One-Dimensional Resampling of Digital Imagery to Sparse Output Grids

M. D. Craig and A. A. Green
CSIRO Division of Mineral Physics and Mineralogy, North Ryde, NSW 2113, Australia

ABSTRACT: Although the problem of resampling digital imagery to a rotated grid can be solved using only one-dimensional interpolation, this solution is limited to the case where the output grid is sufficiently dense within the original image. The present paper extends the one-dimensional method to sparse output grids. Because of the less rapid increase in the number of samples that must be convolved to produce equivalent results, it is in sparse resampling that the computational advantages of one-dimensional over two-dimensional interpolation are most significant.

INTRODUCTION

IN TWO-DIMENSIONAL RESAMPLING of digital imagery, output lines are constructed successively by two-dimensional interpolation of input pixel values. Where large-angle rotation is involved, many adjacent input lines are needed for the construction of a single output line, an effect that can lead to inordinate disk access in the minicomputer processing of large images.

Friedmann (1981a,b) has shown how this problem may be overcome, at the expense of two 90° rotations of the entire image. He gives an algorithm in which single-pass, two-dimensional interpolation is replaced by a three-pass operation involving only one-dimensional interpolations. In the (theoretical) case of imagery acquired through Nyquist-rate sampling in two orthogonal directions, the resampling avoids both truncation and aliasing of input frequencies. However, it is assumed that the output grid is sufficiently dense within the input image for such retention of frequencies to be possible.

The present paper aims to extend the algorithm to the alternative case of sparse output grids. Whatever method one follows, a large convolving kernel must then be used to ensure radiometric accuracy of the resampled image. However, for a sequence of output grids of increasing sparseness, the kernel size increases much more rapidly for a two-dimensional kernel than for the equivalent one-dimensional kernel (i.e., quadratic versus linear variation). Efficiency in resampling to a sparse output grid is, thus, one of the greatest inherent advantages of the one-dimensional method.

THE FREQUENCY DOMAIN

As in Friedmann's work (Friedmann, 1981a,b), the analysis depends on considerations in Fourier space. The resampling process attempts to avoid aliasing, while retaining input scene detail as well as possible. We perpetuate Friedmann's assumption that the spatial frequencies present in the data are those that would result from Nyquist-rate sampling on a square grid.

Sparse resampling can be effected by combining the Friedmann process (dense resampling involving rotation) with a process for sparse resampling of an image parallel to itself. However, two combinations are possible:

(a) First, resample (with smoothing of input frequencies) to a sparse grid which is parallel to the input grid, and within which the output grid is dense; then apply Friedmann's algorithm.
(b) Alternatively, resample (with rotation) to a dense grid parallel to the final output grid; then effect size reduction and smoothing in a subsequent parallel resampling operation.

Although geometrically equivalent, the alternative procedures of rescaling before or after rotation differ in the spatial frequencies that can be retained, as considered in greater detail below. Thus, in the case of a 45° rotation where the total number of pixels is halved, excessive smoothing must be performed in the first step of method (a) if aliasing is to be avoided. Because of the consequent loss of scene detail, method (b) is superior in this instance.

We treat exclusively the case of resampling between two square grids, and begin by establishing notation. Our convention regarding Fourier transforms is that, if (s_1, s_2) and (t_1, t_2) denote Cartesian coordinates in frequency space, and in image space, respectively, then

$$\Phi(s_1,s_2) = \iint \phi(t_1, t_2) \exp\{-2\pi i\,(s_1 t_1 + s_2 t_2)\}\, dt_1\, dt_2$$

is the direct Fourier transform of the function $\phi(t_1, t_2)$. The spectrum for an image sampled on a given lattice is then a sum of spectra obtained by translating the original spectrum to all points of the *reciprocal* lattice.

Further, let r denote the pixel spacing for the output grid, measured in units of input pixel spacing, and write θ ($0° \leq \theta \leq 45°$) for the absolute angle of rotation between the input and output grids. The output grid is considered to be *dense* when $r < \cos\theta$, and *sparse* when $r > \cos\theta$. Either class can be widened to incorporate $r = \cos\theta$.

Finally, we write $\mathcal{L}(\mathbf{L}_1, \mathbf{L}_2)$ for the lattice $\mathbf{ZL}_1 + \mathbf{ZL}_2$ generated by linearly independent vectors $\mathbf{L}_1$, $\mathbf{L}_2$ ($\mathbf{Z}$ = integers), and $\mathcal{P}(\mathbf{L}_1,\mathbf{L}_2,c)$ for the parallelogram

$$\{s_1\mathbf{L}_1 + s_2\mathbf{L}_2;\ |s_1|, |s_2| < c\},$$

where $c > 0$. The frequencies (s_1,s_2) within this parallelogram are selected through multiplication by its characteristic (indicator) function, i.e., the function with value unity at interior points and zero at exterior points, represented by a "pill-box" surface.

