Analysis of image registration noise due to rotationally dependent aliasing[2]

Harold S. Stone[†] and Bo Tao[‡] and Morgan McGuire[*]

[†] *NEC Research Institute; and* [‡] *Streaming21; and* [*] *Brown University*

E-mail: hstone@research.nj.nec.com

This paper investigates factors that degrade the precision of image registration based on phase correlation. The major sources of error are interpolation error and rotationally dependent aliasing. The latter error stems from the fact that the discrete-Fourier transform does not commute with the rotation of sampled-images, whereas in the continuous domain the corresponding operations do commute. We show through a series of examples how much the various sources of error contribute to phase-correlation registration, and we demonstrate constructive techniques for improving precision and signal to noise ratio in the registration process.

Since rotationally dependent aliasing is exacerbated by the presence of high frequencies, the examples demonstrate that the use of a Blackman window removes spurious high frequencies in the spectral leakage created by the image boundary and greatly reduces aliasing effects. Since remaining aliasing effects are strongest in the low frequencies of the Fourier transform, their affects can be reduced to a negligible amount by removing frequencies within a radius of $N/4$ of the Fourier domain origin. A third technique is to perform phase correlation over half the Fourier plane rather than over the full plane, which more than doubles the signal-to-noise ratio of phase correlation. For an example image, the combination of techniques improved the signal-to-noise ratio from 6.6 to 115.9 and raised the phase-correlation peak from 0.300 to 0.835, which are substantially higher values than previously reported.

Key Words: aliasing, image searching, image registration, Fourier transform, scale invariance, translation invariance, rotation invariance

1. INTRODUCTION

Image registration is the process of aligning two similar images of the same scene so that points from one scene lie in the same positions as corresponding points in the other. For alignments involving only rotation, translation, and changes of scale, phase-correlation techniques reported in the image-registration literature have been shown to be quite powerful [2, 1, 4, 5, 6, 7, 8, 10, 13, 14, 17]. Phase correlation typically has excellent signal-to-noise ratio (SNR), and produces very sharp correlation peaks. Moreover, scale factors and rotational angles can be discovered independently of translation, thereby greatly reducing the complexity of registering with respect to translation, rotation, and change-of-scale concurrently [6]. With respect to translaton-only registration, a noniterative phase-correlation algorithm achieved an average absolute error in position of less than 0.01 pixels [15], and is competitive in precision with the leading iterative pixel-domain interpolation scheme[16].

To understand phase-correlation registration, recall that the Fourier representation of a 2D image contains complex values at each point in the spatial frequency domain. A translation of the original image changes only the phase angle of the complex Fourier coefficients and has no effect on the magnitude of those coefficients. Hence, the Fourier phase by itself gives the translation parameters, and can be used to register images [10]. A rotation of the image about any point rotates the Fourier representation (both phase and magnitude) around the axes origin. Therefore, by comparing the Fourier-transform magnitude of a translated and rotated image to the Fourier-transform magnitude of the original image, one can determine the rotational difference, undo the rotation, and then find the translational difference between the images. A similar generalization of this idea allows one to find scale-factor differences by observing quantities that are not affected by translations or rotations [6, 14, 17]. This process has been given different names in the literature, most notably FMI-SPOMF (Fourier-Mellin invariant, symmetric phase-only matched filter) in [6]. These invariants and their detection through phase corre-

lation have been useful in watermarking of images that is robust with respect to scale and rotational changes of images [18].

Our concern in this paper is the basic precision of the FMI-SPOMF process and its related techniques. The process rests strongly on the premise that the Fourier transform of a rotated image is equal to the rotated transform of the image. In this paper we show that premise is actually not true in a fundamental way for finite discrete images. The discrete Fourier transform of a rotated sampled image differs from the rotated discrete Fourier transform of the same image by artifacts caused by rotationally dependent aliasing that are in addition to any artifacts produced by interpolation error within the rotation. The study in this paper reveals that errors from rotationally dependent aliasing can be rather substantial if not surpressed, and that they can be controlled by standard windowing and filtering techniques. We show that a finely tuned phase-correlation algorithm produces excellent signal-to-noise ratio with virtually no artifacts from rotationally dependendent artifacts. The phase-correlation peak in the example is an order of magnitude higher than the peak reported in [14] for a phase correlation under both scale and rotational changes.

Lucchese *et al.* [11] report a variation of phase correlation in which they improve precision by projecting the 2D Fourier domain radially into a 1D Fourier domain that varies only with rotational angle. Their phase correlation function, prior to a final refinement, has a false peak at the origin with a height about 0.61, and a correct peak with a height of about 0.25. To find the correct registration to high precision, they refine the initial answer by taking into account the process possible presence of one or more false peaks. The algorithms backs out the rotations corresponding to each phase-correlation peak in seeking the rotation that best explains the observations. Our work shows that the false peak at 0 degrees rotation is likely to be due to rotational dependent aliasing.

Section 2 of this paper contains the mathematical background and shows the presence of rotationally dependent aliasing. Section 3 reviews phase correlation and normalized correlation, and explains the effect of false peaks on the height of a correct phase-correlation peak. The experimental data appear in Section 4. Section 5 contains a summary and conclusions.

2. FOURIER TRANSFORMS AND IMAGE ROTATIONS

In this section, we examine the commutativity of Fourier transforms and image rotation. We show that commutativity fails for finite sampled images, although it holds for infinite continuous images. The discussion considers three cases — infinite continuous images, finite continuous images, and finite sampled images.

2.1. Infinite continuous case

This case is well-known and is included here because it forms the basis of analyses below.

Let $f(x, y)$ be a two-dimensional image of infinite extent defined on a continuous space. Its Fourier transform $F(\omega, \nu)[f(x, y)]$ is:

$$F(\omega, \nu)[f(x, y)] = \left(\frac{1}{K}\right) \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(x, y) e^{-j(\omega x + \nu y)} dx dy \tag{1}$$

where K is a normalizing constant. To simplify notation, we use vectors to represent the coordinate pairs in the pixel and Fourier domains. Let $\mathbf{x}$ to denote the column vector $(x, y)^t$, where $\mathbf{x}^t$ denotes the transpose of $\mathbf{x}$, and let $\mathbf{w}$ to denote the row vector (ω, ν). Then we can rewrite Eq. (1) as

$$F(\mathbf{w})[f(\mathbf{x})] = \left(\frac{1}{K}\right) \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(\mathbf{x}) e^{-j\mathbf{w}\mathbf{x}} d\mathbf{x} \tag{2}$$

Let $\mathbf{R}_\theta$ be the rotation matrix

$$\mathbf{R}_\theta = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}. \tag{3}$$

The column vector $\mathbf{x}_\theta = \mathbf{R}_\theta \mathbf{x}$ is $\mathbf{x}$ rotated in the positive (counter clockwise) direction by θ degrees. Rotating an image $f(\mathbf{x})$ in the positive θ direction produces

the image $f(\mathbf{x}_{-\theta})$. Note that $\mathbf{w}_{-\theta} = \mathbf{w}\mathbf{R}_\theta$ because $\mathbf{R}_\theta^{-1} = \mathbf{R}_\theta^t$. Also, $\mathbf{w}\mathbf{x}_\theta = \mathbf{w}\mathbf{R}_\theta\mathbf{x} = \mathbf{w}_{-\theta}\mathbf{x}$. The following proof shows that the Fourier transform commutes with image rotation in the infinite continuous-image case.

$$\begin{aligned} F(\mathbf{w})[f(\mathbf{x}_{-\theta})] &= \left(\frac{1}{K}\right)\int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} f(\mathbf{x}_{-\theta})e^{-j\mathbf{w}\mathbf{x}}d\mathbf{x} \\ &= \left(\frac{1}{K}\right)\int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} f(\mathbf{x})e^{-j\mathbf{w}\mathbf{x}_\theta}d\mathbf{x}_\theta \\ &= \left(\frac{1}{K}\right)\int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} f(\mathbf{x})e^{-j\mathbf{w}_{-\theta}\mathbf{x}}d\mathbf{x} \\ &= F(\mathbf{w}_{-\theta})[f(\mathbf{x})] \end{aligned} \tag{4}$$

The second to last line uses the fact that $d\mathbf{x}_\theta = d\mathbf{x}$ because $\mathbf{R}_\theta$ is a unitary matrix.

2.2. Finite continuous case

In this section we show that the discrete Fourier series transform and image rotation commute to within interpolation error when images are continuous and finite in extent. The interpolation error arises in the Fourier domain.

