Sampling Concerns in Scanline Algorithms

Martin Fleury and Adrian F. Clark

M. Fleury and A. F. Clark are with the Department of Electronic Systems Engineering, University of Essex, Wivenhoe Park, Colchester, CO4 4SQ, U.K.

Abstract

This paper investigates the problem of providing suitable interpolation for scanline algorithms. These algorithms are of interest as they are parallelizable. A structure for analysing the problem is given. The theory in regard to resampling is developed in the context of a scanline algorithm for image rotation. The theory is compared to results arrived at in practice. Alternative interpolation schemes are discussed, including the use of a cubic Hermite interpolator. The paper points to theoretical limitations of scanline algorithms.

Keywords

Resampling, affine transforms, scanline algorithms, parallelization

I. Introduction

Whenever batches of images need to be aligned then an image may need to be rotated. Rapid if not real-time processing of such images may require alternative rotation algorithms that compromise on fidelity. Examples from recent studies are: in Figure 1 Image 1, a Magnetic Resonance (MR) image, is from a study of changes in breast volume during the menstrual cycle [1]; and Figure 1 Image 2, a computer tomography (CT) image, is from a study on the relative amount of fat and lean tissue in a sheep [2]. However, our analysis is not dependent on the image chosen. Therefore, the choice of examples is primarily to indicate applications. An essential component of a rotation algorithm is interpolation. Theoretical analysis of interpolation schemes are also relevant to volume rendering of medical images such as MR brain scans or CT cranial scans [3].

Other techniques will require comparison of successive images, which may have been affected by intervening patient motion. One such case is in the field of digital radiology [4]. Resampling occurs when two images are registered, e.g., [5]. An example of registering electrophoresis images (from differing genetic strains of the malarial parasite) is in [6]. Where batches of images are processed there is a time constraint which might be addressed by parallel processing. This study was in part a by-product of the need for such processing by a genetic research laboratory.

A fast way of performing an image rotation, whatever the processor, is by a scanline algorithm. A 'scanline' algorithm is one in which each step of the algorithm is confined to operations on scan or raster lines of an image. Scanline algorithms have been developed for

affine transforms or warps [7] and in particular for image rotations [8, 9]. These algorithms are suitable for hardware implementation [10, 11] since they break the geometric transform into a number of simple repeatable operations which might be represented by components such as shift registers. Look-Up Tables (LUTs) can be employed to avoid trigonometric calculations. Because the number of trigonometric multiplications is reduced, scanline algorithms are also attractive to software implementations. It may have previously gone unnoticed [12] that scanline algorithms for rotation may also provide a suitable decomposition for parallel implementation on an MIMD-style message-passing multicomputer. Alternative proposed parallel implementations of affine transforms that are SIMD-based or at least fine-grained may lack flexibility. Scanline algorithms have received attention in the graphics literature [13, 14], and have potential applications in remote sensing [15] and computer vision in addition to the medical field. However, repeated resampling after each stage of the scanline algorithm may introduce unacceptable distortion. For graphical applications, the effects visible to the eye, such as jagged edges resulting from aliasing, are of concern, but for image processing applications still more rigorous criteria exist. Repeated resampling may also arise in image compression by a coding tree mechanism, such as binary tree predictive coding [16].

Since feature points are commonly used for alignment between successive images, it is important to be aware of the image blur introduced by resampling. This manifests itself and is conveniently analysed in the frequency domain representation of an image. The convenience largely arises because of the Convolution Theorem [17] whereby convolution in one domain is multiplication in the other. Another technique used is to work in the spatial domain by means of difference images for selected test images. This technique is less general, though careful construction of test images can be helpful [18]. In [19], the question of the response to differential phase shift in the image was considered in the spatial domain. That study, though thorough, was concerned with bi-level images and used a restricted model for the image capture process. Fundamental and general work in the area of image reconstruction was completed in [20, 21] and [22]. The problem of repeated resampling was not explicitly examined, though it may occur for multi-pass registration, in computer graphics [23] and in the present application.

This short study is concerned with one particular scanline algorithm, image rotation. It became apparent when re-sampling images, after checking correct operation of the algorithm, that it did not supply consistent resampling in the output image, though this was not obvious unless a test image was used. Two questions arise: how can this distortion be characterized and are there measures that can correct the distortion? For instance, it is possible to vary the interpolation kernels that are used to reconstruct the image, though used indiscriminately this may not improve the situation. In this paper, a theoretical model is expounded and developed in some detail to characterize a lower bound to the distortion arising, but it should be emphasized that this analysis is nugatory if the algorithm employed imposes a high upper bound to the distortion. In other words, the theory may *not* explain the results if the algorithm introduces substantial distortion additional to that of repeated resampling.

Section II gives more detail on the algorithm used, especially in regard to sampling concerns. This section also analyses the repeated resampling problem in the spatial domain, establishing a lower bound. Section III expounds the theory behind *correct* frequency domain analysis of radiometric error and extends this theory to the case of repeated resampling. Section IV gives the results of experiments to see whether the lower bound could at all be reached. Error measurements were made as a means of indicating resampling scheme rankings, including use of a Hermite interpolator. Finally, Section V provides a summary and draws some conclusions.

II. Resampling a Scanline Algorithm

Though it may be inconvenient for a hardware implementation by having more than two skews, the thrice-skew image rotation algorithm introduced in [9] avoids rescaling[1] and provides an isomorphic function for all angles in the first quadrant.[2] The algorithm decomposes the rotation into a horizontal skew, followed by a vertical skew, followed by another horizontal skew:

[1]Rescaling would run the danger of aliasing through undersampling.

[2]In [24], a twice-skew algorithm is described, actually ascribable to [7], which may run into the difficulties noted. This paper notices the problems with resampling but does not provide an error analysis. It does, however, give a good analysis of the errors that would arise from approximating the interpolating position so as to use an LUT for the cubic convolution coefficients (further described in this section).

$$\begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} 1 & -\tan\frac{\theta}{2} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ \sin\theta & 1 \end{bmatrix} \begin{bmatrix} 1 & -\tan\frac{\theta}{2} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}. \quad (1)$$

Each row in the initial horizontal skew is skewed by a fixed amount, which has the effect of imposing an integer-spaced resampling grid upon the row. Indeed, it is possible to see these algorithms as basically resampling algorithms. Since the amount of skew is incremented for each row, succeeding rows will have a different resampling grid imposed upon them. Commonly used 2-D interpolation kernels, as defined by International Standard 12087 [25], are separable across up to five dimensions. Examples are linear, sinc, cubic convolution, and cubic B-spline.[3] In the subsequent vertical skew, suppose that it were possible to retain the neighboring pixels within the support of the interpolation filter. It would then be possible to perform a second resampling in the vertical direction, which would complete one operation of a 2-D interpolation. The final interpolation operation for the second horizontal skew will constitute a repeated resampling. Notice that this second resampling, though performed as a row-by-row 1-D interpolation, can also be regarded as a second 2-D interpolation, where the resampling grid is perfectly aligned with the original grid in the vertical direction. This is because most interpolation kernels have value unity when the sampling and resampling grids are aligned. Ignoring the difficulty of the non-alignment of rows and columns after each of the skews, it is possible to analyse this problem as one of a 2-D resampling followed by a second 2-D resampling. In other words, we have a lower bound. In Section III the problem is analysed as just two resamplings, where the analysis is confined to one dimension, assuming the results carry on to two dimensions.[4]

[3]For applications such as image magnification, but *not* normally for rotation, separability is used by means of an intermediate image.