We first briefly recapitulate the theory underlying resampling to a dense output grid. Thus, Figure 1 shows bases $(\mathbf{i}_m)$, $(\mathbf{j}_m)$ for the input and output grids ($m = 1,2$), and $(\mathbf{k}_m)$ for the "intermediate" grid, formed by the intersections of output lines with input columns. The input basis is orthonormal, while the output system (obtained by a rotation $\pm\theta$ and expansion r) is merely orthogonal.

Figure 2 exhibits reciprocal bases $(\mathbf{I}_m)$, $(\mathbf{J}_m)$, $(\mathbf{K}_m)$ in frequency space. Here $\mathbf{I}_m = \mathbf{i}_m$, orthonormal systems being self-reciprocal. Also $|\mathbf{J}_m| = |\mathbf{j}_m|^{-1} = r^{-1}$ so $\mathbf{J}_m = r^{-2}\,\mathbf{j}_m$. Moreover, as indicated in the diagram, we have $\mathbf{K}_1 = \mathbf{J}_1$, $\mathbf{K}_2 = \mathbf{I}_2$.

The input spectrum is assumed to occupy the central square $\mathcal{P}(\mathbf{I}_1, \mathbf{I}_2, 1/2)$ in Figure 2, together with its translates by all vectors of the lattice $\mathcal{L}(\mathbf{I}_1, \mathbf{I}_2)$. Within-column resampling from the input grid $\mathcal{L}(\mathbf{i}_1, \mathbf{i}_2)$ to the intermediate grid $\mathcal{L}(\mathbf{k}_1, \mathbf{k}_2)$ corresponds to replicating values within $\mathcal{P}(\mathbf{I}_1,\mathbf{I}_2, 1/2)$ about all points of the reciprocal intermediate grid $\mathcal{L}(\mathbf{K}_1, \mathbf{K}_2)$. The replicated spectra do not overlap. For, as $r < \cos\theta$ by assumption, the distance $d = \frac{1}{2}\sec\theta$ in Figure 2 is less than $\frac{1}{2}|\mathbf{K}_1| = \frac{1}{2}r^{-1}$, or half the separation between centers of vertically neighboring squares; and no square has been moved in relation to its horizontal neighbors.

PHOTOGRAMMETRIC ENGINEERING AND REMOTE SENSING,
Vol. 53, No. 5, May 1987, pp. 501–506.

0099-1112/87/5305–501$02.25/0

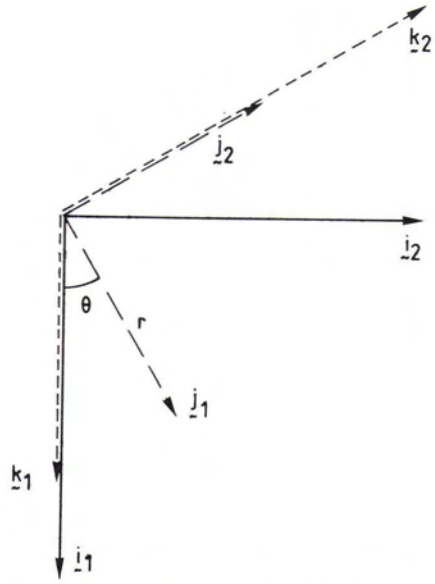

FIG. 1. Basis vectors for input, intermediate, and output sampling lattices—dense case.

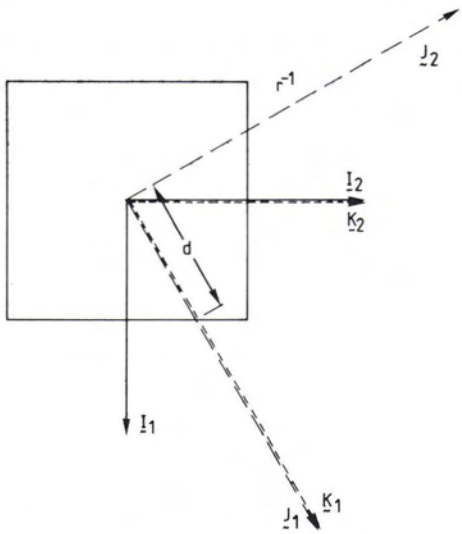

FIG. 2. Reciprocal bases in the frequency domain.

