Publications of the Astronomical Society of the Pacific
104: 285–289, 1992 April

The Nyquist Criterion in CCD Photometry for Surface Brightness

ROBERT L. WILDEY
Department of Physics and Astronomy, Northern Arizona University, Box 6010, Flagstaff, Arizona 86011
Received 1991 September 16; accepted 1992 January 8

ABSTRACT. Astronomers routinely violate the directive to sample surface brightness with at least twice the frequency of the highest spatial frequency of the Fourier transform of the continuous image, when doing direct CCD imaging. It is reasonably speculated that this practice is rationalized on the basis that the CCD does not actually *sample* the surface brightness at periodic intervals, but instead integrates the surface brightness over contiguous regions (the CCD pixels). It is herein derived that this mode of sampling changes the form of aliasing error, but the aliasing error is nevertheless present when undersampling occurs. The very nature of the error betrays the possibility of detecting its presence, *a priori.* It would be of value to develop an active optical apodizer to accommodate a given CCD, in terms of the Nyquist criterion, without the need to abandon either the full light-gathering power of the telescope or the plate scale at the chosen observing station.

1. INTRODUCTION

1.1 Observational Background

The use of digital imaging techniques in the photometry for the surface brightness of extended objects introduces a new source of systematic uncertainty into the corresponding error budget: aliasing due to undersampling. The error is pernicious in its lack of contribution to a visibly obvious increase in image noise level, as it takes the form of a failure of simple interpolation techniques within the digitized image and the aliasing of fine structure within the continuous image into systematic patterns of a broader nature. Aliasing error is something like a beat phenomenon between the sampling frequency and higher frequencies in the information spectrum, and as such is a form of systematic error at the grid points themselves of the digitized image. Astronomers are inclined to feel vindicated in ignoring the maxim formulated by Nyquist, that the light signal must be sampled at least twice as frequently, in traversal of the continuous image, as the cresting in the finest sinusoid in its Fourier decomposition, if this aliasing error is to be avoided. One principle reason for this omission is that adherence to proper sampling theory often requires a very high focal ratio, so high that the exposure time is increased to impractical duration. Of course this consideration does not cause the error to vanish, and generally impacts only deep sky photometry, interestingly increasing exposure times to values once tolerated in the era of the photographic emulsion. In lunar and planetary astronomy the use of diaphragmed apertures to achieve focal ratios in Nyquist compliance is not only feasible but has sometimes actually been practiced (Wildey 1977, 1978). A full-moon exposure with a CCD is still under 1 s at a focal ratio of 80 and a spectral bandwidth of 200 Å, at wavelength 5000 Å. An exposure for Jupiter is nominally only 60% longer. Light is a terrible thing to waste, but collecting extra photons in the form of partly scrambled signal may be a bad bargain.

A second justification for ignoring the Nyquist criterion may be offered in terms of the dominance of atmospheric seeing over the fundamental diffraction of the telescope's aperture. But if compliance with the Nyquist criterion is brought about through the use of an off-axis aperture in a diaphragm at the entrance pupil, thus degrading the diffraction limit, the resolution will be sufficiently low that seeing will not be dominant most of the time. In addition, even if the telescope's resolving power and light-gathering power is preserved in bringing about Nyquist compliance, through the use of a repeater lens to increase effective focal length, it must be remembered that seeing, unlike diffraction, attenuates and shifts the complex amplitudes at higher spatial frequencies rather than omitting them. Image restoration techniques are feasible where seeing is concerned, because the information is there. But no form of image processing can compensate for undersampling.

Perhaps the principle reason for the failure to comply with the recommendation emerging from Nyquist's theorem is the *accurate* perception of the invalidity of its application, that the premise upon which it is based and even the language of its development and conclusion are incompatible with the conditions of CCD photometry. For in fact the CCD does not *sample* periodically the continuous image, but instead integrates over each member in a uniform series of contiguous intervals. May not such breadth of integration cancel out the aliasing errors that would be implied in a formally correct application of Nyquist's theorem? The principle contribution of the present paper is to derive a properly modified form of Nyquist's theorem that is rigorously applicable to CCD photometry. In doing this, it has been discovered that such aliasing errors are not canceled out by spatial integration, though they are significantly changed in form, and that the emergent criterion remains unmodified in expression.