For this derivation, we assume that $f(\mathbf{x})$ is defined within a square of size $T \times T$ centered at the origin. We use the discrete Fourier series transform of $f(\mathbf{x})$ in this case, which is defined on points that are integer multiples of $\omega_0 = 2\pi/T$. Let $\mathbf{n} = (n, m)$ for integers n and m, $-\infty < n < \infty$, $-\infty < m < \infty$. The discrete Fourier series transform of $f(\mathbf{x})$ is given by

$$F(\mathbf{n})[f(\mathbf{x})] = \left(\frac{1}{K_1}\right)\int_{-T/2}^{+T/2}\int_{-T/2}^{+T/2} f(\mathbf{x})e^{-j\omega_0\mathbf{n}\mathbf{x}}d\mathbf{x} \tag{5}$$

The inverse transform reconstructs $f(\mathbf{x})$ from the basis functions $e^{j\omega_0\mathbf{n}}$ for values of $\mathbf{n}$ whose components range over all the integers.

$$f(\mathbf{x}) = \mathbf{F}^{-1}\left(\mathbf{F}(\mathbf{n})[f(\mathbf{z})]\right) = \left(\frac{1}{K_2}\right)\left(\sum_n\sum_m F(\mathbf{n})[f(\mathbf{z})]\right)e^{j\omega_0\mathbf{n}\mathbf{x}} \tag{6}$$

Constants K_1 and K_2 are normalizing constants. Note that Eq. (6) is periodic with period 2π in each coordinate of $\mathbf{x}$. Hence, the reconstruction described by Eq. (6) tiles the plane with copies of $f(\mathbf{x})$, and the tiling is aligned with the $\mathbf{x}$-axes. This

is well known and is exemplified by Fig. 1. Figure 1(a) shows a finite continuous image. When we reconstruct the image from its discrete Fourier series, we obtain the infinite tiled plane in Figure 1(b) (with only 9 of the tiles depicted).

To avoid effects caused by rotating nonzero pixels out of the image rectangle, we now assume that $f(\mathbf{x}) = 0$ for $\mathbf{x}$ outside of the circle $\| \mathbf{x} \| > T/2$. The discrete Fourier series appears to commute with image rotation mathematically because this case seems to be parallel to Eq. (4).

$$\begin{aligned} F(\mathbf{n})[f(\mathbf{x}_{-\theta})] &= \left(\frac{1}{K_1}\right) \int_{\|\mathbf{x}\| \le T/2} f(\mathbf{x}_{-\theta}) e^{-j\omega_0 \mathbf{n}\mathbf{x}} d\mathbf{x} \\ &= \left(\frac{1}{K_1}\right) \int_{\|\mathbf{x}\| \le T/2} f(\mathbf{x}) e^{-j\omega_0 \mathbf{n}\mathbf{x}_{\theta}} d\mathbf{x}_{\theta} \\ &= \left(\frac{1}{K_1}\right) \int_{\|\mathbf{x}\| \le T/2} f(\mathbf{x}) e^{-j\omega_0 \mathbf{n}_{-\theta}\mathbf{x}} d\mathbf{x} \\ &= F(\mathbf{n}_{-\theta})[f(\mathbf{x})] \end{aligned} \tag{7}$$

But the operations actually do not commute. The derivation in Eq. (7) depends on a continuous space $\mathbf{n}$. The reconstruction, however, requires that the $\mathbf{n}$ be discrete samples. When we invert the Fourier series transform in Eq. (7) using basis functions $e^{j\omega_0 \mathbf{n}}$, we obtain the reconstruction

$$\left(\frac{1}{K_2}\right) \left(\sum \sum F(\mathbf{n})[f(\mathbf{z}_{-\theta})] \right) e^{j\omega_0 \mathbf{n}\mathbf{x}} = f(\mathbf{x}_{-\theta}), \tag{8}$$

which is a rotated tile in the original basis, as shown in Fig. 3. The reconstruction basis functions $e^{j\omega_0 \mathbf{n}}$ are orthogonal with respect to the original basis, but with respect to the rotated basis. To reconstruct the function from the rotated-basis Fourier series, we must use the rotated basis functions $e^{j\omega_0 \mathbf{n}_{-\theta}}$. These give rise to the reconstruction

$$\left(\frac{1}{K_2}\right) \left(\sum \sum F(\mathbf{n}_{-\theta})[f(\mathbf{z})] \right) e^{j\omega_0 \mathbf{n}_{-\theta}\mathbf{x}} = f(\mathbf{x}), \tag{9}$$

which is the original tile in a rotated basis. This reconstruction appears in Fig. 4. Note that the reconstructed images in both Figs. 3 and 4 are periodic in T, with T aligned with the basis functions used for reconstruction.

Although commutativity does not hold as an exact equality, it holds to within interpolation in the following sense. Rotate the Fourier series of an image. Interpolate the samples $F(\mathbf{n}_{-\theta})$ from the points lying on the grid $\mathbf{n}_{-\theta}$ to samples $F(\mathbf{n})$ that lie on the grid $\mathbf{n}$. Invert $F(\mathbf{n})$ using Eq. (8). This produces a tiled plane of rotated images to within interpolation error. There is no interpolation error when θ is a multiple of 90 degrees because the grids $\mathbf{n}$ and $\mathbf{n}_{-\theta}$ are identical. Figures 3 and 4 confirm this because they are identical to each other for these rotations.

Interpolation error alone explains the failure of commutativity. There is no error from aliasing because the Fourier series coefficients are not periodic. The discrete Fourier series has an infinite number of coefficients. Its inverse, a continuous image, has an infinite number of points in pixel space.

2.3. Finite discrete case

Commutativity fails in a fundamental way for discrete Fourier transforms because aliasing effects depend on rotation. Different rotations of an image alias differently in the discrete Fourier domain.

For the discrete case, we let $\mathbf{x} = (x, y)^t$ range over N^2 grid points such that $-N/2 \leq x < N/2;\ -N/2 \leq y < N/2$, where N is even. Similarly, $\mathbf{n} = (n, m)$ ranges over N^2 grid points such that $-N/2 \leq n < N/2;\ -N/2 \leq m < N/2$. The discrete Fourier transform $G(\mathbf{n})[f(\mathbf{x})]$ is given by

$$G(\mathbf{n})[f(\mathbf{x})] = \left(\frac{1}{K}\right) \sum_{\mathbf{x}} f(\mathbf{x}) e^{-j\omega_0 \mathbf{n}\mathbf{x}} \tag{10}$$

where $\omega_0 = 2\pi/N$. Note that $G(\mathbf{n})$ is periodic in N if we extend the function from the original tile with N^2 points to the full infinite plane of discrete grid points. Frequencies outside the range $-N/2 \leq n < N/2$ map modulo N into the original range, and thereby produce aliasing effects associated with sampling that are often observed in discrete Fourier transforms [12]. The reconstruction of the original function with basis functions $e^{-j\omega_0 \mathbf{n}}$ produces a periodic tiled plane. Consequently,

Eq. (10) is the inverse transform of a periodic function whose tiles are aligned with the $\mathbf{x}$ axes.

As done earlier, we assume that the image function is zero outside a circle of radius $N/2$ in order to remove discrepancies caused by nonzero pixels being rotated out of the image tile. We can rewrite Eq. (10) in a form similar to Eq. (2) by using the function $comb(x, y) = \sum_{r=-\infty}^{\infty} \sum_{s=-\infty}^{\infty} \delta(x - r, y - s)$. (In the literature, this is also called the "bed of nails" function and "impulse sheet" for 2D functions, and the "impulse train" function for 1D functions.) Letting $\mathbf{x}$ range over the continuous plane instead of discrete grid points, we can rewrite Eq. (10) as:

$$G(\mathbf{n})[f(\mathbf{x})] = \left(\frac{1}{K}\right) \int_{\|\mathbf{x}\| \leq N/2} comb(\mathbf{x}) f(\mathbf{x}) e^{-j\omega_0 \mathbf{nx}} d\mathbf{x}. \tag{11}$$

From Eq. (11) we can develop equations for the rotated transform of an image and for the transform of a rotated image. The rotated transform is:

$$\begin{aligned} G(\mathbf{n}_{-\theta})[f(\mathbf{x})] &= \left(\frac{1}{K}\right) \int_{\|\mathbf{x}\| \leq N/2} comb(\mathbf{x}) f(\mathbf{x}) e^{-j\omega_0 \mathbf{n}_{-\theta} \mathbf{x}} d\mathbf{x} \\ &= \left(\frac{1}{K}\right) \int_{\|\mathbf{x}\| \leq N/2} comb(\mathbf{x}_{-\theta}) f(\mathbf{x}_{-\theta}) e^{-j\omega_0 \mathbf{nx}} d\mathbf{x}. \end{aligned} \tag{12}$$

The key observation is that the change of variables rotates the *comb* function relative to the original axes. Multiplication in the pixel domain becomes convolution in the Fourier domain, which yields:

$$\begin{aligned} G(\mathbf{n}_{-\theta})[f(\mathbf{x})] &= \left(\frac{1}{K}\right) F(\mathbf{n})[comb(\mathbf{x}_{-\theta})] \otimes F(\mathbf{n})[f(\mathbf{x}_{-\theta})] \\ &= \left(\frac{1}{K}\right) comb(\mathbf{n}_{-\theta}) \otimes F(\mathbf{n})[f(\mathbf{x}_{-\theta})]. \end{aligned} \tag{13}$$

where "$\otimes$" denotes convolution. Eq. (13) uses the fact that the Fourier transform of $comb(\mathbf{x})$ is $comb(\mathbf{n})$ [3].