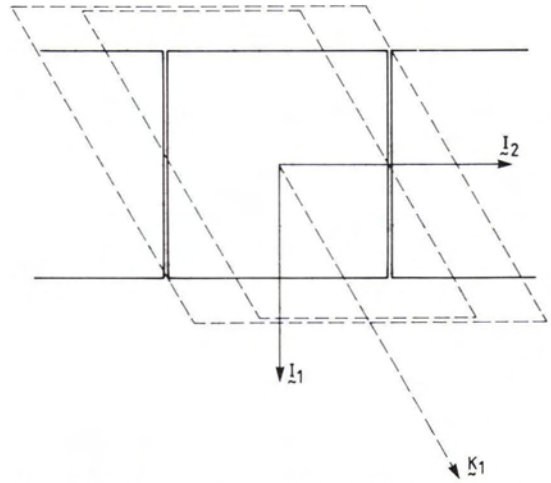

FIG. 3. Friedmann's modified horizontal resampling filter.

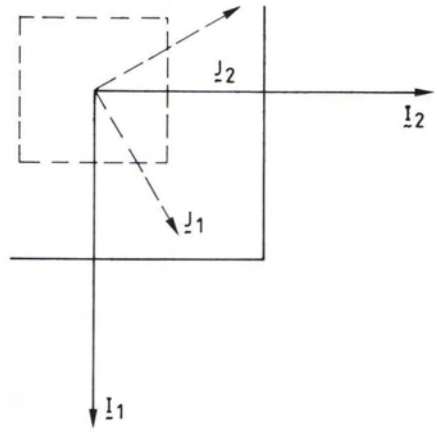

FIG. 4. Filter for parallel resampling before rotation.

Similar considerations apply to the subsequent horizontal resampling. In this case we are selecting frequencies within $\mathcal{P}$ ($\mathbf{K}_1$, $\mathbf{K}_2$, 1/2) (and translating by $\mathcal{L}$ ($\mathbf{J}_1$, $\mathbf{J}_2$)). This region is the inner parallelogram in Figure 3. As drawn here (for the case of counter-clockwise rotation), it excludes input frequencies in the upper right and lower left of the input specturm. The parallelogram may be laterally enlarged to recapture these frequencies (outer parallelogram in Figure 3), but will then intrude still farther into horizontally neighboring squares, resulting in aliasing.

As Friedmann (1981a, b) observed, this problem disappears if we first increase the horizontal separation between the squares by a factor equal to the ratio, $w = 1 + \tan\theta$, of parallelogram widths. The lattice reciprocal to $\mathcal{L}$ ($\mathbf{I}_1$, $w\mathbf{I}_2$) being $\mathcal{L}$ ($\mathbf{i}_1$, $w^{-1}\,\mathbf{i}_2$), this increase entails reduction by the factor w in the horizontal spacing for input pixels, i.e., a preliminary oversampling of the image in this direction.

We proceed now to the case of sparse output grids. Here, some assumed input frequencies need elimination, to preclude aliasing. One-dimensional processing admits only limited forms of "tailoring," but there is still some freedom in choosing the shape of the region of retained frequencies. *Frequency selection by rectangles parallel to the input or output grids corresponds to the respective methods ((a) and (b) above) of parallel resampling before or after rotation.* It seems less natural to rely on frequency-domain filters that align with the reciprocal intermediate lattice, $\mathcal{L}$ ($\mathbf{K}_1$, $\mathbf{K}_2$).

Whichever method is chosen, we may suppose the central region of retained frequencies to be both convex and symmetric. To avoid aliasing, this region must not overlap its translates by vectors of the lattice $\mathcal{L}$ ($\mathbf{J}_1$, $\mathbf{J}_2$). The condition is that it contain in its interior only a single point of the lattice $\mathcal{L}$ (½ $\mathbf{J}_1$, ½ $\mathbf{J}_2$) (*viz.*, the origin — see Appendix). The largest admissible retention region of a given type will be termed a *critical* region for $\mathcal{L}$ ($\mathbf{J}_1$, $\mathbf{J}_2$).

The next two sections analyze the resampling process for critical square regions parallel to the input or output grids, respectively. We then compare the two methods for "superior" retention of input frequencies. It will be noted that the rotation and parallel resampling stages are combined in a single operation, not separate, as described in (a) and (b) above.