1.2 Theoretical Background

Nyquist's theorem is elaborated in many different textbooks, e.g., Downing (1961), in relation to signal processing theory. The theorem transcends dimensionality. It may involve an independent variable to be thought of as time, and hence be one dimensional (Downing 1961), or two independent variables that are coordinates in an optical image, and hence be two dimensional (Wildey 1967 ab, 1973). In continuous spectrometric imaging the problem is three dimensional (Young 1974) (the concept of the image "cube"). But the logic that underlies the Nyquist criterion is independent of dimension, so that if it is stated in one dimension its nature in n dimensions is obvious. Formally stated, it can be proved that if a continuous function is known only in the form of a tabulation at evenly spaced discrete values of its independent variable, there is a simple condition which, if satisfied, guarantees that there is zero loss of information in going from the continuous function to its tabulation. That condition is that the sampling frequency of the table (the number of tabulations of the function per unit interval of the independent variable) must equal or exceed *twice* the highest frequency at which information exists in the Fourier transform of the continuous function. If there is no information loss, the inescapable conclusion is that the exact value of the function at an arbitrary value of its independent variable can be obtained from the table. This may be thought of as a form of interpolation, however, it is ultimate in high order, as the entire table is involved for the inference of functional value at one point in the continuum of values.

The mathematical equivalent of a table of values of a function is the product of the function with a sum of progressively shifted Dirac delta functions. If the function is the distribution of flux in the focal plane of a telescope, then the corresponding table is the array of reduced surface brightnesses derived from the time-integrated signals recorded by an array of detectors whose individual point-spread functions (spatial responsivity functions) are infinitesimally narrow compared with their spacings. Because a CCD is more nearly an array of flat-topped vertically sided functions that abut one another (an array of contiguous "boxcar" functions), there is no justification for applying the Nyquist criterion in the above form in order to determine whether one's CCD images possess aliasing errors, or may not be reliably interpolated with a precision that increases as higher-order interpolation formulas are used. A visceral sentiment that a sampling frequency selected according to the Nyquist criterion must surely be correct within a factor of 2, even if not erroneous, is not of adequate precision. Nor is there more than vague satisfaction in the notion that the Airy disk should just nominally cover an individual pixel of the CCD. The subject deserves deeper and more rigorous inquiry.

2. THE PRESENT THEORY

2.1 Modeling the Signal in Image Space

In the following treatment, we will keep arguments transparent to the need for normalization constants in taking Fourier transforms for the sake of simplicity. We will also avoid the unwieldy notation required for two-dimensional dependence, inasmuch as the extension of applicability from one dimension to two is completely straightforward, so long as one is content with Cartesian coordinates.

Let a one-dimensional continuous image formed by a hypothetical telescope be represented by the function $f(x)$. The function thus defined is not the original scene, then, as its Fourier spatial spectrum has already been apodized in the form of the Fraunhofer diffraction of the telescope's entrance aperture. Next, let the *continuous* spatial signal stream corresponding to the dwell of a general transducing device be $s(x)$. It is mathematically useful to the present argument to define such a function even though it corresponds in any practical sense to image acquisition by area scanning rather than the simultaneous operation of detectors in an array. The relation between $s(x)$ and $f(x)$ will then be given by

$$s(x)=\int_{-\infty}^{\infty} r(x')f(x-x')dx',$$

$$s(x)=r(x)*f(x), \tag{1}$$

where $r(x)$ is the point-spread function of the detector, normalized to integral unity. We have defined the integral operation of convolution above, in terms of the asterisk, for subsequent brevity of notation.

Insofar as the electron wells of a CCD are operationally identical, the digital image that it acquires can be related to Eq. (1) by stating that such image is equivalent to $s(x)$ being given only at discrete positions, $x=n\Delta$, where n is a real integer and Δ is pixel separation. Then $r(x)$ is now the nominal point-spread function of a pixel (we are not yet enforcing that Δ is also the pixel size). A functional representation of the digitized image, which we will call $h(x)$, can be easily contrived,

$$h(x)=\sum_{n=-\infty}^{\infty} \delta(x-n\Delta)s(x),$$

$$h(x)=\sum_{n=-\infty}^{\infty} \delta(x-n\Delta)[r(x)*f(x)]. \tag{2}$$

For the moment, we will ignore the implication that the infinite limits on the summation suggest an infinitely large format for the digital image (see the Appendix). Thus $h(x)$ is the digital version of $s(x)$. The use of the Dirac delta function, inasmuch as the final pixels are thereby infinitely narrow, expresses the result in terms of pure information, as ready for processing on a digital computer, whereas an optical rendering would be represented by an additional convolution of $h(x)$ with such as a boxcar function. No conceivably useful form of processing for present considerations demands such additional operation. We reemphasize that CCD operation is already incorporated into $s(x)$, and that this use of the Dirac delta function signifies only the discreteness of a tabulation and *not* an idealized point-spread function.