The transform of the rotated image is:

$$\begin{aligned} G(\mathbf{n})[f(\mathbf{x}_{-\theta})] &= \left(\frac{1}{K}\right) \int_{\|\mathbf{x}\| \leq N/2} comb(\mathbf{x}) f(\mathbf{x}_{-\theta}) e^{-j\omega_0 \mathbf{nx}} d\mathbf{x} \\ &= \left(\frac{1}{K}\right) comb(\mathbf{n}) \otimes F(\mathbf{n})[f(\mathbf{x}_{-\theta})]. \end{aligned} \tag{14}$$

Eqs. (13) and (14) obviously differ in the orientation of the *comb* function. Because the *comb* function determines the aliasing axes, the two equations have different aliasing effects. In essence, the order of aliasing and rotation is reversed in the two equations, and these operations do not commute. The aliasing in Eq. (13) is equivalent to collapsing frequencies of $F(\mathbf{n})[f(\mathbf{x})]$ on the original axes, followed by a rotation of the axes. The aliasing and rotation in Eq. (14) occur in a reversed order. The Fourier transform rotates first (to create $F(\mathbf{n})[f(\mathbf{x}_{-\theta})]$), and then collapses along the original axes. Thus, the aliasing collapses frequencies along axes that are aligned differently with respect to the Fourier transform axes.

Examples of rotationally dependent aliasing appear in Figs. 5 through 9. Figure 5 shows the discrete Fourier transform of Fig. 1(a). The transform plot is a log magnitude plot in which large magnitudes are dark and small magnitudes are light. We use log magnitude in all of the transform plots to show more clearly the structure of the transform at off-peak coordinates.

Figure 6 is the discrete Fourier transform of Fig. 2. which is a rotation of Fig. 1. Note that the principal components of Fig. 6 are rotated axes that appear in Fig. 5, but Fig. 6 also contains off-diagonal lines not present in Fig. 5. These lines are artifacts of rotationally dependent aliasing. Figure 7 explains their source. Aliasing superimposes the 25 tiles of the infinite transform space in Fig. 7(a) into a single tile of Fig. 7(b) by collapsing coefficients modulo N along rotated axes. It is clear that the off-diagonal lines in Fig. 7(b) lie exactly where we see off-diagonal lines in Fig. 6. Aliasing also exists in Fig. 5, but the aliasing in this case collapses the vertical and horizontal axes onto themselves along the original axes.

Figure 8(a) is the rotated transform of Fig. 5, with 0s rotated into the transform from outside the tile. The inverse of this transform is a periodic tiled plane. Inversion of Fig. 8(a) with the basis functions $e^{j\omega_0\mathbf{n}_{-\theta}}$, aligns tiles to the rotated axes, and each tile contains the original figure. However, image registrations invert transforms with basis functions $e^{j\omega_0\mathbf{n}}$, and produce a periodic tiled plane whose tiles

are aligned to the original axes. The tiles in this reconstruction exhibit aliasing in the pixel domain not present in the original image. The aliasing is evident in Figure 8(b), the inverse transform of Fig. 8(a). The aliasing in Fig. 8(b) is explained by Fig. 9, which shows a rotated tiled plane superimposed on a plane tiled along the original axis orientation. The rotated transform is a transform of the rotated tiling. Inverting with respect to the original axes, is equivalent to collapsing the rotated tiled plane along the original axes into a single tile, which produces the tile in Fig. 9(b). Observe that this is the outline of the image in Fig. 8(b).

These examples confirm the mathematics of the aliasing. Because aliasing is rotationally dependent, correlations of rotated Fourier magnitudes of images have an inherent precision loss.

3. NORMALIZED CORRELATION AND PHASE CORRELATION

The data in the next section reveal that phase correlation produces a correlation peak that is relatively low compared to the ideal peak of 1.00, and is much lower than the peak produced by normalized correlation computations. The data also show that rotationally dependent aliasing produces a false correlation peak as well as the correct peak. The false peak may be taller than the peak at the correct registration position. This section reviews the notation and definitions for normalized correlation and phase correlation, and explains why the presence of a false peak reduces the height of the peak at the correct position.

For simplicity, we use a one-dimensional example. The discussion generalizes trivially to two-dimensional images. Let $\mathbf{x}$ be a column vector of length N. Let $\mathbf{F}$ be a Fourier-transform matrix such that $\mathbf{X} = \mathbf{F}\mathbf{x}$ is the Fourier transform of $\mathbf{x}$. The (n, m)th element of $\mathbf{F}$ is $f_{n,m} = W^{nm}$ where $W = e^{-j2\pi/N}$ is an Nth root of unity. The nth row of $\mathbf{F}$ consists of successive powers of W^n.

Let $\mathbf{x}_1$ and $\mathbf{x}_2$ be real valued vectors of length N. The *normalized correlation* of $\mathbf{x}_1$ and $\mathbf{x}_2$ is defined to be

$$\frac{(\mathbf{x}_1^t \mathbf{x}_2)/N - \bar{\mathbf{x}_1}\bar{\mathbf{x}_2}}{\sqrt{\mathrm{Var}(\mathbf{x}_1)\mathrm{Var}(\mathbf{x}_2)}} \tag{15}$$

where $\bar{\mathbf{x}}$ is the mean of $\mathbf{x}$, and $\mathrm{Var}(\mathbf{x})$ is the variance of $\mathbf{x}$. The magnitude of normalized correlation cannot exceed unity, and the absolute value of a correlation is a measure of the quality of the correlation. Let $\mathbf{X}_1$ and $\mathbf{X}_2$ be column-vector Fourier transforms of vectors $\mathbf{x}_1$ and $\mathbf{x}_2$. The *Phase correlation* of $\mathbf{X}_1$ and $\mathbf{X}_2$ is the inverse Fourier transform of the vector

$$\left(\mathbf{X}_1 .* \hat{\mathbf{X}}_2\right) ./ \left(\| \mathbf{X}_1 \| .* \| \mathbf{X}_2 \|\right), \tag{16}$$

where the operation ".∗" is pointwise multiplication in Matlab notation, "./" denotes pointwise division, $\hat{\mathbf{X}}$ is the complex conjugate of $\mathbf{X}$, and $\| \mathbf{X} \|$ denotes the vector of magnitudes of the complex components of $\mathbf{X}$. In theory, some denominators of Eq. (16) can be identically zero, thereby leading to undefined coefficients. But in practice, this is extremely rare because of the presence of noise and small perturbations. If $\mathbf{x}_1$ and $\mathbf{x}_2$ differ only by cyclic translation by an integral number of pixels in the signal domain, then $\mathbf{X}_1$ and $\mathbf{X}_2$ differ only in the phase of their complex coefficients. The inverse Fourier transform of Eq. (16) in this case is an impulse at position n in the signal domain.

Under ideal circumstances, the height of a phase-correlation peak in the signal domain is unity. This follows because the phase correlation function should be a single peak at the the position that corresponds to the correct displacement and should be zero otherwise. In the Fourier domain, the DC coefficient is 1 because of normalization, and this is equal to the average of the time domain representation of the same function. Hence, the ideal peak height in the time domain has to be unity to meet the DC constraint. If there are two or more peaks in the phase-correlation function in the time domain, the sums of those peaks must be unity.