As a useful abbreviation for a frequently occurring quantity, write

$$p = r^{-1}(\cos\theta + \sin\theta).$$

PARALLEL RESAMPLING BEFORE ROTATION

The critical square parallel to $\mathcal{L}$ ($\mathbf{I}_1$, $\mathbf{I}_2$) is $\mathcal{P}$ ($\mathbf{I}_1$,$\mathbf{I}_2$, ½ r^{-1} cos θ), its sides bisecting $\mathbf{J}_1$ and $\mathbf{J}_2$ (see Figure 4). Vertical resampling

proceeds as before, except that the space-domain convolving kernel must be broader than the one used in dense resampling. The breadth factor will equal the ratio of numbers of input to output lines, or the ratio, $r \sec \theta$, of the lengths of $\mathbf{k}_1$, $\mathbf{i}_1$. It is also equal to the ratio of vertical projections of input and output frequency regions (Figure 4). The construction of a suitable space-domain filter is treated later.

As regards horizontal resampling (third pass), we note that the smallest parallelogram encompassing the critical square, and aligned with the reciprocal intermediate grid $\mathcal{L}(\mathbf{J}_1,\mathbf{I}_2)$, has width

$$(1 + \tan \theta)\, r^{-1} \cos \theta = p.$$

If greater than unity, this quantity becomes (by analogy with the case of dense resampling) the "oversampling factor" for the first pass of the processing. But it may also, on occasion, be less than unity; in all cases the first pass should result in a horizontal pixel spacing equal to p^{-1}. This pass has the additional task of horizontal low-pass filtering. The kernel breadth factor is $r \sec \theta$, as in vertical resampling.

PARALLEL RESAMPLING AFTER ROTATION

The critical square parallel to $\mathcal{L}(\mathbf{J}_1,\mathbf{J}_2$ is $\mathcal{P}(\mathbf{J}_1,\mathbf{J}_2, ½)$. We can no longer select $\mathbf{K}_1$ equal to $\mathbf{J}_1$ because points of $\mathcal{L}(\mathbf{J}_1,\mathbf{I}_2)$ differing by $\pm\mathbf{J}_1$ have a vertical separation, $r^{-1} \cos \theta$, less than p, the vertical projection of the critical square. Thus, $\mathbf{K}_1 = c\mathbf{J}_1$ where the constant, c, must be chosen greater than unity. Note that $cr^{-1} \cos \theta$ is the vertical projection of $\mathbf{K}_1$. There are, then, two cases to consider ($p = 1$ can be included in either case).

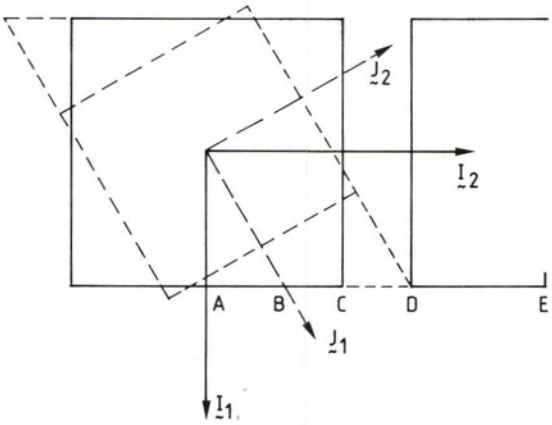

FIG. 5. Filter for parallel resampling after rotation, $p > 1$ case.

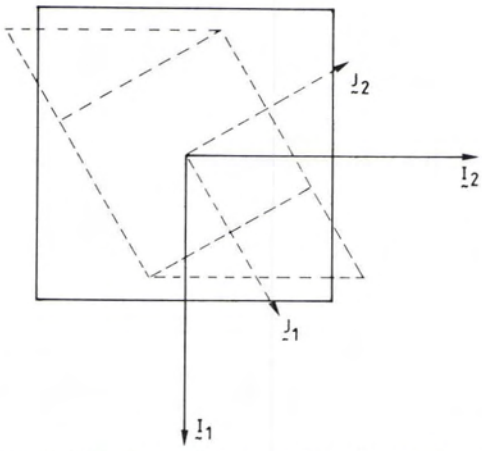

FIG. 6. Filter for parallel resampling after rotation, $p < 1$ case.

(a) $p > 1$.

When the critical square (dashed square in Figure 5) is not interior to the central input square, the minimal choice of c is given by $cr^{-1} \cos \theta = 1$, i.e., $c = r \sec \theta$. The oversampling factor in the first pass (obviously greater than unity, from geometrical considerations; see Figure 5) equals

$$AB + BD + DE = \frac{1}{2} \tan \theta + \frac{1}{2} r^{-1} \sec \theta + \frac{1}{2}$$

$$(\text{note } DE = AC)$$

$$= \frac{1}{2}(1 + pr^2) r^{-1} \sec \theta.$$

This quantity is chosen to prevent intrusion of horizontally neighboring input squares by the parallelogram filter used in horizontal resampling.

