed looking back at figure 2, where all the three methods shown, UDFT, 2DGF and LSPD, have results that mostly lie within the size of the black circle, centered at the real shifts. This sort of result is the expected one for a good method under very little aliasing, which in this case was obtained by using the Aerial

View test image at 100% aliased conditions, which at the end not necessarily indicates that the image have a total aliased condition because maybe the original image is oversampled, or have very little high frequency content. This is not the case with the Resolution Chart image where a similar performance is obtained for 50% aliasing, as shown when comparing tables 1 and 2.

Finally, and in order to show the incidence of using or not the preprocessing steps, we present a graphical comparison in figure 4, at least for the UDFT method where both preprocessing have noticeable different effects, but no one can improve the result obtained when using both together as in figure 3a).

5. CONCLUSION

In this paper, an evaluation of the subpixel registration accuracy using different phase correlation based methods under aliased conditions was carried on. Results show that all methods suffer from aliasing in different proportions, but the ones based on the one and two dimensional Gaussian fit of the phase correlation function are the less affected by higher degrees of aliasing. This result is important, for example, when applying phase correlation registration methods for performing superresolution on undersampled imagery. Further work may include comparing more registration techniques for a large collection of aliased images, and develop a method for measuring the aliasing effects for other motion models rather than pure translational.

REFERENCES

- [1] B. Zitova and J. Flusser, "Image registration methods: a survey," *Image and Vision Computing*, vol. 21, no. 11, pp. 977–1000, October 2003.
- [2] C. D. Kuglin and D. C. Hines, "The phase correlation image alignment method," in *Proc. Int. Conference on Cybernetics and Society*, 1975, pp. 163–165.
- [3] B. Reddy and B. Chatterji, "An fft-based technique for translation, rotation, and scale-invariant image registration," *IEEE Trans. on Image Processing*, vol. 5, no. 8, pp. 1266–1271, 1996.

Table 3: Overall Mean and Standard Deviation of the Shift Estimation Error for all methods using both images and with all the different mixes of preprocessing (0% Aliasing not shown due to space limitations)

Aliasing	50%	100%	200%	400%	Average
UDFT	0.0484 ± 0.0146	0.0508 ± 0.0206	0.0953 ± 0.0387	0.1182 ± 0.0544	0.0768 ± 0.0306
FSPE	0.0633 ± 0.0643	0.0624 ± 0.0790	0.1021 ± 0.1112	0.2916 ± 0.2561	0.1225 ± 0.1181
1DGF	0.0266 ± 0.0117	0.0309 ± 0.0145	0.0597 ± 0.0266	0.0909 ± 0.0573	0.0495 ± 0.0253
2DGF	0.0294 ± 0.0102	0.0309 ± 0.0117	0.0553 ± 0.0242	0.0733 ± 0.0456	0.0466 ± 0.0218
LSPD	0.0207 ± 0.0123	0.0386 ± 0.0197	0.1268 ± 0.0687	0.2471 ± 0.1150	0.0911 ± 0.0459