[4]In two dimensions there will be an additional anisotropic effect but this will not necessarily appear in lower orders of the transfer function's Taylor expansion. For instance, as interpolation kernels are even and separable, one might make a comparison with the expansion of a Binomial filter (with $b_x^2 = \frac{1}{4}[1\ 2\ 1]$), where one finds the Fourier transform of the 2D filter and its Taylor expansion are:

$$\begin{aligned} \hat{b}(\nu) &= \tfrac{1}{4}(1+\cos(\pi\nu_x))(1+\cos(\pi\nu_y)) \\ &= 1 - \tfrac{\pi^2(\nu_x^2+\nu_y^2)}{4} + \tfrac{\pi^2\nu_x^2\nu_y^2}{4} + \tfrac{\pi^4(\nu_x^4+\nu_y^4)}{48} + O(6), \end{aligned} \quad (2)$$

when the first anisotropic term is of order four. In fact, examination of the Taylor expansions in the frequency domain of the interpolation kernels used in this study and given in [21] shows that the anisotropic terms are only

The effect of the first horizontal skew followed by the vertical skew is to distort the region of support of the interpolating kernel, which is a circle (cf. [26]), to become an ellipse. (Strictly, the ideal kernel is $\mathrm{sinc}(x, y)$ and its region of support is infinite. A circle represents a region of support which is isotropic and an ellipse represents an anisotropic region of support.) The response changes in the direction that the ellipse departs from circularity. A suitable test image [11] is a grid composed of lines that are one pixel in width. This test image introduces fixed high frequency features which are difficult for the interpolating kernel to cope with. In Figures 2–4, an original test image, the result after using linear interpolation with the thrice skew algorithm, and a difference image are shown for a rotation of 50°. The comparison was made with an image that was formed by directly resampling with linear interpolation from output image back to input image by means of an inverse transform[5]. Rotation was made in an anti-clockwise fashion and a first step is to check the operation of both algorithms by use of a protractor on the printed image. The computer code should also be scrutinized. In particular, the origin of the rotation should coincide for both algorithms, when on an integer grid with even-sized images the center of the image will be offset by one pixel in both the x and y directions. It will be apparent from the difference image that the distortion occurs principally in one direction. There are also differences arising from the differing treatment of image boundaries. In general, distortion increases with increasing angle of rotation within the first quadrant. This is consistent with the nature of the algorithm. The first horizontal skew involves a factor of $\tan\theta/2$, which monotonically increases the distortion of the support region in the angle range $0° \leq \theta < 90°$. The vertical skew involves a factor of $\sin\theta$ that introduces the principal distortion seen in Figure 4. For this algorithm the vertical distortion will always be greater than the horizontal distortion. Some twice-skew scanline algorithms

at most a fourth order effect and, for the cubic convolution kernel with $\alpha = -0.5$, are an eighth-order effect. The Taylor expansion for a 1D cubic convolution interpolation kernel is $\hat{r}(\nu_x) = 1 - \frac{4}{15}(2\alpha + 1)(\pi\nu_x)^2 + \frac{1}{35}(16\alpha + 1)(\pi\nu)^4 + \ldots$. Set $\alpha = -0.5$ and use the separability of the kernel to see the effect.

[5]Where there is global memory access, it is possible to pass along the output image on a row-by-row basis, locating the source of each output pixel by means of the reverse rotation (screen scanning). This method is easier to organize and computationally more efficient than going from input image to output (source scanning). Source scanning can result in holes in the image from overlapping pixels. The result of screen scanning represents the best that can be achieved for a rotation algorithm with linear interpolation. Rotating an image by optical means would introduce more error from transferring the result to the computer.

give the possibility of altering the order of the horizontal and vertical skew, thus enabling the direction of error to be controlled.

It is also possible to use a common test image (Figure 5 Image 1), equivalent in 2D to a swept frequency scan in 1D [27]. As this image is circularly symmetric, the rotated version by a thrice-skew algorithm can be compared with the original as well as with the result of a one-pass algorithm. The intention of this type of image is to give a visual check which should be confirmed by frequency domain analysis. The generating equation is

$$f(r) = \begin{cases} A\cos(2\pi f_m r^2/P) & r \leq P/2 \\ A\cos(2\pi(f_m P/2 - f_m(r-P)^2/P)) & P/2 \leq r \leq 3P/2. \end{cases} \tag{3}$$

$f(r)$ is the intensity at radius r. A 512×512 image was generated with A, the amplitude, set to 255 in the image. P is the period in which the spatial frequency passes from 0 to the maximum frequency, f_m. The period in Figure 5 is 320 (the origin is in the center). The maximum frequency was chosen to be 0.25 because on our equipment (graphical display unit supplied with a SPARCstation 5) moiré patterns disturbed the display at higher frequencies. Readers may wish to repeat the experiment at a frequency approaching 0.5 cycles/sample interval. Unlike a two-pass scanline algorithm [7], in comparing the one-pass version (Figure 5 Image 2) with the thrice-skew version (Figure 5 Image 3), or indeed with the original, the display did *not* show significant differences. Figure 5 Image 4 (a rotation at 50°) is included for comparison with the earlier test image. The root mean square (RMS) error (see Section IV) plots between the original image and the rotated versions are given in Figure 6. Clearly this measure will depend on the proportion of background revealed, but still for the most part the thrice-skew RMS error is greater than the one-pass version. The error for the thrice-skew version does not have the cyclic pattern of the one-pass version. The error curve between the two rotation methods is also given, treating the one-pass version as the base image. Within the resolution of the plot, the relative error does not alter appreciably across the angle range. Bear in mind that the frequency increases linearly from the centre of the test image, within the period of equation 3. A profile taken across the image will allow the relative attenuation at a given frequency. Comparative profiles for a rotation of 50° are shown in Figure 7 for the reader's perusal. The profiles are taken horizontally through the origin.