Hence, the presence of false peaks in the time domain, must diminish the value of the phase-correlation peak at the correct displacement position.

It is well-known that the Fourier transforms of real images have conjugate symmetry, so that the radial representation of the Fourier transform magnitude has period 2 as the rotation angle moves through 2π radians. If we phase-correlate both periods of the Fourier magnitude, the phase correlation produces two copies of the same function spaced π radians apart. Since the sum of the peaks is 1.00, no peak has a height greater than 0.50. By phase correlating only half of the radial representation of the Fourier magnitude, the maximum height of a peak doubles, and noise power diminishes. This technique was used in [11] and we use it in the remainder of this paper.

The data in the next section show examples of false peaks and give a sense of how much rotationally dependent aliasing can lower the height of peaks at correct positions.

4. SOURCES OF REGISTRATION ERROR

Given that the Fourier transform does not commute with rotation, the next question is to determine how much noncommutativity degrades phase correlation. In this section we take a brief look at errors due to the following processes:

1. interpolation error in the pixel and Fourier domains,
2. noncommutativity of rotation and discrete Fourier transform,
3. low frequencies in the discrete Fourier transform,
4. high frequencies in the discrete Fourier transform,
5. conversion from polar to rectangular coordinates, and
6. use of phase correlation versus use of normalized correlation.

In the remainder of this section, the term "correlation" by itself refers to normalized correlation. Phase correlation values are reported for correlations of 180-degree transforms as mentioned earlier.

The sample figure used in this experiment is an aerial photograph of an agricultural scene shown in Fig. 10(a). Note that the image is cropped to a centered circle that lies within the image rectangle. This allows us to correlate pairs of circular rotations of the image within the rectangle without aligning image pixels to background pixels or nonexistent pixels.

The image in Fig. 10(a) has sharp boundaries around the circumference, which produce spectral leakage in the spatial frequency domain. The spurious energy tends to concentrate in the high frequencies which, in turn, tends to produce strong rotationally dependent aliasing artifacts. It is well known that windows of various types can reduce or eliminate the spectral leakage caused by boundary effects [9]. Figure 10(b) shows the example figure after windowing with a Blackman window [9]. Note that the center of the image in Fig. 10(b) is normal and the image fades to black gradually as distance increases away from the image center. In this example we expect the windowed image to produce a negligible amount of precision loss from aliasing and expect that the unwindowed image of Fig. 10(a) will be highly aliased and exhibit a larger loss of registration precision.

The first experiment is a vivid demonstration of rotationally dependent aliasing. This experiment rotates the sample image by multiples of 15 degrees in the pixel plane, transforms to the Fourier domain, derotates back to the original orientation, and does an inverse transform. The results in Fig. 11 show the angular dependency on aliasing, and also shows that no aliasing occurs for rotations of a multiple of 90 degrees, which agrees with theory. The unwindowed set in Fig. 11(a) exhibits extreme aliasing, and the widowed set in Fig. 11(b) exhibit virtually no visible effects, although close inspection of the images reveals that aliasing is present.

In all experiments that follow, the rotational angular resolution is 360/256 degrees. The experiments use the agricultural image in its original orientation and oriented at $\theta_0 = 29.53$ degrees, which is 21 increments of the angular resolution. Figure 12(a) shows the 2D Fourier transform of the windowed image, and shows

that some of the coefficients are masked out to prevent them from participating in the correlation computations.

There are two regions masked out. The outer region is removed because rotations of these regions do not map onto valid transform coefficients in all rotations. The inner region is masked out because the magnitude of aliasing effects is greatest in a region near the origin of the Fourier plane. Eliminating these coefficients greatly improves the precision of the correlation algorithms. The participating coefficients lie within a radius of $N/2$ of the central peak, and outside a radius that varies in the experiments from $N/32$ and $N/4$.

Figure 12(a) exhibits strong horizontal and vertical structures. These structures are related to the circular outline, rather than to the example image. Figure 12(b), the discrete Fourier transform of a black disk of size $N/2$, demonstrates that those structures are artifacts of the circular shape. The discrete transform in Fig. 12(b) is not circularly symmetric, and the lack of circular symmetry is due to aliasing. The Fourier transform of a rotation of the example image has two components—one is a rotation of the Fourier components of the content of the example, and the other is the collection of Fourier components produced by the circular boundary, which is aligned with the unrotated axes.

Consider the magnitudes of the Fourier transform of the example image and the Fourier transform of a rotated version of the same image. Rotate the magnitude of one of the transforms by various angles and correlate it with the magnitude of the other transform. The correlations exhibit two different kinds of correlations peaks. One kind of peak occurs where the transforms of the image content align with each other. This occurs at an angle equal to the rotational difference of the images, and at that angle plus 180 degrees. The second type of peak occurs where the aliasing structures align, and these occur at multiples of 90 degrees regardless of the relative rotation of the images. Figure 12(b) also indicates by the shading that the magnitude of the aliasing is greatest near the center frequency of the

transform. Hence, removing a circular region of frequencies centered at the origin tends to remove coefficients that contribute most to false-correlation peaks.

The first experiment does correlations of Fourier magnitudes of two images. One image is the original rotated by θ_0. The other image is the original image rotated by $\theta_0 + \theta$ in the pixel domain, and then rotated by $-\theta$ in the Fourier domain, where θ varies over the unit circle. The resulting transforms are perfectly aligned, and should yield unity correlations in the absence of aliasing. Figure 13 plots the results of this experiment for the images, both with and without Blackman windowing. The major characteristic of the plots is the periodicity of 90 degrees, and the peak heights of 1.0 for θ a multiple of 90. Both aliasing and interpolation errors have this characteristic. The correlations tend to diminish as the low-frequency exclusion radius increases. The degradation is locally greatest in Fig. 13(a) at $\theta = -\theta_0$ and 90 degree offsets of this shift. At this particular angle, no rotational interpolation occurs in the pixel domain. The only interpolation that occurs at this angle is within the rotation of $+\theta_0$ in the Fourier domain. Apparently, correlation improves slightly when it involves two interpolations instead of one. The plots in Fig. 13(b), correlations of the Blackman windowed image, exhibit far less degradation than do those of Fig. 13(a), as we expect.

The four experiments in Figs. 14-17 illustrate four different ways to carry out image correlations. The four methods introduce progressively greater sources of correlation distortion. All four experiments correlate the Fourier magnitude of the example image oriented at θ_0 with Fourier magnitudes of rotations of the original image. Table 1 lists the correlation peak heights, SNR, and the ratio of peak height to false-peak height in the series of figures. The signal-to-noise ratio tabulated is the ratio of peak height to the root-mean-square value of correlations from points outside of the peak.

Figure 14 is a plot of the baseline experiment, which contains interpolation error only. This experiment rotates the image in the pixel domain, and then computes

the transforms. It is computationally expensive because it requires a Fourier transform for each of 256 rotations. This experiment clearly shows why low-frequency exclusion is required. The plot for a radius of 0 yields near unity correlations for all angles, and is not useful for finding the correct angular displacement. This is due to the very large central peak in the Fourier transform, which dominates the correlation computation. Discrimination improves for a radius $N/32$ and is excellent for a radius of $N/4$. These data confirm that Fourier-magnitude correlations can find the orientation difference of two images. The minor peaks visible in the $N/4$ curve are at 90 degrees with respect to the correct correlations. They are due to correlations between x and y-axis artifacts present in the transforms.

The next experiment introduces aliasing error. It transforms the original image, and rotates the transform in the frequency domain. It performs 256 rotations, but does only a single transform. The results appear in Fig. 15. Note the degradation apparent in the top curve. The peaks in that curve are at multiples of 90, and are totally erroneous. The same peaks are present but are less dominant in the other two plots in the figure. The main idea conveyed by this figure is that Fourier magnitude can be used effectively with rotations of Fourier transforms to find the rotational difference of two images. Though aliasing lowers the peak for $N/4$ compared to its counterpart in Fig. 14, the peak is still high enough to be useful for angular registration. The false peaks virtually disappear in the $N/4$ curve in Fig. 15(b), which confirms that they are artifacts of aliasing.

Figure 16 shows the effects of converting from rectangular to polar coordinates. This experiment converted the transforms of the image, both rotated by θ_0 degrees and unrotated, into polar coordinates with an angular resolution of 360/256 degrees. It computed normalized correlations in the polar coordinate system. These correlations are normalized correlations of one transform with cyclic shifts of the other transform, and can be done very efficiently by making use of the Convolution Theorem. The total computational cost for doing correlations for all rotations is of

the order of $N \log N$ where N is the number of coefficients in the Fourier representation. The additional degradation in this figure is due to coordinate conversion, including the interpolation errors in the new coordinate system. The plot shows correlations for only two low-frequency exclusion radii.