The oversampling and vertical resampling stages entail no smoothing of input frequencies. But in horizontal resampling there is a scale reduction given by the ratio of lengths of $\mathbf{j}_2$ and $\mathbf{k}_2$, namely, r divided by the projection parallel to $\mathbf{j}_2$ of the input column spacing after oversampling. We thus require a filter breadth factor equal to $½ (1 + pr^2)$.

A final, fourth, pass must be implemented to reduce the number of output lines by the factor c. This vertical resampling of the image is one between strictly parallel grids, and can be effected economically without the need of further 90° rotations. Of course, to diminish the range of vertical frequencies, we must broaden the convolving kernel by the factor c.

(b) $p < 1$.

If the input square contains the critical region (Figure 6) we choose instead $cr^{-1} \cos \theta = p$, i.e., $c = 1 + \tan \theta$. The filter in vertical resampling will be expanded by the factor p^{-1}. Horizontal resampling is used in the first pass to alter the input column spacing. The proper separation, which may be either greater or less than unity, is the reciprocal of

$$\frac{1}{2} p \tan \theta + \frac{1}{2} r^{-1} \sec \theta + \frac{1}{2} p = \frac{1}{2}(1 + p^2r^2) r^{-1} \sec \theta.$$

The filter breadth factor in this first pass is again p^{-1}. Horizontal resampling in the third pass requires expansion of the convolving kernel, this time by the factor $½ (1 + p^2r^2)$. Finally, there is a fourth pass for parallel vertical resampling. The filter breadth factor is c, equal to the ratio of input to output lines for this operation.

It will be noted that, in both (a) and (b), vertical resampling is to an intermediate grid whose rows are parallel to the final output rows, only more numerous in the ratio c. In case (b) we make effective use of the oversampling pass to minimize the number of columns in the image prior to the main horizontal resampling. Likewise, vertical resampling effects some reduction in the number of rows. But there is always a fourth pass to ensure the correct number of final output lines. (The reader, hoping to eliminate this pass, is referred to the italicized sentence on page 502.)

AREAS OF FREQUENCY RETENTION REGIONS

The critical square for parallel resampling before rotation (Figure 4) has area $(r^{-1} \cos \theta)^2$.

In parallel resampling after rotation, the area of the critical square is $|\mathbf{J}_1|\,|\mathbf{J}_2| = r^{-2}$. This is also the area of the region of retained frequencies (Figure 6) when $p < 1$. But when $p > 1$ we must subtract the area of the four triangular regions within the critical square, but external to the input square (Figure 5).

To determine this area, recall that

$$s_1 \cos \phi + s_2 \sin \phi = R$$

is the equation of the straight line, in the Fourier plane, having

(R, ϕ) as polar coordinates of the foot of the perpendicular from the origin. Thus,

$$s_1 \cos\theta + s_2 \sin\theta = \frac{1}{2} r^{-1}$$

$$-s_1 \sin\theta + s_2 \cos\theta = \frac{1}{2} r^{-1}$$

are the equations of the two perpendicular sides of the critical square that meet in the first (i.e., bottom right) quadrant. By finding the coordinates of intersection of these lines with each other, and with the line $s_2 = ½$, we obtain $(p - 1)^2/\sin 2\theta$ for the combined area of the four excluded regions.

When $p < 1$, parallel resampling is clearly preferable after rotation (r^{-2} versus $(r^{-1} \cos\theta)^2$). Figure 7 illustrates the case $r = \sqrt{2}$, $\theta = 45°$. Of the input frequencies (outer square), we can save half (diamond-shaped region) through parallel resampling after rotation, the translates by $\mathcal{L}$ (J_1, J_2) of the retained frequencies completely covering the plane. With parallel resampling before rotation (inner square in Figure 7), only one quarter of the input frequencies is retained.

In the case $p > 1$ we find that there is a transition of preference between the two methods, given by the solution of the equation

$$(r^{-1} \cos\theta)^2 = r^{-2} - (p - 1)^2/\sin 2\theta.$$

This equation reduces to

$$r = \cos\theta + \sin\theta - \sin\theta \sqrt{\sin 2\theta}$$

where the ambiguous sign must be minus, because $p < 1$ implies $r < \cos\theta + \sin\theta$. We shall not add a further illustration for this case.