A number of experiments were made with three interpolation kernels, which make a compromise between computation and accuracy.[6] Accuracy is defined in a least-squares sense, though if verisimilitude is required then one should take account of the human visual response, which acts as an integrator at certain frequencies and as a differentiator at others [30]. Recent interpolation schemes tolerate a greater computational burden than is the case here [31, 32].

The nearest neighbor kernel is:

$$r(x) = \mathrm{rect}(x) = \begin{cases} 1, & |x| \leq \frac{1}{2} \\ 0, & \text{otherwise,} \end{cases} \tag{4}$$

with frequency domain representation $\mathrm{sinc}\nu$. Linear interpolation uses a Bartlett window:

$$r(x) = \mathrm{tri}(x) = \begin{cases} 1 - |x|, & |x| < 1 \\ 0, & \text{otherwise,} \end{cases} \tag{5}$$

with frequency domain representation $\mathrm{sinc}^2\nu$. Finally, the cubic convolution function [33, 22] is

$$r(x) = \begin{cases} (\alpha+2)|x|^3 - (\alpha+3)|x|^2 + 1, & |x| < 1 \\ \alpha(|x|^3 - 5|x|^2 + 8|x| - 4), & 1 \leq |x| < 2 \\ 0, & \text{otherwise,} \end{cases} \tag{6}$$

where $-3 \leq \alpha \leq 0$ is a free parameter in the continuity equations that are used to establish the spline coefficients.[7]

[6] It is well known [28] that, provided a signal has been sampled at least at the Nyquist frequency, which is twice the maximum frequency component in the original signal, then it is possible in principle to reconstruct it perfectly. This would be achieved by multiplying by a rect function in the frequency domain. Unfortunately, in the spatial domain the equivalent sinc function ($\frac{\sin(\pi x)}{\pi x}$) requires global support. The Taylor series for the sinc function converges slowly and truncation results in ringing known as the Gibbs phenomenon [29], since truncation introduces a discontinuity. It may be possible to reduce the ringing by multiplying by other than a rect function. For raised cosine interpolation kernels and others not considered here see [13].

[7] Notice the use of a spline, or piecewise polynomial formula, which localizes the effect of altering the magnitude of any of the interpolation points. The linear interpolation formula can be found from the Lagrangian interpolation formula, but Lagrange polynomials of odd order are odd functions [28]. In general, Lagrangian interpolating polynomials are not guaranteed to converge to the image intensity function and if they do the approximation to the derivative of the intensity function is poor [34]. This is unsatisfactory for subsequent processing such as edge detection by spatial convolution. Additionally, the spline procedure, if constrained to imitate the properties of the sinc function, also removes the discontinuities that would otherwise be formed by truncating sinc.

If $\alpha = -0.75$ then continuity up to the second derivative is established, which is necessary to equate curvatures. However, α also moulds the low frequency response [20]. Observation of image spectra establishes that most of the image energy is in the low frequency region, a fact used by orthogonal-transform compression algorithms. In (strictly) the normalized wave-number domain ($\nu = \frac{1}{\lambda}$, with λ the wavelength), $\nu = [0, 1)$[8] using $\alpha = -0.5$, the choice of this study, ensures that the response below about $\nu = 0.2$ is good.[9] The cubic convolution kernel has Fourier transform (where the 'hat' notation, $\hat{r}$ indicates the Fourier transform of r)

$$\hat{r}(\nu) = \frac{3}{(\pi\nu)^2}[\mathrm{sinc}^2\nu - \mathrm{sinc}2\nu] - \frac{2\alpha}{(\pi\nu)^2}[3\mathrm{sinc}^2 2\nu - 2\mathrm{sinc}2\nu - \mathrm{sinc}4\nu]. \tag{7}$$

The form of this transform can be established by using the property of the Fourier Derivative Theorem ($\hat{f}'(\nu) = 2\pi i\hat{f}(\nu)$), provided that the transform of rect to sinc is accepted.[10] At $\nu = 0.5$ the $\alpha-$dependent term drops out giving a degradation of -6dB on an ideal low-pass filter. Since aliasing in practice arises from rejection at around the Nyquist frequency, the implication is that all candidates for α have a similar response in this respect. The output from the cubic convolution function may well require rescaling or possibly clipping and clamping, but this source of error is not considered here.[11]

All resamplings will in practice suffer from rounding error, which can be modelled by a uniform distribution between 0 and 1 (cf. [37] or from number theory [38]), with mean-square error of 1/12. Sample frequency-ripple will not be a problem where an interpolation kernel has the condition

$$\sum_{n=-\infty}^{\infty} r(x - n) = 1, \tag{8}$$

[8]Since in $\int_{-\infty}^{\infty} f(x)\exp(2\pi ix\nu)dx$, x is normalized in units of the sampling interval.

[9]Using $\alpha = -1.0$ makes the response below $\nu = 0.2$ marginally worse than for linear interpolation, though due to a local minimum, the response at around $\nu = 0.3$ is superior to the other interpolation kernels considered. $\alpha = -1$ may be considered for a more visually pleasing effect as it emphasizes edges.

[10]For the transforms of powers of sinc the calculus of residues should be used with typically a semi-circular contour and a pole at zero [35]. Computation is facilitated by making use of Euler's formula for $exp(2\pi i\nu x)$, the properties of even and odd functions and trigonometric simplification. Note that L'Hôpital's rule establishes $sinx/x \to 1$ in the limit. For a more rigorous approach using Lebesgue integration see [36]. Equation 7 can also be reached via the Laplacian Derivative Theorem [3].

[11]A solution to avoid the clamping problem, arising from negative side lobes in the interpolation kernel, is to introduce a DC shelf and a solution to avoid clipping is to reduce the dynamic range of the image intensity function.

as for the interpolation kernels considered here, when once-applied. Taking the Fourier transform of both sides of equation 8 yields

$$\hat{r}(\nu)\mathrm{comb}(\nu) = \sum_{n=-\infty}^{\infty} \mathrm{sinc}(\nu - n), \tag{9}$$

with $\mathrm{comb}(\nu) = \sum_{n=-\infty}^{\infty} \delta(\nu - n)$ (δ is the Kronecker function). At $\nu = 0$ the left-hand side is a sample of $\hat{r}(\nu)$ at integer-spaced intervals, while the right-hand side is a single spike. Thus the DC value at $\hat{r}(0)$ is not aliased, that is a constant signal is exactly reproduced (flat field response). An analysis of the forms of reconstruction error is given in [23], though a mathematical treatment is avoided.