2.2 Analysis in the Fourier Domain

We now wish to take the general complex Fourier transform of both sides of Eq. (2). With ω as frequency in the sense of radians per unit interval of x, we follow the convention of using the capital letter corresponding to the lower-case function designation to signify its Fourier transform. The Fourier transform of the sum is the sum of individual Fourier transforms. Thus

$$H(\omega)=\sum_{n=-\infty}^{\infty} \int_{-\infty}^{\infty} \{\delta(x-n\Delta)[r(x)*f(x)]\}e^{i\omega x}\,dx. \tag{3}$$

An important theorem of Fourier analysis is now useful (e.g., Bracewell 1986):

> The Fourier transform of a product is the convolution of the individual Fourier transforms. The Fourier transform of a convolution is the product of the individual Fourier transforms.

Applying this theorem twice, once in each form, we have

$$H(\omega)=\sum_{n=-\infty}^{\infty} \{e^{i\omega n\Delta}*[R(\omega)F(\omega)]\},$$

$$H(\omega)=\sum_{n=-\infty}^{\infty}\int_{-\infty}^{\infty}[R(\omega')F(\omega')]e^{in\Delta(\omega-\omega')}\,d\omega'. \tag{4}$$

[That $\exp(i\omega n\Delta)$ is the Fourier transform of $\delta(x-n\Delta)$ is a matter of straightforward integration.] In the explicit form of the convolutions of Eq. (4), we see that the summation and integration can be commuted and terms not containing n can be factored out of the sum. This yields

$$H(\omega)=\int_{-\infty}^{\infty}R(\omega')F(\omega')\sum_{n=-\infty}^{\infty}e^{in\Delta(\omega-\omega')}\,d\omega'. \tag{5}$$

The series sum on the extreme right-hand side of the integrand of Eq. (5) can be regarded as the Fourier series representation of some function we will label g, whose argument is $(\omega-\omega')$. The peculiarities of this Fourier series are that its coefficients are all the same and that its periodicity in $(\omega-\omega')$ is $2\pi/\Delta$. Given the general rule for determining Fourier coefficients on a periodic interval, as applied to this situation,

$$C_n\propto\int_{-\pi/\Delta}^{+\pi/\Delta}g(y)e^{iny\Delta}\,dy=\text{const independent of } n, \tag{6}$$

the nature of $g(y)$ emerges. If C_n is independent of the value of n, it is thereby independent of the weighting given to $g(y)$ by the weighting function $\exp(iny\Delta)$ at all values of y for which the value of n influences the weighting. That is, the weighting of $g(y)$ is immaterial at all values of y except $y=0$. This is only possible if $g(y)$ is zero everywhere except at $y=0$, where it is sufficiently infinite to produce a finite contribution to an integral. In other words,

$$g(y)=\delta(y), \tag{7}$$

the Dirac delta function. As a Fourier series representation of $\delta(y)$, however, the periodicity on a repeating interval of $2\pi/\Delta$ must follow. [If the limit is taken in Eq. (6) as the periodic interval becomes infinite, the well-known Fourier integral representation of the Dirac delta function emerges (Mathews and Walker 1965).]

As a result of immediately preceding considerations, we conclude that

$$\sum_{n=-\infty}^{\infty}e^{iny\Delta}=\sum_{m=-\infty}^{\infty}\delta\left(y-\frac{2\pi m}{\Delta}\right), \tag{8}$$

in which the summation on the right-hand side for terms other than $m=0$ provides the requisite periodicity. Relabeling the independent variable from y to $(\omega-\omega')$, and inserting from Eq. (8) into Eq. (5), we have

$$H(\omega)=\int_{-\infty}^{\infty}R(\omega')F(\omega')\sum_{m=-\infty}^{\infty}\delta\left(\omega-\omega'-\frac{2\pi m}{\Delta}\right)d\omega'. \tag{9}$$