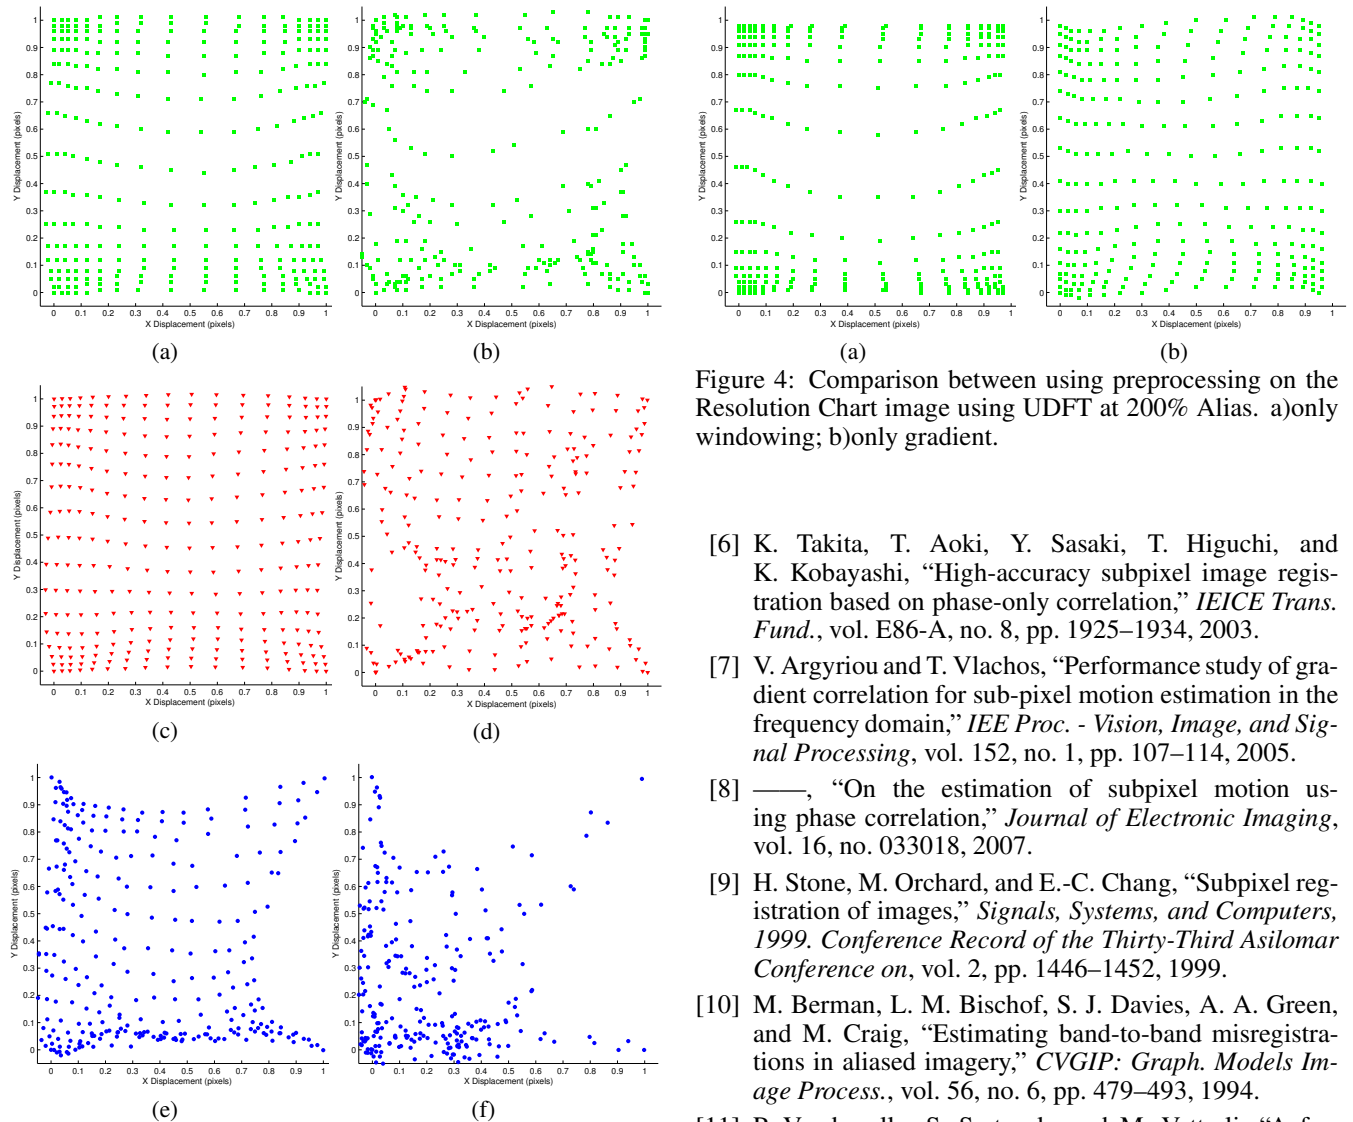


Figure 3: Results for the Resolution Chart image. a)UDFT 200% Alias and b)400% Alias; c)2DGF 200% Alias and d)400% Alias; e)LSPD 200% Alias and f)400% Alias.

- [4] B. Marcel, M. Briot, and R. Murrieta, "Calcul de translation et rotation par la transformation de fourier," *Traitement du Signal*, vol. 14, no. 2, pp. 135–149, 1997.
- [5] H. Foroosh, J. Zerubia, and M. Berthod, "Extension of phase correlation to subpixel registration," *IEEE Trans. on Image Processing*, vol. 11, no. 3, pp. 188–200, 2002.

Figure 4: Comparison between using preprocessing on the Resolution Chart image using UDFT at 200% Alias. a)only windowing; b)only gradient.

- [6] K. Takita, T. Aoki, Y. Sasaki, T. Higuchi, and K. Kobayashi, "High-accuracy subpixel image registration based on phase-only correlation," *IEICE Trans. Fund.*, vol. E86-A, no. 8, pp. 1925–1934, 2003.
- [7] V. Argyriou and T. Vlachos, "Performance study of gradient correlation for sub-pixel motion estimation in the frequency domain," *IEE Proc. - Vision, Image, and Signal Processing*, vol. 152, no. 1, pp. 107–114, 2005.
- [8] —, "On the estimation of subpixel motion using phase correlation," *Journal of Electronic Imaging*, vol. 16, no. 033018, 2007.
- [9] H. Stone, M. Orchard, and E.-C. Chang, "Subpixel registration of images," *Signals, Systems, and Computers, 1999. Conference Record of the Thirty-Third Asilomar Conference on*, vol. 2, pp. 1446–1452, 1999.
- [10] M. Berman, L. M. Bischof, S. J. Davies, A. A. Green, and M. Craig, "Estimating band-to-band misregistrations in aliased imagery," *CVGIP: Graph. Models Image Process.*, vol. 56, no. 6, pp. 479–493, 1994.
- [11] P. Vandewalle, S. Ssstrunk, and M. Vetterli, "A frequency domain approach to registration of aliased images with application to super-resolution," *EURASIP Journal on Applied Signal Processing*, vol. 2006, pp. Article ID 71 459, 14 pages, 2006.
- [12] M. Guizar-Sicairos, S. T. Thurman, and J. R. Fienup, "Efficient subpixel image registration algorithms," *Optics Letters*, vol. 33, no. 2, pp. 156–158, 2008.
- [13] F. Humblot, B. Collin, and A. Mohammad-Djafari, "Evaluation and practical issues of subpixel image registration using phase correlation methods," in *Int. Conf. on Physics in Signal and Im. Proc.*, 2005, pp. 115–120.