III. The Effect of Repeated Resampling

A. Interpolation Theory

Interpolation is the fitting of a curve to a sampled set of data, after which the curve is resampled, usually at a different set of points. Frequently in practice the continuous curve, representing the reconstructed image or in 1-D the scanline, is a theoretical construct as only the values of the interpolation curve at the resampled point are considered. 1-D interpolation[12] to achieve resampling can be represented as the convolution in the spatial domain of an interpolation function with a sampled signal:

$$g_r(x) = g_s(x) * r(x), \tag{10}$$

where $r(x)$ is the interpolation kernel, $g_s(x)$ is the sampled signal, and $g_r(x)$ is the theoretical reconstructed image. The original signal is the signal after sensor degradation has been introduced, as commonly represented by convolution with a point spread function (PSF). That is to say, the original signal is $g(x) = g_i(x) * h(x)$, with $g_i(x)$ the input signal and $h(x)$ the point spread function of the sensor system. For simplicity, the image functions and interpolation functions are assumed to be defined in the real field, when applied in the spatial domain. In [21] the mean radiometric or frequency-related error is defined, under certain circumstances, as

[12] An alternative approach to the one developed here would make use of the z-transform, which provides a natural representation of a sampled signal [28].

$$E[\epsilon^2(x)] = \int_{-\infty}^{\infty} |g(x) - g_r(x)|^2 dx. \tag{11}$$

This represents the combined error from sampling and then reconstructing with a designated interpolation kernel. The following develops in somewhat more detail the theory first given in [21]. By Rayleigh's theorem [17] $E[\epsilon^2]$ has the same value in the frequency domain:

$$E[\epsilon^2(x)] = \int_{-\infty}^{\infty} |\mathcal{F}(g(x) - g_r(x))|^2 d\nu, \tag{12}$$

with $\mathcal{F}$ a Fourier transform operator, which is linear. Sampling is commonly modelled by means of

$$g_s(x) = g(x)\mathrm{comb}(x) = \sum_{n=-\infty}^{\infty} g(x)\delta(x-n), \tag{13}$$

where the Kronecker function is

$$\delta(x) = \begin{cases} 1, & x = 0 \\ 0, & \text{otherwise} \end{cases} \tag{14}$$

Equation 12 can be re-written as

$$E[\epsilon^2] = \int_{-\infty}^{\infty} \left| \mathcal{F}\left(g(x) - \left(\sum_{n=-\infty}^{\infty} g(x)\delta(x-n) \right) * r(x) \right) \right|^2 d\nu, \tag{15}$$

which becomes

$$E[\epsilon^2] = \int_{-\infty}^{\infty} \left| \hat{g}(\nu) - \mathcal{F}\left(\sum_{n=-\infty}^{\infty} g(x)\delta(x-n) \right) \hat{r}(\nu) \right|^2 d\nu. \tag{16}$$

Using the properties of the impulse generalized function we have

$$E[\epsilon^2] = \int_{-\infty}^{\infty} \left| \hat{g}(\nu) - \sum_{n=-\infty}^{\infty} \hat{g}(\nu - n)\hat{r}(\nu) \right|^2 d\nu. \tag{17}$$

Expanding the modulus gives a daunting expression:

$$\begin{aligned} E[\epsilon^2] \;=\; & \int_{-\infty}^{\infty} |\hat{g}(\nu)|^2 - \hat{g}(\nu) \sum_{n=-\infty}^{\infty} \hat{g}^*(\nu - n)\hat{r}^*(\nu) - \hat{g}^*(\nu) \sum_{\infty}^{\infty} \hat{g}(\nu - n)\hat{r}(\nu) \\ & + |\hat{g}(\nu)|^2 \sum_{\underline{n}=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} \hat{r}(\nu - \underline{n})\hat{r}^*(\nu - n) d\nu. \end{aligned} \tag{18}$$

However, if $|\nu| \leq 0.5 = \nu_{nyquist}$, then all terms not involving $n = 0$ drop out, except in the double summation term, where only terms with $n = \underline{n}$ survive. The integral is zero outside these limits if and only if the signal itself is band-limited[13] at the critical frequency. After judicious substitutions, one arrives at:

$$\begin{aligned} E[\epsilon^2] &= \int_{-\nu_{nyquist}}^{\nu_{nyquist}} |\hat{g}(\nu)|^2 \left[1 - \hat{r}^*(\nu) - \hat{r}(\nu) + \sum_{n=-\infty}^{\infty} |\hat{r}(\nu - n)|^2\right] d\nu \\ &= \int_{-\nu_{nyquist}}^{\nu_{nyquist}} |\hat{g}(\nu)|^2 \left[1 - 2\hat{r}(\nu) + \sum_{n=-\infty}^{\infty} |\hat{r}(\nu - n)|^2\right] d\nu \qquad (19) \\ &= \int_{-\frac{1}{2}}^{\frac{1}{2}} |\hat{g}(\nu)|^2 e^2 d\nu \qquad (20) \end{aligned}$$

with $r(x)$ real and even (with transform real and even). e^2 is an image-independent measure, which is also called the radiometric error function. In some cases, for the purpose of calculating the function $E[\epsilon^2]$, equation 20 is convenient, as only one function, the transform, is required in the computer code. Notice that negative and positive terms are needed for the calculation of the infinite summation of equation 19 as the terms are not symmetric. Comparison with the closed form of the summation, developed next, reveals that convergence is rapid.

Through the Convolution Theorem and since the comb function is its own Fourier transform, it is the case that

$$\begin{aligned} \sum_{n=-\infty}^{\infty} |\hat{r}(v - n)|^2 &= \hat{r}^2(\nu) * \mathrm{comb}(\nu) \\ &= \int_{-\frac{1}{2}}^{\frac{1}{2}} r(x) * r(x) \sum_{n=-\infty}^{\infty} \delta(x - n)\, exp(-2\pi i x \nu) dx \\ &= \sum_{n=-\infty}^{\infty} r(x) * r(x)(n) \exp(-2\pi i n \nu) \\ &= r(x) * r(x)(0) + \sum_{n=1}^{\infty} r(x) * r(x)(n) \cos(2\pi n \nu), \qquad (21) \end{aligned}$$

where $r(x) * r(x)(n)$ is the self-convolution of $r(x)$ evaluated at n. The terms to be considered extend to the support region of the convolution integral, which is best established

[13]This naturally means band-limited *before* sampling takes place.

graphically. Equation 21 has the advantage, for power series expansions, that the summation involves only expansions of the cosine function but otherwise calculation of the closed-form for other than elementary kernels is onerous.