In summary, for

$$\cos\theta < r < \cos\theta + (1 - \sqrt{\sin 2\theta}) \sin\theta,$$

parallel resampling before rotation is superior, while parallel resampling after rotation is to be preferred for all larger values of r.

Graphs of the extreme members in the inequality above are shown in Figure 8. The range of values of r within which parallel resampling before rotation gives better results is seen to be very narrow. In fact, the maximum vertical separation between the curves, occurring near $\theta = 13.7°$, is less than 0.0762.

As this case will probably seldom occur, it is perhaps not worth the trouble to make special provision for it. Sparse resampling would then always be carried out by parallel resampling after rotation. The analysis will at least make clear why, in that case, the oversampling factors for dense and sparse resampling given above do not come into coincidence in the transition case, $r = \cos\theta$.

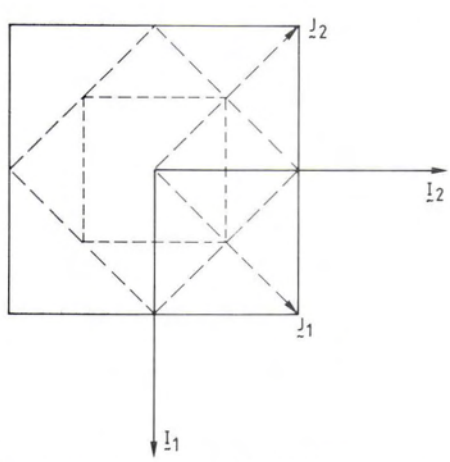

FIG. 7. Comparison of retained frequencies in a case of sparse resampling involving 45° rotation.

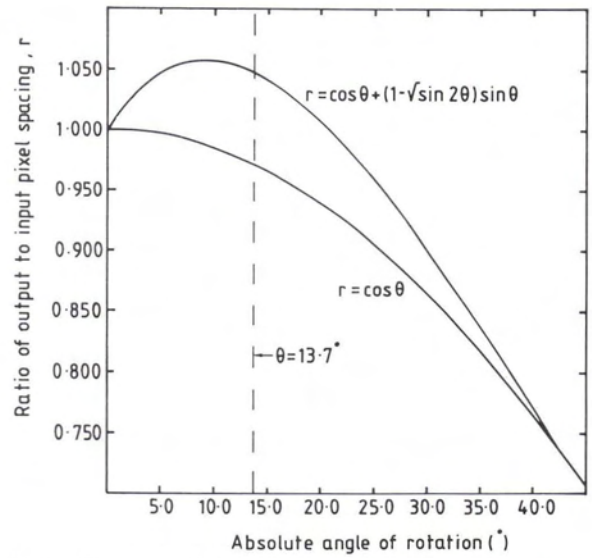

FIG. 8. Variation of ratio of output to input spacing with absolute angle of rotation. The lower curve divides the regions of dense (below) and sparse resampling. Parallel resampling after rotation is superior in the unbounded region above the top curve.

FILTERS FOR SPARSE OUTPUT

We consider the construction of suitable convolving kernels for sparse resampling. These filters resemble the one used for cubic convolution, but are used for averaging $2n$ neighboring input pixels when $n > 2$. In particular, the $2n$-point interpolator has the same support as the function

$$F(x) = \begin{cases} \text{sinc}\,(2x/n) & |x| \leq n \\ 0 & |x| > n \end{cases}$$

and is proportional to it for integral values of x.

The main design problem is to ensure correct interpolation of uniform scenes. Nearest-neighbor and linear interpolation methods obviously perform this function. So too does cubic convolution, though it is less obvious why. We show that a cubic spline collocated with $F(x)$ at integer points has the required property provided that the derivative at $x = \pm n$ is zero.

To this end, let the spine (which we assume to be an even function like F) be given by

$$f_k(x) = g_k(x - k + 1),\ k - 1 \leq x \leq k,$$

where $1 \leq k \leq n$ and

$$g_k(x) = a_k + b_k x + c_k x^2 + d_k x^3,\ 0 \leq x \leq 1.$$

The value interpolated at a pixel position with fractional part x, within a scene of uniform brightness 1, will be

$$\Sigma[g_k(x) + g_k(1 - x)]$$
$$= \Sigma[2a_k + b_k + c_k + d_k + (x^2 - x)(2c_k + 3d_k)].$$