In Eq. (9) we may take $R(\omega')F(\omega')$ under the summation sign and then interchange integration and summation to obtain some directly do-able integrals:

$$H(\omega)=\sum_{m=-\infty}^{\infty}\int_{-\infty}^{\infty}R(\omega')F(\omega')\delta\left(\omega-\frac{2\pi m}{\Delta}-\omega'\right)d\omega',$$

$$H(\omega)=\sum_{m=-\infty}^{\infty}R\left(\omega-\frac{2\pi m}{\Delta}\right)F\left(\omega-\frac{2\pi m}{\Delta}\right). \tag{10}$$

2.3 Aliasing as Harmonic Overlap

Thus we see that the effect produced by discretely sampling the continuous function $s(x)$ [to produce $h(x)$] is that it accomplishes a result whose Fourier transform is not only the Fourier transform of $s(x)$, but the sum of that transform with an infinite number of versions of itself whose origins are shifted from zero to the sampling frequency and all of its harmonics. Given that fact, it must now follow that if, and only if, the shifted versions of $S(\omega)=R(\omega)F(\omega)$ do not overlap, one can formulate a method of exactly "interpolating" the value of $s(x)$ by taking the Fourier transform of $h(x)$, multiplying the result by a product of Heaviside step functions (to form a boxcar function) in the form of

$$H_e\left(\frac{\pi}{\Delta}-\omega\right)H_e\left(\frac{\pi}{\Delta}+\omega\right),$$

thus isolating the Fourier transform of $s(x)$, then taking the inverse Fourier transform of that result and evaluating it at the desired value of x. The foregoing condition is identifiable with the Nyquist criterion, but applied to the function $[r(x)*f(x)]$, not $f(x)$, as would be the case in a straightforward application of Nyquist's theorem as though a CCD were an array of Dirac delta functions. The characteristics of the CCD, as the specifier of $r(x)$ and $R(\omega)$, have not yet been introduced into the argument. That is now the appropriate point to raise.

2.4 The CCD as Periodic Convolver

For a CCD image, at least the vast majority of nonexperimental varieties in which the dimension of the photon-capture well greatly exceeds the wavelength of light, there is a simple relation between the point-spread function and the sampling interval. The point-spread function is a boxcar function whose width is the sampling interval. Consequently, its Fourier transform is simply

$$R(\omega)=\int_{-\Delta/2}^{\Delta/2}e^{i\omega x}\,dx,$$

$$R(\omega)=\frac{\sin(\omega\Delta/2)}{(\omega\Delta/2)}=\mathrm{sinc}(\omega\Delta/2). \tag{11}$$

The zeros of $R(\omega)$ occur at $\omega_n=2\pi n/\Delta$ for all *nonzero* integers n. As we address the issue of overlapping harmonics, we note that if this $R(\omega)$ from Eq. (11) is substituted into Eq. (10), with no constraints on $F(\omega)$, the center frequency (zero) of the dc component of the right-hand side of Eq. (10) is superposed with the nth zero of the nth harmonic, and all harmonics are present. We further note, more specifically, that at one half the sampling frequency of the CCD, one sample per (2Δ) or π/Δ radians per unit of x, which is where $S(\omega)$ must reach and stay at zero if there is to be no overlap from the first harmonic, and hence aliasing errors are to be avoided, the contribution is the same from $R(\omega)$ and $R[\omega-(2\pi/\Delta)]$ at a value of $200/\pi=64\%$ of maximum. This state of affairs is indicative of aliasing. It also precludes the simple interpolation of $s(x)$, hence the reasonable inference of $f(x)$ to a definable precision and resolution. Such an interpolation produces a result that cannot be precisely characterized, given a free choice of $f(x)$, but its difference from the correct result may be crudely approximated as a random variable whose mean stationary amplitude is equal to the rms difference between adjacent pixels in the image and whose correlation length is about one-fourth, more or less, of pixel separation.

Therefore, the assurance of the necessary purity of the signal, $s(x)$, is not to be gained from the behavior of the CCD. It must be supplied by insuring that $F(\omega)$ cuts off absolutely at $\omega=\pi/\Delta$, thus guaranteeing the nonoverlap of the harmonics on the right-hand side of Eq. (10), irrespective of the behavior of $R(\omega)$. This condition can only be reasonably imposed by assuring that the optical transfer function of the telescope cuts off at $\omega=\omega_c=\pi/\Delta$. This becomes a constraint on the focal ratio of the telescope imposed by the choice of CCD.