The last experiment, depicted in Fig. 17, shows the equivalent results for phase correlation. This experiment computes the phase-correlation function of the Fourier magnitudes in polar coordinates. It differs from the prior experiment only in the computation performed on the polar-coordinate representations of the transforms. The computational complexity is approximately equal to the complexity of normalized correlation. As pointed out in [6], rotations of the magnitude of the Fourier transform in rectangular coordinates are equivalent to translations of the function expressed in polar coordinates. Hence, phase-correlation techniques that can find a translation by which two images differ [10] can also find the translation by which two Fourier magnitudes in polar coordinates differ. This translation is the rotational difference of the original transforms. The same reasoning extends to finding differences in scale [6, 14, 17]. As mentioned earlier, the peak heights in Fig. 17 double when using 180-degree half circles of transforms instead of the full 360-degree transforms shown in Fig. 12(a). The false peaks in Fig. 17(a) for $N/32$ are similar to those observed in [11].

It is clear from the figures and the data in Table 1 that rotationally dependent aliasing degrades the correlation data. To grasp the aliasing contribution, consider, for example, the data for a low-pass cutoff of $N/4$ without a Blackman window in Figs. 14(a) through 17(a). Interpolation error drops the peak from 1.00 to 0.852. Aliasing reduces it to 0.667. Polar-coordinate conversion reduces it to 0.5829. The phase-correlation peak falls to 0.452. For Blackman windowed data with a low-pass cutoff of $N/4$, the sequence of peaks is 0.923, 0.935, 0.950, and 0.835. This suggests that the noise due to interpolation, coordinate conversion and aliasing

causes negligible peak degradation on the normalized correlation peak, but the phase correlation peak is affected somewhat more by these noise sources.

In spite of the lower correlation peaks, phase correlation enjoys superior signal-to-noise ratio. With a low-pass cutoff of N/4, the ratio is not less than 22 and is as high as 115 in this experiment, whereas the best SNR that any other scheme produced was just above 7. False peaks are a potential problem when the true peak and false peak lie close to each other. Normalized correlation may yield higher precision than does phase correlation in such cases.

In Fig. 17(a), the relatively low values of the correct peaks and the relatively high values of the false peaks are consistent with the analysis in the previous section. Rotationally dependent aliasing has two impacts on these graphs. It directly lowers the correct peak, and indirectly lowers the correct peak through the false peaks that it creates. Nevertheless, the signal-to-noise ratio for phase correlation is exceptionally high for the $N/4$ graph in Fig. 17(a).

Overall, the experimental data in Figs. 14-17 are consistent with our observation that rotationally dependent aliasing is a noise source in Fourier-magnitude correlations, and our data help quantify the degradation due to that noise. By Blackman windowing the sharp boundary of the image we removed substantial high frequency energy that is not part of the image. The reduction of that energy together with a low-pass cut-off filter of $N/4$ reduced aliasing noise to an extent that it became insignificant in all forms of correlations.

5. CONCLUSIONS

The main result in this paper is the discovery, explanation, and quantification of rotationally dependent aliasing. Experimental data shows that such aliasing degrades the precision of phase correlation. Aliasing effects in our experiments are quite vivid. Sharp boundaries produce high frequencies, which in turn cause substantial aliasing, and thereby create observable rotationally dependent degrada-

tion. Smoothing the boundaries removes much of the high frequency energy, which substantially reduces aliasing, and increases the quality of phase correlation.

Historically, phase correlation has proven itself to be computationally efficient and accurate when used to register images with respect to translations. It works well in that context in spite of aliasing in the Fourier domain, because aliasing has identical effects on the Fourier transforms of images that differ only by translation. Rotations change the situation because phase correlations of aliased rotated transforms depend on the individual orientations of the images. Aliasing has two effects on phase-correlations – it diminishes the peak at the correct orientation and creates false peaks at multiples of 90 degrees.

Our experiments tested normalized correlation as well as phase correlation, and show that normalized correlation may be a good alternative to phase correlation in some cases because it is less sensitive to aliasing noise. It has inherently better precision when orientations differ by an angle close to a multiple of 90 degrees. The computational cost of normalized correlation is about equal to that of phase correlation. In practice, many additional sources of noise impact the precision of image registration. Correlation peaks may be lower than those in our controlled experiments, and background noise can be higher. Consequently, one needs to examine the results of many more experiments in order to judge accurately the relative qualities of phase-correlation and normalized correlation.

Another alternative that has been discussed in the literature is the use of prefilters and windows on the images in order to reduce effects of aliasing. For example, [6] proposes a Hanning window filter to apply to the Fourier transform magnitudes before the coordinate conversion. The Blackman window used in this paper confirms that windowing has strong positive effects on registration precision. We have shown that aliasing effects can be reduced by cutting off low frequencies as well as smoothing the image boundary to remove high frequencies. While this type of filtering reduces aliasing, the precision of image registration is often dependent on

tight registration of sharp edges interior to an image. Precision degrades when these edges are softened. Hence, two opposing forces are in contention regarding filtering. Alias reduction tends to encourage high-pass filtering; precision tends to discourage high-pass filtering.

ACKNOWLEDGMENT

The authors are grateful to Prof. Michael Orchard of Princeton University and to referees of an early revision of this paper for their helpful comments. We also thank Rob Wolpov and Steve Martucci of NEC Research for their careful reading of the manuscript. The aerial photograph is used with the permission of Positive Systems.

REFERENCES

1. S. Alliney, G. Cortelazzo, and G. A. Mian. On the registrations of an object translating on a static background. *Pattern Recognition*, 29(1):131–141, January 1996.
2. S. Alliney and C. Morandi. Digital image registration using projections. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 8(2):222–233, March 1986.
3. R. N. Bracewell. Two dimensional aerial smoothing in radio astronomy. *Australian J. Physics*, 9(4):297–314, 1956.
4. L. Brown. A survey of image registration techniques. *ACM Computing Surveys*, 24(4):325–376, 1992.
5. S. H. Chang, F. H. Chen, W. H. Hsu, and G. Z. Wu. Fast algorithm for pointer pattern matching; invariant to translations, rotations, and scale changes. *Pattern Recognition*, 30(2):311–320, February 1997.
6. Q.-S. Chen, M. Defrise, and F. Deconinck. Symmetric phase-only matched filtering of Fourier-Mellin transforms for image registration and recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(12):1156–1168, December 1994.
7. B. Dasgupta and B. N. Chatterji. Fourier Mellin transform based image matching algorithm. *Journal of the Institution of Elect. and Telecommunications Engineers*, 42(1):3–9, 1996.
8. E. De Castro and C. Morandi. Registration of translated and rotated images using finite Fourier transforms. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, PAMI-9(5):700–703, 1987.
9. F. J. Harris. On the use of windows for harmonic analysis with the discrete Fourier transform. *Proceedings of the IEEE*, 66(1):51–83, January 1978.
10. C. D. Kuglin and D. C. Hines. The phase correlation image alignment method. In *IEEE 1975 Conference on Cybernetics and Society*, pages 163–165, September 1975.
11. L. Lucchese, G. M. Cortelazzo, and C. Monti. High resolution estimation of planar rotations based on Fourier transform and radial projections. In *Proc. ISCAS 97*, volume 2, pages 1181–1184, June 1997.
12. A. V. Oppenheim and R. W. Schafer. *Digital Signal Processing*. Prentice-Hall, Englewood Cliffs, NJ, 1975.
13. D. V. Papadimitriou and T. J. Dennis. Stereo disparity using phase correlation. *Electronics Letters*, 30(18):1475–1477, September 1994.
14. B. S. Reddy and B. N. Chatterji. An FFT-based technique for translation, rotation, and scale-invariant image registration. *IEEE Trans. on Image Processing*, 3(8):1266–1270, August 1996.
15. H. S. Stone, M. Orchard, E.-C. Chang, and S. A. Martucci. A fast direct Fourier-based algorithm for subpixel segistration of images. *submitted to IEEE Transactions on Geoscience and Remote Sensing*, July 2000.
16. P. Thévenaz, U. E. Ruttimann, and M. Unser. A pyramid approach to subpixel registration based on intensity. *IEEE Trans. on Image Processing*, 7:27–41, January 1998.