Now, the cubic function $g(x)$ for which

$$g(0) = A,\ g(1) = B,\ g'(0) = C,\ g'(1) = D$$

is

$$g(x) = A + Cx + (-3A + 3B - 2C - D)x^2 + (2A - 2B + C + D)x^3.$$

Writing $a_k = A_k$, $b_k = C_k$, and so on, we obtain, for the interpolated value, the expression

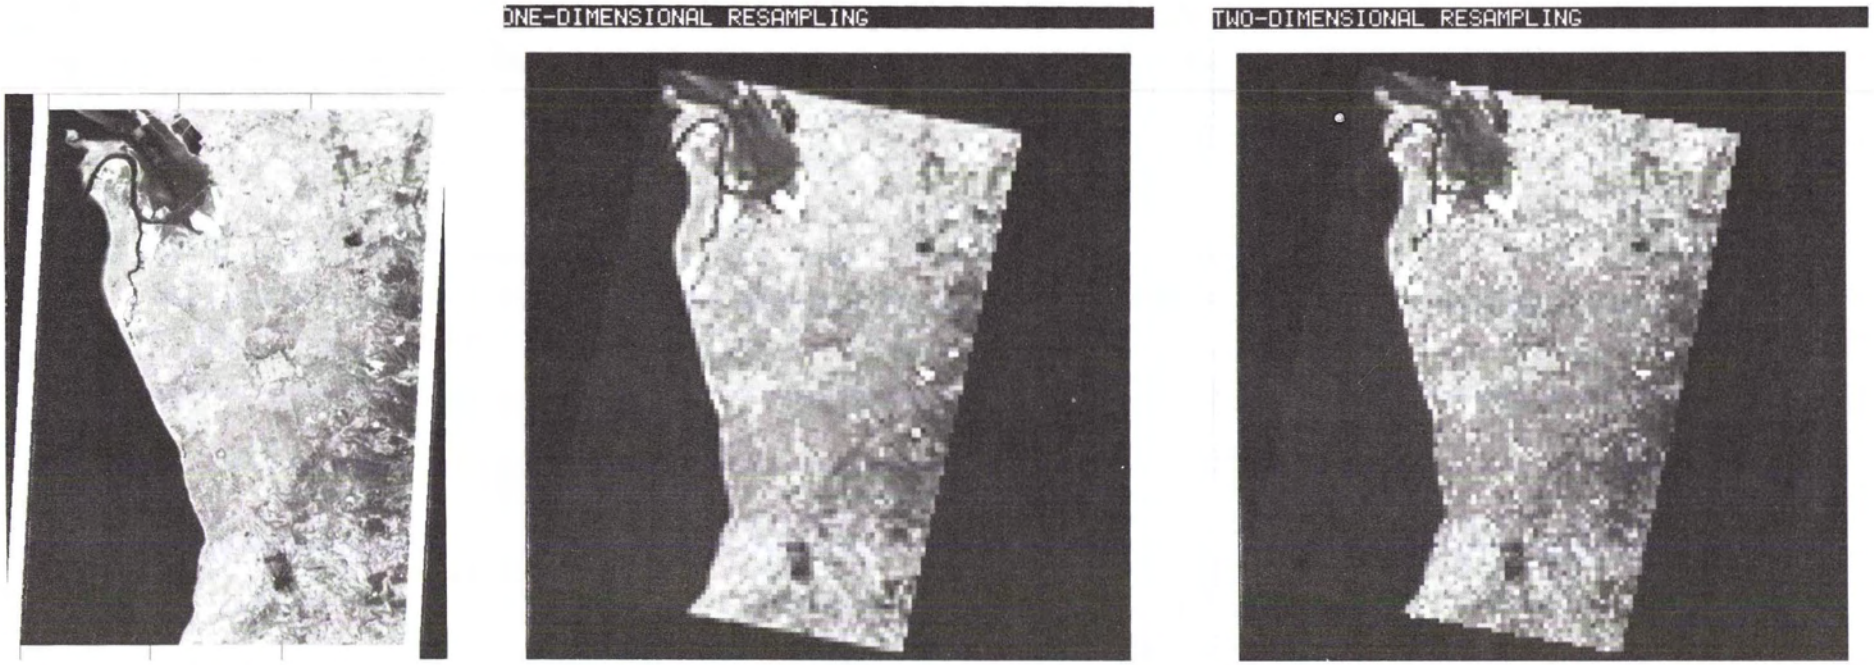

(a) (b) (c)

FIG. 9. Adelaide, South Australia, resampled with reduction. (a) Subscene from MSS band 5 image 1150-000-5. (b) One-dimensional resampling. (c) Two-dimensional resampling (standard 16-point cubic convolution, no smoothing).