2.5 Optical Transfer Function of the Telescope

The optical transfer function of the telescope is a scaled autoconvolution of the aperture as a cylindrical boxcar function, which cuts off at a spatial frequency, in cycles per radian of scene/image subtension at the effective lens/mirror nodal point, equal to the objective diameter divided by the wavelength of light (e.g., Smith 1963). Let the effective focal length be f, the objective lens/mirror diameter be D, the focal ratio be F, and the wavelength of light be λ. The focal plane scaling factor is the reciprocal of the focal length. In terms of radians rather than cycles, the telescope imposes the condition

$$\omega_c=\frac{2\pi D}{\lambda f}=\frac{2\pi}{\lambda F}. \tag{12}$$

But we have made a case for setting this cutoff at π/Δ. Consequently, we require

$$F=2\Delta/\lambda. \tag{13}$$

2.6 Theoretical Summary

Equation (13) expresses the same result that would have been achieved had we required the optical transfer function to cut off at one-half the sampling frequency of the CCD, i.e., as though a CCD were an array of Dirac delta functions and the Nyquist criterion could be straightforwardly applied. However, such an approach was not justifiable and the argument presented has been quite different.

3. DISCUSSION

In Table 1 we have randomly chosen a collection of 7 CCDs currently accessible to the astronomical community. The set of corresponding focal ratios has been based on a wavelength of 5000 Å. Only the smallest pixel size available (which tends to indicate a more temperamental CCD) enables compliance with the Nyquist criterion at a typical telescope's Coudé focus, operated at full aperture. The Tektronix 512×512 CCD with 27 μm pixel size would be the more likely candidate for deep-sky imaging photometry at a prime or Newtonian focus, just as in former times signal-to-noise considerations would have dictated a coarse-grain emulsion over a fine-grain emulsion for such work. With a likely focal ratio of from 3 to 5, the result would be undersampling by a factor of from 22 to 36. To take the Hale telescope as an example, the sampling frequency would be about one cycle per second-of-arc of the celestial sphere, while the cutoff of the optical transfer function is about 30 cycles per arcsec. It can be calculated

TABLE 1
Commonly Available CCDs

Manufacturer	Format	Δ (μm)	F
Videk (Kodak)	1320×1035	6.8	27
Thompson, T. T.	1024×1024	12	48
Ford, RCA	512×512	20	80
GEC8602	576×385	22	88
Tektronix	512×512	27	108
TI/VPIM	1024×1024	18	72
TI4849	582×390	22	88

that a strong periodic pattern in object space at a characteristic spatial frequency of 1.02 cycles per arcsec, if the foregoing sampling frequency be taken exactly, will be imaged, in part, as a periodic pattern of almost a minute of arc between the brightness crests, with scarcely a diminution of 50% over the original amplitude of the unaliased frequency. Strictly speaking, the result is not a photometric error, but a morphometric error, but that nuance provides little consolation. In the example just cited, should the telescope be diaphragmed to 15 cm to conform to sampling theory, the exposure time will increase by a factor of 900. If light is not wasted, by use of 4× repeating optics at the Coudé focus, the long exposure will still be required for a single frame, in addition to which that frame will be but 19 arcsec across. What is needed is an active optical device that will cleanly reapodize the image, without diminished entrance pupil, somewhere ahead of the CCD.

The foregoing discussion has been thus far limited to the aliasing error that can appear even at the precise grid points of the digital image. There is also the matter of the precision of interpolation. For unbounded variation of the sampled function in a strictly correct context for application of Nyquist's original theorem, the size of error in a failure of linear interpolation caused by undersampling is itself unbounded. But for the nature of the sampling process that is characteristic of a CCD, the interpolation error is clearly relatively subdued and more nearly like the random variable discussed earlier in terms of its approximate boundedness. It remains a mode of characterization of the error associated with undersampling.

Compliance with the Nyquist criterion, in terms of interpolability, serves concrete practical considerations in certain programs. For example, the lunar photometry program proposed by Kieffer and Wildey (1985) of modeling reflection of sunlight from the lunar surface so that the Moon can be used to radiometrically and stereometrically calibrate spacecraft and satellite imagery *after launch* requires intercomparison of an extremely large number of photometric images of the Moon over a range in libration and phase requiring years to accumulate. It is a logistical necessity that every image be interpolated and accumulated on a master selenographic grid for proper modeling.