17. J. Z. Wang, G. Wiederhold, O. Firschein, and S. X. Wei. Wavelet-based image indexing techniques with partial sketch retrieval capability. In *Proceedings of Advances in Digital Libraries*, pages 13–24, May 1997.

18. M. Wu, M. L. Miller, J. A. Bloom, and I. J. Cox. A rotation, scale and traslation resilient public watermark. In *1999 ICASP*, volume 4, page 2065, March 1999.

TABLE 1

Results of Correlation Experiments

Blackman Window	Noise Source	Exclusion at $N/32$			Exclusion at $N/4$		
		Peak Height	SNR	Peak / False Peak	Peak Height	SNR	Peak / False Peak
No	Interpolation	0.991	1.096	1.071	0.852	2.406	1.401
	and Aliasing	0.821	1.106	0.911	0.667	2.809	1.583
	and Coordinate Conversion	0.750	1.118	0.824	0.582	2.899	1.668
	and Phase Correlation	0.300	6.558	0.907	0.452	22.742	3.925
Yes	Interpolation	0.976	2.053	1.413	0.923	6.440	2.610
	and Aliasing	0.979	2.137	1.434	0.935	7.285	2.718
	and Coordinate Conversion	0.992	1.959	1.307	0.950	7.012	2.732
	and Phase Correlation	0.751	73.975	13.753	0.835	115.879	None

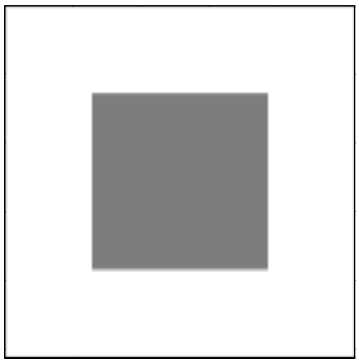

Fig. 1(a): An example image within a rectangular outline.

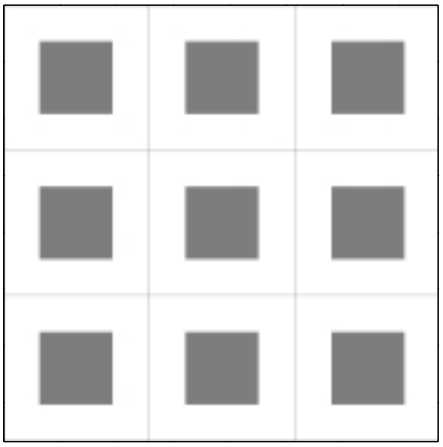

Fig. 1(b): The image in an infinite tiled plane.

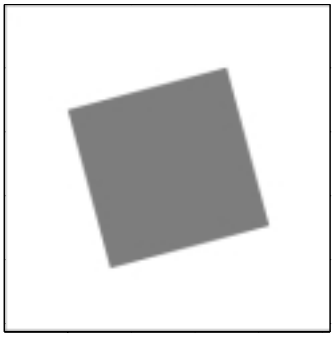

Fig. 2: A rotation of the image in Fig. 1(a).

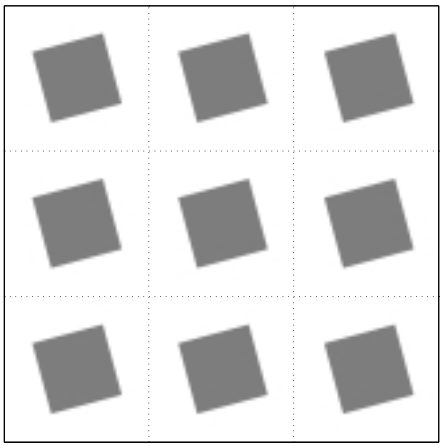

Fig. 3: The rotated image in a tiled plane.

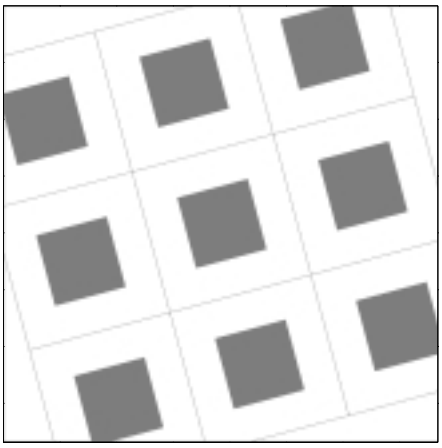

Fig. 4: The rotated tiled plane.

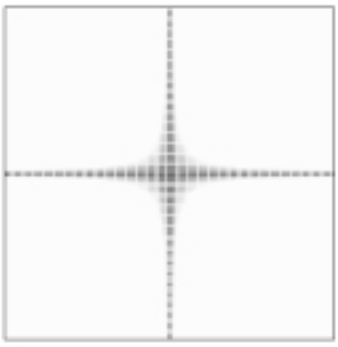

Fig. 5: Transform of Fig 1(a).

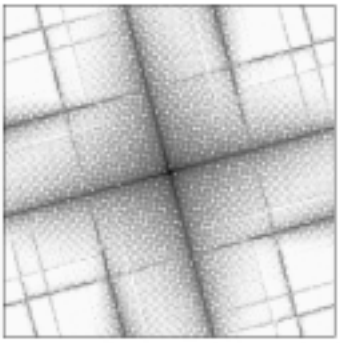

Fig. 6: Transform of Fig. 2.

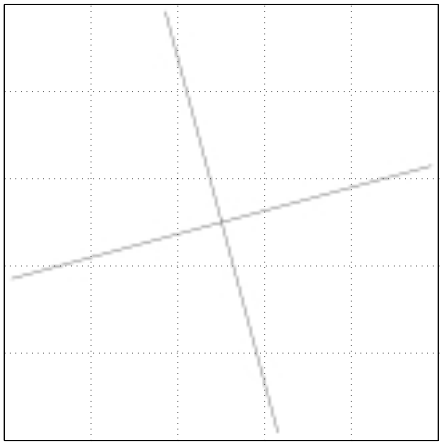

Fig. 7(a): Tiled rotated infinite transform plane of Fig. 5.

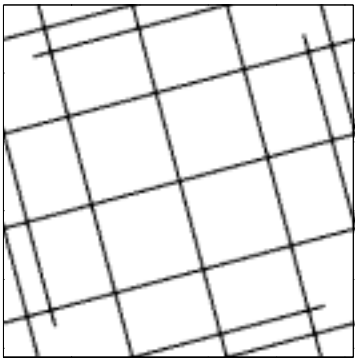

Fig. 7(b): The rotated transform plane collapsed into a single tile.

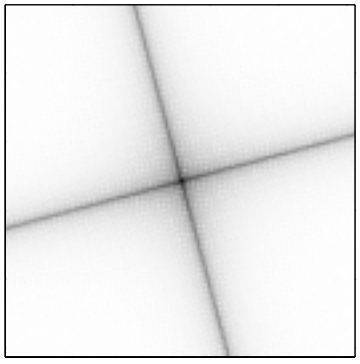

Fig. 8(a): Rotated transform of Fig. 5.

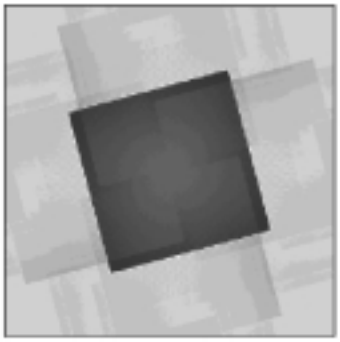

Fig. 8(b): Inverse transform of Fig. 8(a).

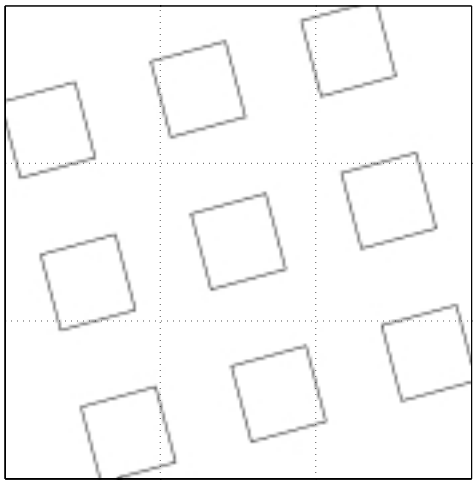

Fig. 9(a): Tiled rotated infinite pixel plane of Fig. 1(b).