B. Incorporating Image Phase-Dependency

The radiometric error function

$$e^2(\nu) = 1 - 2\hat{r}(\nu) + r(x) * r(x)(0) + 2\sum_{n=1}^{\infty} r(x) * r(x)(n)\cos(2\pi n\nu) \tag{22}$$

is an indicator of the blur introduced into the frequency response of the reconstructed image, but its development in Section A has not accounted for image phase. Differential image-phase noise occurs because of the relative position of an image feature to the resampling grid. Contrary to the application of a convolution kernel for the purpose of correcting blur introduced by the sensor, applying an interpolation function is sensitive to image phase through the nature of the sampling process.[14] The analysis confirming that equation 22 is an average which accounts for image phase proceeds as follows. As $\epsilon^2(u)$ is periodic in phase-shift u for unit sampling interval, it can be represented [20] by a Fourier series:

$$\epsilon^2(u) = \sum_{m=-\infty}^{\infty} a_m \exp(2\pi umi), \tag{23}$$

with $a_m = \int_0^1 \epsilon^2 \exp(-2\pi umi)du$ and assuming $E[\epsilon^2] = a_0$. The Fourier coefficients are found to be (refer to Appendix A)

$$a_m = \int_{-\infty}^{\infty} [\delta(m) - \hat{r}(\nu) - \hat{r}^*(\nu - m) + \sum_{n=-\infty}^{\infty} |\hat{r}(\nu - n)|^2]\ \hat{g}(\nu)\hat{g}^*(\nu - m)d\nu, \tag{24}$$

with $m = 0, 1, 2 \ldots$. Re-arranging the order of integration and summation in the term for a_0 gives

$$\begin{aligned} E[\epsilon^2] \quad = \quad & \int_{-\infty}^{\infty}\int_0^1 \sum_{m=-\infty}^{\infty}[\delta(m) - \hat{r}(\nu) - \hat{r}^*(\nu) \\ & + \sum_{n=-\infty}^{\infty} |\hat{r}(\nu - n)|^2]\ \hat{g}(\nu)\hat{g}^*(v - m)\exp(2\pi umi)dud\nu \end{aligned} \tag{25}$$

[14]Otherwise one might simply take the frequency response of the interpolation kernel as characterizing the response.

All terms vanish (through the periodicity of the complex exponential), except for $m = 0$, to give the same error term as equation 20. However, this analysis is applicable only to a uniform distribution, where $E[\epsilon^2] = a_0 = \int_0^1 \epsilon^2(u)du$. In general, there is no reason to suppose that the phase-shift will be other than randomly distributed, but referring to the common overhead views of rectilinear field boundaries, such as in the images required for remotely-sensed archæology [15], this hypothesis is difficult to sustain. If u is uniformly distributed, then the sum given in equation 23 less a_0 will have zero mean when it is averaged over all phase-shifts. The additional terms will vanish too if the image is band-limited to $|\nu| < 0.5$, since equation 24 becomes zero in this case (because of the non-overlapping convolution term).

C. An Extension to Repeated Resampling

Let us now turn to the case of repeated resampling. As outlined in Section II, assume interpolation takes place twice in succession. If there are any sharp edges in the sampled image, it is likely that the image before sampling was not band-limited (or the sample set was created artificially without pre-filtering). Thus, on interpolating for a second time, the phase-related error is given by equation 23, but equation 20 can still be relied upon to give the mean error. The order of events is: sample the original signal; interpolate with the first kernel; resample at some offset m; interpolate a second time to reconstruct the signal. (The reconstructed image is actually held in resampled form.) m is the fractional difference between the original sampling grid and the first resampling grid. The reconstructed signal is now given by

$$g_r(x) = \left((g_s * r_1(x)) \sum_{n=-\infty}^{\infty} \delta(x - m - n) \right) * r_2(x), \ m \in [0, 1), \tag{26}$$

where r_1 is the first interpolation kernel to be applied and r_2 is the second kernel to be applied. Resampling has taken place with unit spacing. Due to the associativity of convolution, equation 26 can be rewritten as

$$g_r(x) = g_s * r_3(x). \tag{27}$$

The Fourier transform of the aggregate interpolation kernel $r_3(x)$ is

$$\begin{aligned}
\hat{r}_3(\nu) &= \hat{r}_2(\nu)\int_{-\infty}^{\infty} r_1(x)\sum_{n=-\infty}^{\infty}\delta(x-m-n)\exp(-2\pi i\nu x)dx \\
&= \hat{r}_2(\nu)\exp(-2\pi i\nu m)\int_{-\infty}^{\infty} r_1(x+m)\sum_{n=-\infty}^{\infty}\delta(x-n)\exp(-2\pi i\nu x)dx \\
&= \hat{r}_2(\nu)\exp(-2\pi i\nu m)\int_{-\frac{1}{2}}^{\frac{1}{2}}\sum_{n=-\infty}^{\infty} r_1(n+m)\exp(-2\pi i\nu n)dx \\
&= \hat{r}_2(\nu)\exp(-2\pi i\nu m)\sum_{n=-\infty}^{\infty} r_1(n+m)\exp(-2\pi i\nu n). \qquad (28)
\end{aligned}$$

$\hat{r}_3(\nu)$ in general has an imaginary component and, needless to say, only approximates the benign features associated with once-applied kernels. The new kernel can be used in equation 23, which is rewritten as

$$\begin{aligned}
e^2(\nu) &= 1-2\mathcal{R}[\hat{r}_3(\nu)]+\sum_{n=-\infty}^{\infty}|\hat{r}_3(\nu-n)|^2 \\
&= 1-2\hat{r}_2(\nu)\left[\sum_{n=-\infty}^{\infty} r_1(n+m)\cos(2\pi\nu(n+m))\right] \\
&\quad +\sum_{n=-\infty}^{\infty}\hat{r}_2^2(\nu-n)\left|\sum_{\underline{n}=-\infty}^{\infty} r_1(n+m)\exp(-2\pi i(\nu-n)\underline{n})\right|^2 \\
&= 1-2\hat{r}_2(\nu)\left[\sum_{n=-\infty}^{\infty} r_1(n+m)\cos(2\pi\nu(n+m))\right] \\
&\quad +\left|\sum_{\underline{n}=-\infty}^{\infty} r_1(n+m)\exp(-2\pi i\nu\underline{n})\right|^2\sum_{n=-\infty}^{\infty}\hat{r}_2^2(\nu-n), \qquad (29)
\end{aligned}$$

where the operator $\mathcal{R}$ returns the real part of a complex number. If we set $m = 0$ in equation 29, then we recover the once-applied result from a twice-applied kernel. This is as might be expected since the reconstructed image after one interpolation is exact at the sample points. The radiometric error for nearest neighbor interpolation is given by

$$e^2(\nu) = \begin{cases} 1-2\mathrm{sinc}(\nu)\cos(2\pi\nu m)+1, & m \le 0.5 \\ 1-2\mathrm{sinc}(\nu)\cos(2\pi\nu(m-1))+1, & m > 0.5. \end{cases} \qquad (30)$$