4. CONCLUSION

The principal conclusion of this investigation is that compliance with the Nyquist criterion is worth pursuing any time that it is remotely feasible to do so. The fact that a CCD samples surface brightness by integrating it over contiguous intervals does not prevent the presence of the aliasing errors that would be expected from an array of infinitesimal spatial samplings, should that sampling be inadequately frequent according to Nyquist's theorem. When aliasing errors are present they do not signify it in the manner of the "snow" of random errors. The pursuit of research in the development of active optical devices to smoothly apodize the transfer functions of telescopes to accommodate the sampling intervals of digital imaging devices, without sacrificing light-gathering power or increasing scale, is to be encouraged.

APPENDIX

In the foregoing derivation, the infinite limits of the sum in Eq. (2) should be actually truncated in order to provide a finite-image format, contrary to the development shown. Infinite limits were accepted ostensibly because of the irrelevance of predictions outside the central region of actual interest. We now deal with the issue in less cavalier fashion. The effect of a truncation at $n=\pm N$ is to replace the $H(\omega)$ of Eq. (10) by its convolution with $[\sin(\omega N\Delta/2)/(\omega N\Delta/2)]$. That this is true can be seen as follows. Firstly, the truncation can be realized by retaining the infinite limits on the sum and multiplying the Dirac delta function, $\delta(x-n\Delta)$, by an even boxcar function of x whose edges are at $x=\pm N\Delta$. Where one formerly had the Fourier transform of $\delta(x-n\Delta)$ involved in the argument, one will now have the Fourier transform of this product, which is the convolution of the Fourier transform of the boxcar function with the Fourier transform of the Dirac delta function. The Fourier transform of the boxcar is the foregoing sinc function of $(\omega N\Delta/2)$. It is not a function of n and can be taken to the left-hand side of the summation sign in Eq. (4), using also the fact that all convolutions are simply commutative and associative. What is left on the right-hand side of this final convolution is the expression ultimately devolved to the $H(\omega)$ of Eq. (10). Q.E.D.

The convolution of the Dirac delta function with any function is simply that function returned unmodified. The sinc function of $(\omega N\Delta/2)$ is negligibly different from the Dirac delta function (spike-spent in an interval small compared to a pixel) so long as $N\gg 1$. That deserves slightly more elaboration. Certainly this sinc function is thoroughly born and dead over a range of ω that is small compared to the range required for significant change of $R[\omega-(2\pi m/\Delta)]$. It would appear that that would have to be true for $F[\omega-(2\pi m/\Delta)]$ as well, if the trivialization of this issue is to remain justified. By nature of the Fourier transform pair relation, abruptness in F transforms to breadth in f. If we arbitrarily say that the effective range of operation of this sinc function is within ten sign changes of either side of its maximum, that is in correspondence to $\Delta\omega=40\pi/N\Delta$. An abrupt change in such a range implies significant amplitude at the corresponding "frequency" (the frequency corresponding to ω is actually x) of $2\pi/(40\pi/N\Delta)=N\Delta/20$. But presumably the image extent is $(\Delta x)=2N\Delta$. Have we just shown that the incorruptible image we wish to process must occupy no more than 1/40th of the region of digitization? Although the answer develops from discrete Fourier transform theory (Pratt 1978) that is beyond the scope of the present paper, the answer is yes only if one insists that the rest of the region be null. For a range of digitization in correspondence with image size, at $\Delta x=2N\Delta$, the rest of the space of x will not be unoccupied but contain periodic reproductions of $f(x)$. But there would not be file coverage in that greater region, and that is consistent with the fact that it contains redundant information.

REFERENCES

Bracewell, R. N. 1986, The Fourier Transform and its Application, 2nd ed., revised (New York, McGraw-Hill), p. 108

Downing, J. J. 1961, in Aerospace Telemetry, ed. H. L. Stiltz (Englewood Cliffs, NJ, Prentice-Hall), p. 83

Kieffer, H. H., and Wildey, R. L., 1985, Photogramm. Eng. Remote Sensing, 51, 1391

Mathews, J., and Walker, R. L. 1965, Mathematical Methods of Physics (New York, Benjamin), p. 97

Pratt, W. K. 1978, Digital Image Processing (New York, Wiley), pp. 121 and 279

Smith, F. D. 1963, Appl. Opt., 2, 335

Wildey, R. L. 1967a, PASP, 79, 220

Wildey, R. L. 1967b, AJ, 72, 884

Wildey, R. L. 1973, ApJ, 186, L47

Wildey, R. L. 1977, Moon, 16, 21

Wildey, R. L. 1978, Science, 200, 1265

Young, A. T. 1974, Methods of Experimental Physics, Part 12A, ed. N. Carleton (New York, Academic), p. 184