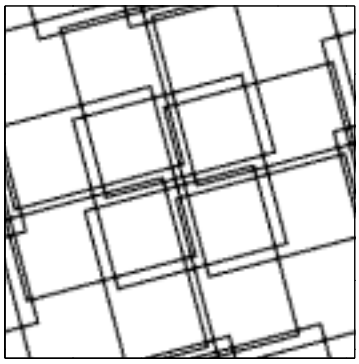

Fig. 9(b): The rotated pixel plane collapsed into a single tile.

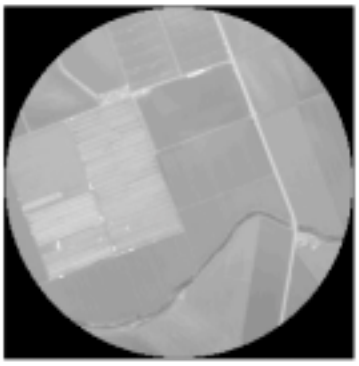

Fig. 10(a): Agricultural aerial photo.

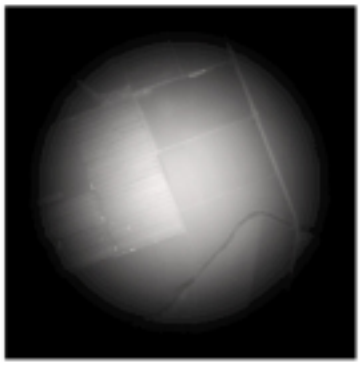

(b): Photo after applying Blackman window.

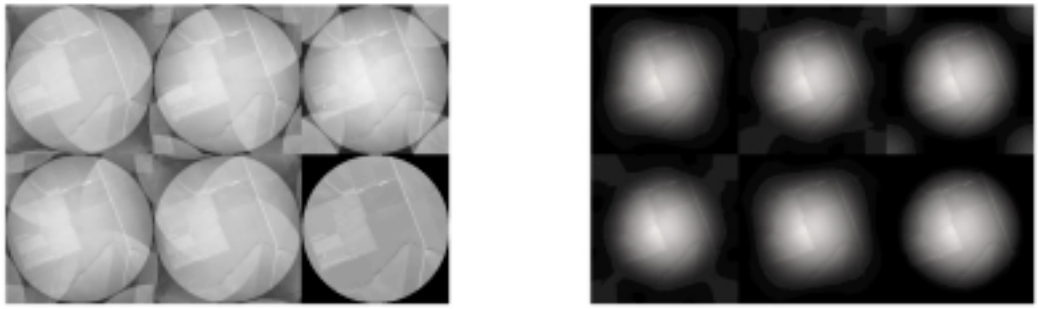

(a) (b)

Fig. 11(a) and (b): Figures created by rotating by θ, transforming, rotating by $-\theta$, and inverse transforming. Angles used from left to right, top to bottom, are 15, 30, ..., 90 degrees. (a) Original image. (b) Windowed image.

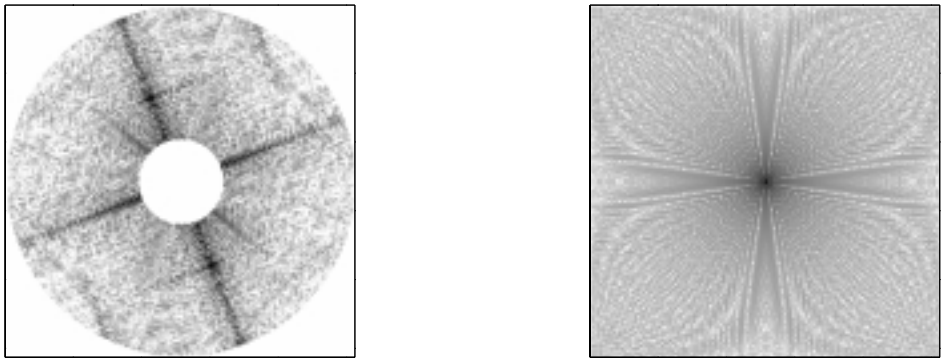

(a) (b)

Fig. 12(a): Coefficients of the Fourier transform used in correlations. The radius of the inner circle varies in the experiments. Fig. 12(b): The Fourier transform of a circular disk.

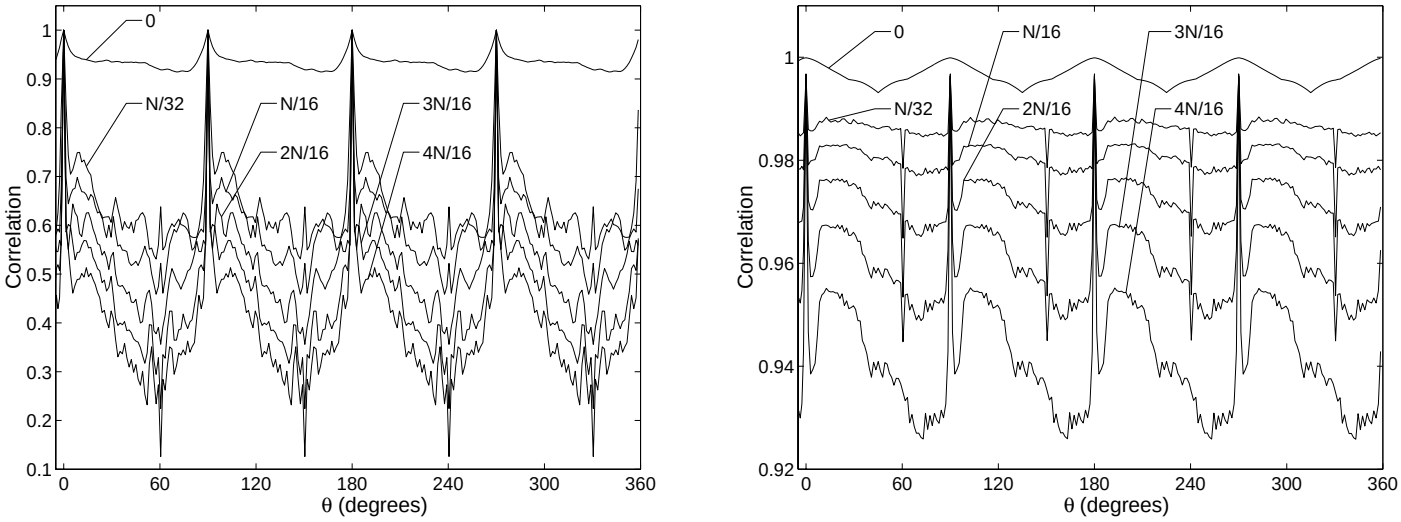

(a) (b)

Fig. 13: Correlations of the transform of a 30°-rotated image with the image rotated $30 + \theta$, transformed, and rotated $-\theta$ for varying θ and low-frequency exclusion radii. All correlations would be unity without aliasing and interpolation error. Plots show correlations for various inner circle radii. (a) Original image. (b) Windowed image.

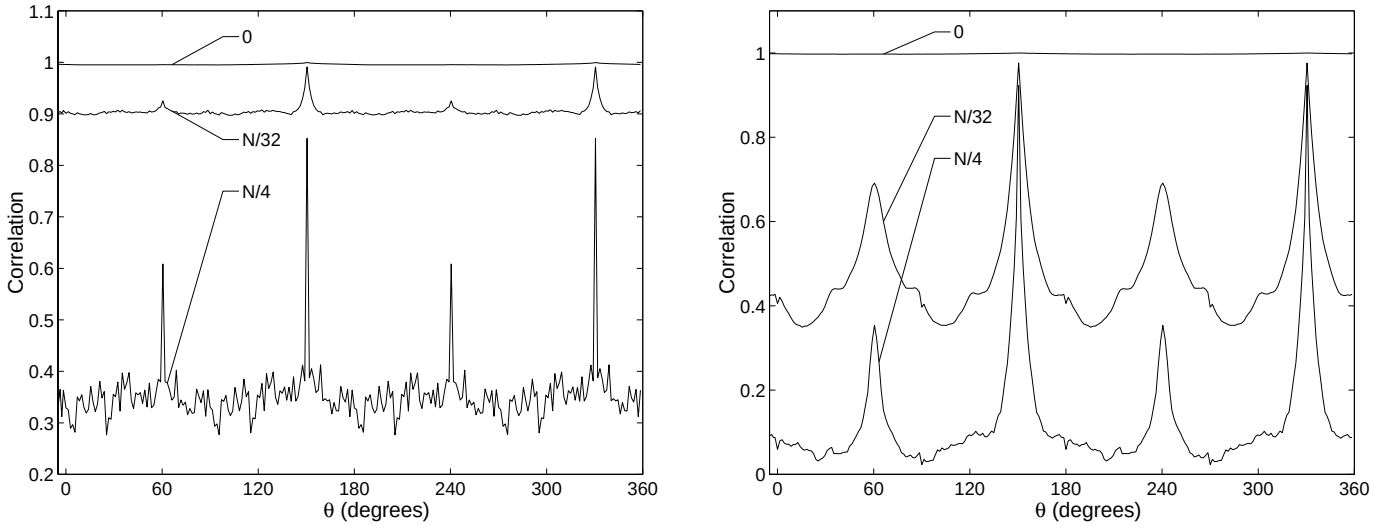

(a) (b)

Fig. 14: Correlation experiment baseline. Rotations in pixel domain. Inner exclusion radii are 0, $N/32$, and $N/4$. Correct registration is at 150.47. (a) Original image. (b) Windowed image.