The theoretical once-applied and twice-applied errors are compared in Figure 8, when the error is greater the further the departure from the original sampling points. Other interpolation kernels when twice applied do not have such a convenient form but are constructed from the terms given in Table I. In Figure 9, linear interpolation is applied twice

in succession and compared to a once-applied interpolation. As m is moved away from the original sampling points, the error, which is symmetrical about 0.5, increases, though for small values of m this is not critical and is less than the effect of using once-applied nearest-neighbor interpolation. Turning to twice-applied cubic-interpolation (Figure 10), the error is reduced in the low frequency region compared to that of linear interpolation. When the three twice-applied interpolation kernels are compared in the region of large error at the same scale (Figure 11), the original order of ranking for once-applied kernels in the low frequency region does not change for twice-applied kernels. It is also possible to use a different interpolation kernel at each stage of the rotation algorithm. For instance, an initial cubic interpolation might be used before the region of support is distorted by the subsequent vertical skew. In Figure 12, a number of interpolation kernel combinations are compared with repeating the same kernel, when it is indicated that using a preliminary cubic interpolation should be superior. The effect of reversing the order of kernel application is examined in Figure 13. Starting with a cubic interpolation is found to be superior, though in the case of cubic convolution and linear interpolation combinations there is little to distinguish the two in the low frequency region.

IV. Results From Applying Interpolation Kernel Combinations

Various combinations of interpolation kernels were applied to various images with a mixture of edges and smoother transitions, Figure 1. The images, with maximum gray-tone range $0 - 255$, were clipped and clamped, though this made little visible difference to the images and, except for cubic convolution, had marginal effect on the results. A standard image (the Cameraman) has been included. It will be recalled that the model of Section III applies to an image with a uniform distribution of edges and not to an image like the test image of Section II, which produces large errors. The most noticable distortion occurs at edge regions in the images. A significant additional area of discrepancy is at the boundaries. This should not be a problem in image processing as the boundaries are generally discarded due to sensor error, but it does weigh any comparison. A further difficulty in making comparisons is if the one pass algorithm uses floating point (f.p.) but the thrice-skewed algorithm were to pass back and forth between f.p. and integer arithmetic. It was found that provided rounding, rather than truncation, was used the

results were virtually identical. A root mean square (RMS) measure was used to compare images, *viz.*

$$\text{RMS} = \sum_{i}^{N} \sum_{j}^{M} \frac{\sqrt{|(f(i,j) - g(i,j)|^2}}{NM}, \tag{31}$$

where f and g are the image intensity functions of the two images being compared, which are of size $N \times M$. It should be borne in mind that the comparison is now between two different methods of sampling, rather than the combined error of interpolation and resampling as described in Section III. Table II gathers the results together, where comparisons have been made between a one-pass routine using successively linear interpolation and cubic interpolation. When using cubic convolution, at scanline boundaries a weighted average of the adjoining pixels was used as described in [22]. It was more convenient when using a two-dimensional interpolation to use wrap-around for cubic convolution but this does not obviate the fact that there were significant differences between the thrice-skew algorithm and what is achievable by a one pass algorithm employing cubic convolution (the comparison is uniformly less favourable when made with cubic convolution rather than linear interpolation). Images with sharper edges, such as image one, score more unfavourably by the RMS measure, while images of a similar type (though not in appearance), such as images two and four, have a similar pattern to the results. The use of nearest neighbor interpolation caused significant distortion except when combined with linear interpolation. Changing the order of the interpolations did not have much effect. Where cubic convolution was included at the second or third skew the results were poorer compared to linear interpolation. The strong implication is that distortion of the region of support towards a stepped ellipse, as identified in Section II, overrides any advantage to be gained by using a more efficient interpolator. For cubic convolution the region of support for each one dimensional interpolation is larger; thus the distortion may be larger. However, in the case of cubic convolution followed by two stages of linear interpolation, for 50° rotations, the results were comparable or better than using just linear interpolation. At the first stage of the algorithm, the region of support is not distorted so that using cubic convolution may give an improved representation before going into subsequent stages.

It might be thought that using an improved interpolation kernel over a region of support of two pixels size would improve matters. Setting $\alpha = 0$ in equation 6 gives a Hermite cubic,

which on consideration of the kernel's frequency response alone might seem an improvement on linear interpolation [5].[15] Table III shows that for the thrice-skew algorithm the Hermite interpolation kernel behaves much like a weaker linear interpolator when compared to one-pass linear interpolation, though it performs more favourably when compared to one-pass cubic convolution. Why this might be is apparent from the error plot for one-pass interpolation (Figure 14) and in particular from the low-frequency error (Figure 15), which agrees with the curves reported in [21]. Hermite interpolation has a similar failing to cubic convolution interpolation with $\alpha = -1$ in this respect. Comparing twice-applied kernels, (Figures 16 & 17) shows that the twice-applied Hermite interpolator is again predicted to be worse at low frequencies.

Table 4 gives the behaviour of repeated interpolation with change of angle. Increasing the angle in the main increases the error, apart from the lacuna at 30 degrees[16]. Angle-dependent error may arise principally from increased distortion of the interpolating filter's region of support as remarked in Section III. The parallel algorithm also increases in cost due to increased communication [12], which implies a restriction to small angle rotations. The comparable results for the Hermite interpolator, which mirror the linear interpolation results, are also included.

V. Conclusion

This paper has drawn attention to the interpolation problem that arises in applying scanline algorithms, by means of a case study of image rotation. Scanline algorithms provide a suitable decomposition for the parallelization of affine transforms. If the algorithms are used for image processing applications where a one pixel error may be significant, such as medical imagery and perhaps in remote sensing, the problem of satisfactory interpolation should be critically examined. Scanline algorithms trade off a reduction in trigonometric

[15]The Hermite interpolating kernels, $\{h_i\}$, are given by:

$$g_r(x) = \sum_{n=-\infty}^{\infty} (g(n)h_0(x-n) + g'(n)h_1(x-n)), \tag{32}$$

where g' is the derivative w.r.t. x [23]. For the present purposes, g' is set to 0 because it is unavailable, but see [39] on this and related issues.

[16]Since $\sin 30^\circ = 0.5$ even lines will experience a skew with no fractional component, so that the new value is not ameliorated by interpolation with neighboring pixels.

calculations for a decrease in image quality.[17] The paper provides an appropriate error analysis framework, after previous work in this field but with additional computational detail. The analysis is extended to the case of an extra resampling. We have shown that satisfactory results from scanline algorithms are dependent on the nature of the images. This dependency may render them unsuitable for complex forms of imagery — notably — some forms of medical imaging.