$$\Sigma(A_k + B_k) + x(1 - x)\,\Sigma(C_k - D_k),\ 1 \leq k \leq n.$$

The coefficient of the non-constant terms is

$$C_1 + \Sigma(C_{k+1} - D_k) - D_n,\ 1 \leq k \leq n - 1.$$

But $C_1 = 0$ by symmetry and $C_{k+1} = D_k$ as the g_k constitute a spline. Thus, $D_n = 0$ is the sole requirement for a value independent of x. The interpolated value will be the same as for $x = 0$, when it is the sum of values of F at integer points. If we divide the convolving function by this value, the resulting kernel will interpolate uniform scenes correctly. For the cubic spline subroutine employed, see Akima (1970).

ILLUSTRATIONS

Figure 9 illustrates resampling of a Landsat scene (1150-000-5) including Adelaide, South Australia. The original scene (Figure 9(a), 512 by 512 pixels) is of MSS data corrected for Earth rotation and aspect ratio distortions. Upon resampling, it is rotated clockwise by 7° and undergoes reduction by a factor of about 4.6. Figures 9(b) and 9(c) show respective outputs from new and old (16-point kernel cubic convolution) resampling procedures. Quite fine details, such as the railway linking the city with Port Adelaide, to the northwest, are visible in Figure 9(b) but obliterated in Figure 9(c).

CONCLUSIONS

The increasing trend for remote sensing data sets to be resampled to fit geocoded data bases means that there is a need for efficient, high-quality resampling procedures. Furthermore, in order for the volume of stored data to remain as small as possible, it is essential that the final output grid be as sparse as possible within the constraints determined by the data-base users.

When these sparse output grids are required for data that must be rotated through quite large angles, the considerations discussed above will become important. Many data sets acquired from platforms in equatorial orbits, and even from some in polar orbits (such as AVHRR data), can come into this category.

One-dimensional resampling techniques must play an important part when resampling to sparse grids because two-dimensional resampling, if it is not to introduce aliasing, must use very large kernel sizes, with attendant increases in computation overheads.

In this paper we have compared the two major alternative methods available for one-dimensional resampling to sparse grids. For values of r such that

$$r < \cos\theta + (1 - \sqrt{\sin 2\theta})\sin\theta,$$

generalization of Friedmann's procedure, with a preliminary size reduction and smoothing, is preferable. For larger values of r, it is better to retain most spatial frequencies until a dense grid has been established parallel to the final output grid.

The procedures above usually involve one more resampling step than the procedure for resampling to a dense grid. These steps have been outlined, along with suitable broad, one-dimensional filters for resampling. We feel that the implementation of these ideas will ensure optimum quality for imagery that must be rotated and resampled to a sparse grid.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of the Australian Mineral Industries Research Association.

REFERENCES

Akima, H., 1970. A new method of interpolation and smooth curve fitting based on local procedures. *Journal of the Association for Computing Machinery*, 17, 589–602.

Friedmann, D., 1981a. Operational resampling for correcting images to a geocoded format. *Proceedings of the 15th International Congress on Remote Sensing of Environment*, Vol. I, 1195–1212.

Friedmann, D., 1981a. Operational resampling for correcting images to a geocoded format. *Proceedings of the 15th International Congress on Remote Sensing of Environment*, Vol. I, 1195–1212.

———, 1981b. Two-dimensional resampling of line scan imagery by one-dimensional processing. *Photogrammetric Engineering and Remote Sensing*, 47, 1459–1467.

(Received 29 November 1984; revised and accepted 19 January 1987)

APPENDIX

Let $\mathscr{L}$ be a point-lattice and R a convex region symmetric about the origin. Then R is disjoint from its translates by vectors of $\mathscr{L}$ if and only if R contains no point, except the origin, of the lattice $\frac{1}{2}\mathscr{L}$.

The proof in both directions is by contradiction. First, if R contains a non-zero vector $\frac{1}{2}\ell$ (ℓ in $\mathscr{L}$) then, by centro-symmetry, it also contains $-\frac{1}{2}\ell$. Thus R, $R + \ell$ both contain $\frac{1}{2}\ell$.

Conversely, if the translate $R + \ell$ contains a vector, $\mathbf{v}$, then R contains $\mathbf{v} - \ell$, hence also $\ell - \mathbf{v}$, again by symmetry. If R itself also contains $\mathbf{v}$ then, by convexity, it must contain the midpoint

$$\frac{1}{2}\mathbf{v} + \frac{1}{2}(\ell - \mathbf{v}) = \frac{1}{2}\ell.$$