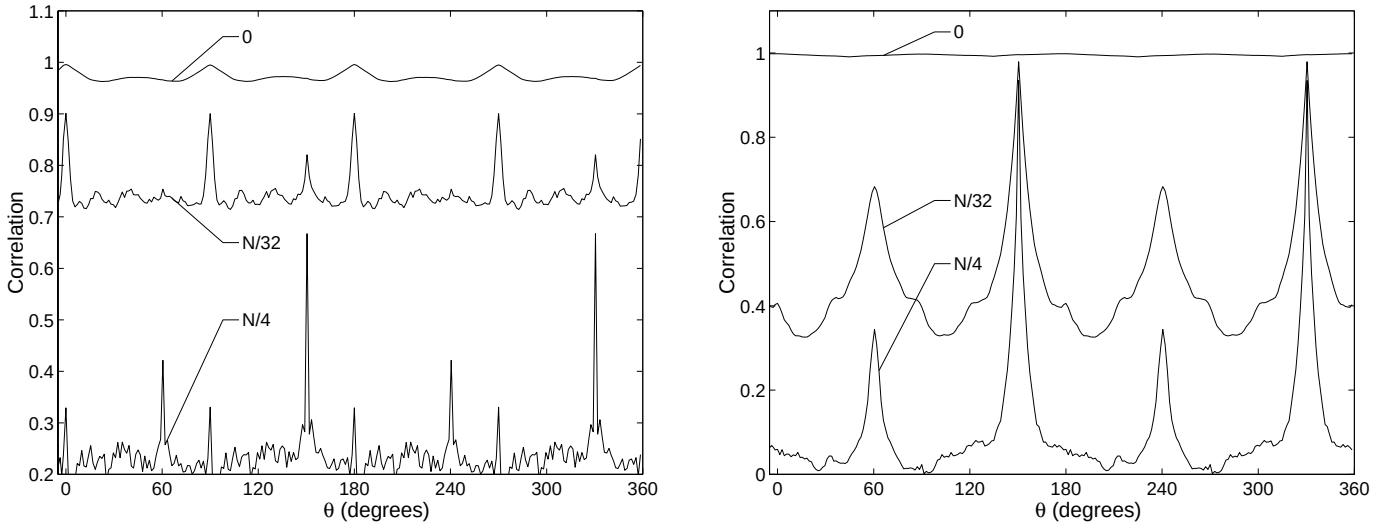

(a) (b)

Fig. 15: Normalized correlations with aliasing errors caused by rotations in the Fourier domain. (a) Original image. (b) Windowed image.

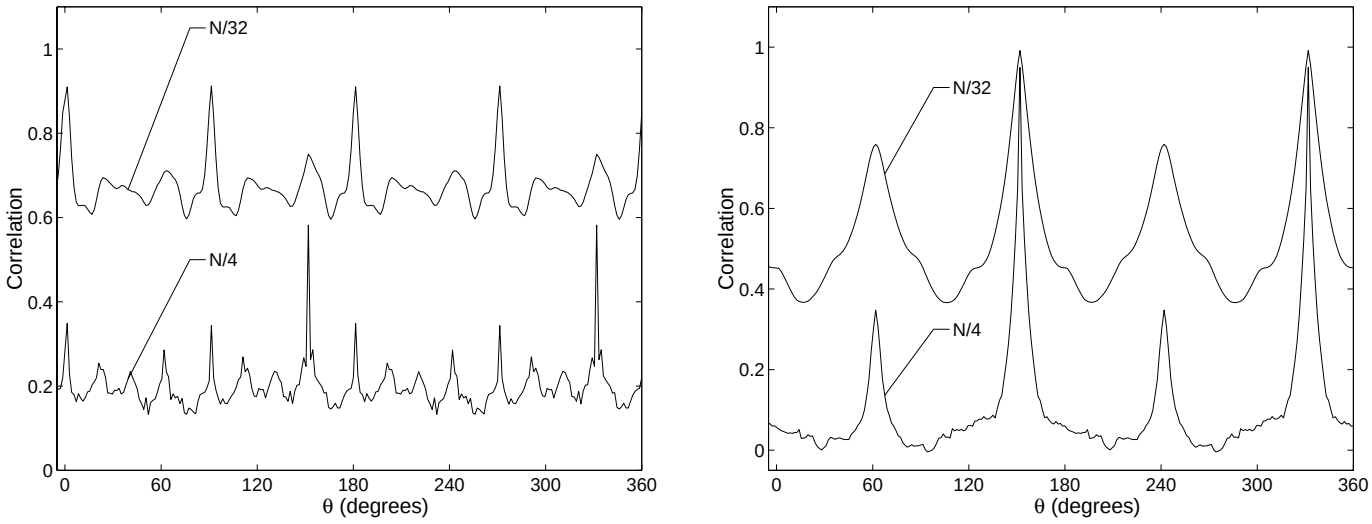

(a) (b)

Fig. 16: Fig. 15 experiment with normalized correlations of transforms after conversion to polar coordinates. (a) Original image. (b) Windowed image.

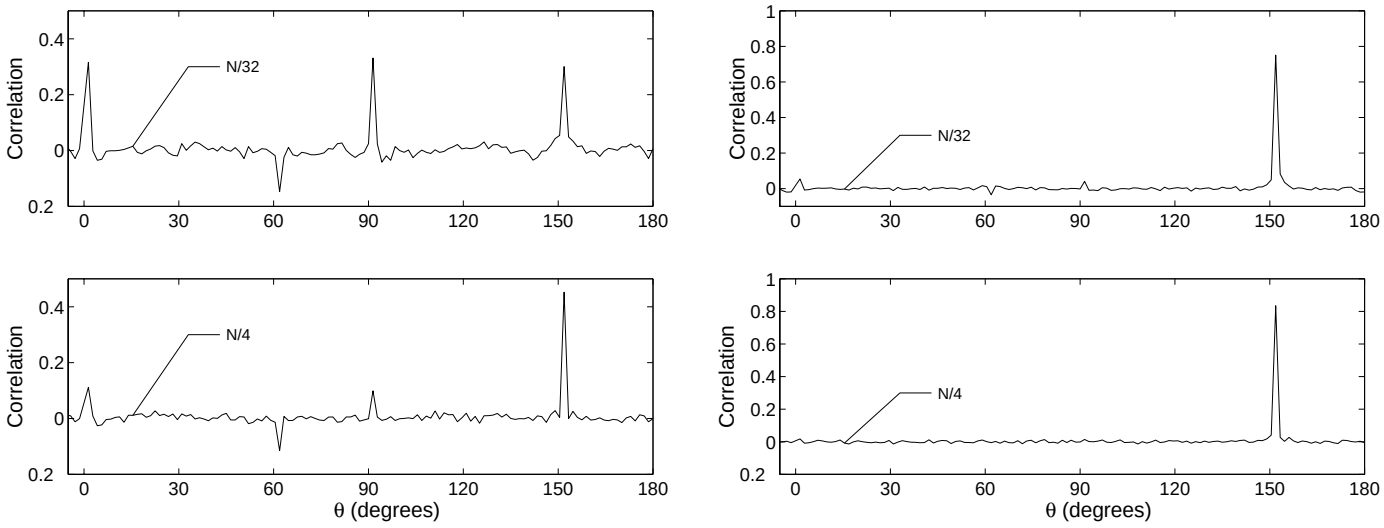

(a) (b)

Fig. 17: Fig. 16 experiment with phase correlations instead of normalized correlations. (a) Original image. (b) Windowed image.