An error analysis of the problem shows that two difficulties arise: the support for the interpolation kernels is distorted and resampling takes place at varying distance from the original sampling point. The distortion of the support region is a significant effect, which increases with the angle of rotation. The theory predicts that the twice-applied cubic convolution interpolation kernel will still be superior, though error will accrue as the resampling point grid moves away from the original sampling grid. For thrice-skew rotation, the measured results did not match this expectation, except when cubic convolution interpolation was confined to the first skew. Using repeated linear interpolation was superior, presumably because the region of support was smaller and therefore less distorted by the skews. By using a purportedly improved interpolator (Hermite cubic) with a small support region it was thought that an improvement might occur. Why this was not the case for the Hermite interpolator was revealed by the error analysis methodology described in the paper. It should be noted that a comparatively large angle of rotation was used in the tests. Where smaller angles of rotation and smoother images are used the measured error is reduced. It remains the case that image rotation by the thrice-skew algorithm will introduce increased image-phase-related error, which is not correctable. The analysis methodology developed here may be appropriate for other scanline algorithms without large distortions of the region of support, for instance those used to interpolate large images which do not easily fit into the memory of the computer.

Further work would be to analyse the error arising from more than two resamples. The radiometric analysis establishes only the mean error and one may well want to establish the error arising if the image phase shift had a Gaussian distribution centered upon a particular location. Another line of attack is to consider whether other than 'off-the-peg'

[17] It is possible to reduce trigonometric calculations, by an incremental method, at the risk of floating-point rounding errors, but the algorithm is not suitable for distributed-memory parallelizations.

interpolation kernels may be appropriate for these algorithms. For those not convinced of the merits of a scanline algorithm for their application but still requiring an improvement on linear interpolation for sequential implementations, one may consider 'best-of-three' point interpolation [40] at a modest increased cost in computation.

Appendix

I. Finding Fourier Coefficients

To find the Fourier coefficients, a_m, set

$$\mathcal{F}(g(x-u) - g_r(x;u)) = \exp(-2\pi u\nu i) \sum_{n=-\infty}^{\infty} [\delta(n) - \hat{r}(\nu)]\ \hat{g}(\nu - n)\exp(2\pi uni), \qquad (33)$$

where the Modulation Theorem has been used to account for the phase-shift by u. If there is phase-dependency,

$$\epsilon^2(u) = \int_{-\infty}^{\infty} |g(x-u) - g_r(x;u)|^2 dx, \qquad (34)$$

which can be compared with equation 11, where there is no phase-dependency. Recalling the expression for a_m, we can rearrange the order of integration and convert to the frequency domain by means of Rayleigh's Theorem:

$$\begin{aligned} a_m &= \textstyle\int_0^1 \epsilon^2(u)\exp(-2\pi umi)du \\ &= \textstyle\int_{-\infty}^{\infty}\int_0^1 |\exp(-2\pi u\nu i)\sum_{n=-\infty}^{\infty}[\delta(n) - \hat{r}(\nu)]\hat{g}(\nu-n)\exp(2\pi uni)|^2\exp(-2\pi umi)dud\nu \\ &= \textstyle\int_{-\infty}^{\infty}\sum_{\underline{n}=-\infty}^{\infty}\sum_{n=-\infty}^{\infty}[\delta(n) - \hat{r}(\nu)][\delta(\underline{n}) - \hat{r}^*(\nu)]\hat{g}(\nu-n)\hat{g}^*(\nu-n)dud\nu, \end{aligned} \qquad (35)$$

when equation 35 is only non-zero for $m = n - \underline{n}$, because of the periodicity of the complex exponential. To reach equation 24, it is necessary to expand the multiplicand in equation 35, when it will emerge that each term has an extra survival condition by reason of the δ functions. Substituting in terms of m will, with suitable re-arrangement, reach equation 24, which was previously given in [20].

Acknowledgement

This work was carried out as part of project IED3/1/2171 ("Parallel Reconfigurable Image Processing Systems"). The authors acknowledge helpful suggestions from the

anonymous referees towards revising this paper for publication.

References

[1] P. A. Fowler, C. E. Casey, G. G. Cameron, M. A. Foster, and C. H. Knight, "Cyclic changes in composition and volume of the breast during the menstrual cycle, measured by magnetic resonance images," *British Journal of Obstetrics and Gynæcology*, vol. 97, pp. 595–602, 1990.

[2] G. Simm, "Selection for lean meat production in sheep," in *Recent Advances in Sheep and Goat Research* (A. W. Speedy, ed.), pp. 193–215, CAB International, 1992.

[3] M. J. Bentum, B. M. A. Lichtenbelt, and T. Malzbender, "Frequency analysis of gradient estimators in volume rendering," *IEEE Transactions in Visualization and Computer Graphics*, vol. VCG-2, pp. 242–253, September 1996.

[4] J. A. Parker, R. V. Kenyon, and D. E. Troxel, "Comparison of interpolating methods for image resampling," *IEEE Transactions on Medical Imaging*, vol. MI-2, March 1983.

[5] E. Maeland, "On the comparison of interpolation methods," *IEEE Transactions on Medical Imaging*, vol. MI-7, pp. 213–217, September 1988.

[6] G. W. Horgan, A. M. Creasey, and B. Fenton, "Superimposing two-dimensional gels to study genetic variation in malaria parasites," *Electrophoresis*, vol. 13, pp. 871–875, 1992.

[7] E. Catmull and A. Smith, "3-D transformations of images in scanline order," *Computer Graphics*, vol. 14, no. 3, pp. 279–285, 1980.

[8] C. Weiman, "Continuous anti-aliased rotation and zoom of raster images," *Computer Graphics*, vol. 14, no. 3, pp. 286–293, 1980.

[9] A. Tanaka, M. Kameyama, S. Kazama, and O. Watanabe, "A rotation method for raster image using skew transformation," in *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 272–277, 1986.

[10] N. Tsuchida, Y. Yamada, and M. Ueda, "Hardware for image rotation by twice skew transformations," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 35, no. 4, pp. 527–531, 1987.

[11] R. Schalkoff and H. Nag, "Decomposition and parallel architecture for the geometric transformation of digital images," *Image and Vision Computing*, vol. 9, no. 5, pp. 275–284, 1991.

[12] M. Fleury, L. Hayat, and A. F. Clark, "On parallelizing grey-scale coordinate transforms," *IEE Part I (Vision, Image and Signal Processing)*, vol. 142, pp. 207–212, August 1995.

[13] G. Wolberg, *Digital Image Warping*. IEEE Computer Society, Los Alamitos, CA, 1990.

[14] J. Foley, A. van Dam, S. Feiner, and J. Hughes, *Computer Graphics: Principles and Practice*. Addison-Wesley, Reading, MA, 1990.

[15] I. Scollar, A. Tabbagh, A. Hesse, and I. Herzog, *Archæological Prospecting and Remote Sensing*. Cambridge University Press, Cambridge, UK, 1990.

[16] J. Robinson, "Efficient general-purpose image compression with binary tree predictive coding," 1994. Accepted for publication in IEEE Transactions on Image Processing.

[17] R. N. Bracewell, *The Fourier Transform and Its Applications*. McGraw-Hill, New York, $2^{\underline{nd}}$ ed., 1986.

[18] B. J. Jähne, *Digital Image Processing*. Springer, Berlin, $2^{\underline{nd}}$ ed., 1993.

[19] I. E. Abdou and K. Y. Wong, "Analysis of linear interpolation schemes for bi-level image applications," *IBM Journal of Research and Development*, vol. 26, pp. 667–680, November 1982.

[20] S. K. Park and R. A. Schowengerdt, "Image sampling, reconstruction, and the effect of sample-scene phasing," *Journal of Applied Optics*, vol. 21, pp. 3142–3151, September 1982.

[21] S. K. Park and R. A. Schowengerdt, "Image reconstruction by parametric cubic convolution," *Computer Vision, Graphics, and Image Processing*, vol. 23, pp. 258–272, 1983.

[22] R. G. Keys, "Cubic convolution interpolation for digital images," *Transactions on Acoustics, Speech and Signal Processing*, vol. ASSP-29, pp. 1153–1160, December 1981.

[23] D. P. Mitchell and A. N. Netravili, "Reconstruction filters in computer graphics," *Computer Graphics*, vol. 22, pp. 221–227, August 1988.

[24] J. Ward and D. R. Cok, "Resampling algorithms for image resizing and rotation," *SPIE Digital Image Processing Applications*, vol. 1075, pp. 260–269, 1989.

[25] *International Standard 12087-2.* International Standard Organization, Geneva, Switzerland, 1994. Part 2, Annex H.

[26] E. R. Davies, "Circularity—a new principle underlying the design of accurate edge orientation operators," *Image and Vision Computing*, vol. 3, no. 3, pp. 134–142, 1984.

[27] T. Thong, "Digital image processing test patterns," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. ASSP-31, pp. 763–766, June 1983.

[28] R. W. Schafer and R. L. Rabiner, "A digital signal processing approach to interpolation," *Proceedings of the IEEE*, vol. 61, no. 6, pp. 692–702, 1973.

[29] P. L. Walker, *The Theory of Fourier Series and Integrals.* Wiley, Chichester, UK, 1986.

[30] W. F. Schreiber and D. E. Troxel, "Transformation between continuous and discrete representations of images: A perceptual approach," *IEEE Transactions on Pattern Recognition and Machine Intelligence*, vol. PAMI-7, pp. 178–186, March 1985.

[31] M. Unser, A. Aldroubi, and M. Eden, "Enlargement or reduction of digital images with minimum loss of information," *IEEE Transactions on Image Processing*, vol. IP-4, pp. 247–258, March 1995.

[32] K. Jensen and D. Anastassiou, "Subpixel edge localization and the interpolation of still images," *IEEE Transactions on Image Processing*, vol. IP-4, pp. 285–295, March 1995.

[33] R. Bernstein, "Digital signal processing of earth observation sensor data," *IBM Journal of Research and Development*, vol. 20, pp. 40–57, 1976.

[34] S. H. Hou and C. A. Andrews, "Cubic splines for image interpolation and digital filtering," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. ASSP-26, pp. 508–517, December 1978.

[35] G. Arfken, *Mathematical Methods for Physicists.* Academic Press, New York, $3^{\underline{rd}}$ ed., 1985.

[36] D. C. Champeney, *A Handbook of Fourier Theorems.* Cambridge University Press, Cambridge, UK, 1987.

[37] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing.* Prentice-Hall, Englewood Cliffs, NJ, 1989.

[38] G. H. Hardy and E. M. Wright, *An Introduction to the Theory of Numbers.* Clarendon Press, Oxford, UK, $5^{\underline{th}}$ ed., 1979.

[39] S. E. Reichenbach and S. K. Park, "Two-parameter cubic convolution for image reconstruction," *SPIE Visual Communications and Image Processing*, vol. 1199, no. 4, pp. 833–840, 1989.

[40] P. R. Smith, "Bilinear interpolation of digital images," *Ultramicroscopy*, vol. 6, pp. 201–204, 1981.

[41] A. A. Glasbey and G. W. Horgan, *Image Analysis for the Biological Sciences.* Wiley, Chichester, UK, 1995.

List of figures

List of tables

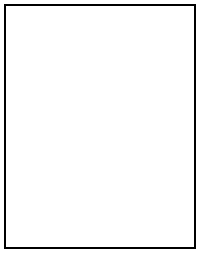

Martin Fleury holds a BA degree in Modern History from Oxford University. After a teaching career, he obtained a BA from the Open University in Maths. and Physics in 1988. He received an MSc in Astrophysics in 1990 from the University of London. He also took an MSc with distinction in Parallel Computing Systems from the University of the South-West in 1991. He is a candidate for PhD degree at the University of Essex. Between 1992 and 1995 Martin Fleury worked in research at the University of Essex on image-processing for reconfigurable parallel systems. He has written a library of distributed systems image-processing software. Martin has published articles in the areas of low-level image-processing algorithms, parallel systems, and the performance estimation of parallel systems. He is currently employed as a Senior Research Officer on a project concerned with parallel embedded signal-processing applications. Martin is a member of the IEEE Computer Society.

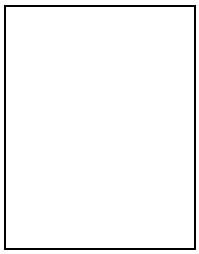

Dr. Adrian F. Clark obtained a BSc degree in Physics from the University of Newcastle upon Type in 1979 and a PhD in image processing from the University of London in 1983. He is a Senior Lecturer in Information Technology in the Department of Electronic Systems Engineering at the University of Essex, which he has been a member of since 1988. He has worked on image restoration, analysis, coding and synthesis techniques applied to a wide variety of both single-band and multi-spectral imagery. His main research interests are in algorithms and software techniques for image and spectral processing, on both serial and parallel hardware. Dr. Clark acted as Convenor of a UK panel of experts involved in the development of the International Standard for Image Processing and Interchange (ISO 12087), and was editor of the first part of that Standard. He is a former Secretary of the British Machine Vision Association and Chairman of Technical Committee 5 ('Benchmarking and Software') of the International Association for Pattern Recognition. He is a Chartered Physicist and a member of the IEEE.