Reprinted from SPIE Vol. 264—Applications of Digital Image Processing to Astronomy
© 1980 by the Society of Photo-Optical Instrumentation Engineers, Box 10, Bellingham, WA 98227 USA

Artifacts in digital images

Jean J. Lorre, Alan R. Gillespie
Image Processing Laboratory, Jet Propulsion Laboratory
Mail Stop 168-522, 4800 Oak Grove Drive, Pasadena, California 91103

Abstract

Three kinds of artifacts unique to digital images are illustrated: aliasing caused by undersampling; interference phenomena caused by improper display of images; and harmonic overtones caused by quantization of amplitudes. Especial attention is given to undersampling when the sample size and interval is the same. This situation is important because it is typical of solid-state cameras.

Introduction

Digital electronic cameras are being used increasingly in a wide variety of fields ranging from medicine to remote sensing of the earth and other planets. Images made by these cameras may contain artifacts, of which the most important result from undersampling.

Discrete sensors must gather data in accordance with the Whittaker-Shannon Sampling Theorem if a scene is to be accurately reconstructed from its image. For vidicon cameras and line scanners, it is often possible to choose point spread function shapes and sample spacings to minimize violation of the sampling theorem, but for the important class of solid-state cameras including line array and two-dimensional array sensors, device fabrication precludes this possibility. For array cameras, the sample dimension cannot exceed the sample spacing, resulting in undersampling. However, solid-state cameras are sufficiently attractive on other grounds that CCD arrays are planned for use on NASA's project Galileo and the Space Telescope, as well as for countless ground-based imaging systems. Therefore, it is important that the consequences of undersampling be widely understood, not only from an analytic basis, but from an image interpreter's viewpoint as well.

Even if the sampling theorem is obeyed, it does not guarantee that a pictorial representation of the sampled data will visually resemble the original band-limited scene - only that the scene may be reconstructed from the image. In practice, the pictorial representation generally does resemble the scene so often that analysts may neglect to consider the exceptions. However, overlooked artifacts can be unintentionally exaggerated during image enhancement, leading to misinterpretation.

The artifacts discussed in this paper are of three kinds: aliasing in CCD images; interference of the sampling function with harmonic constituents of a scene to produce visual artifacts in pictorial representations of the discrete image; and aliasing of harmonic overtones caused by quantization of discretized data. These artifacts can interfere with image interpretation, especially if the scene is strongly ordered or periodic, and can also reduce the effectiveness of enhancements such as edge-recognition or high-pass filtering.

Nevertheless, remedies exist for two of these problems. Undersampling can be made adequate by band-pass filtering the scene before sampling, using a camera lens with a broad PSF. Interference artifacts can be reduced by image resampling, but quantization artifacts cannot be removed.

Sampling and artifacts in CCD images

The sampling theorem[1,2] states that a band-limited function whose Fourier transform is zero at frequencies greater than some frequency S_c is fully specified by values spaced at equal intervals not exceeding $(2\,S_c)^{-1}$, with the exception of harmonic terms with zeroes at the sampling points. Real scenes have extended spectra, but may be band-limited by filtering in the atmosphere, the optics, or the sensor itself. CCD's and other solid-state sensors cannot band-limit a signal properly because the sample dimension is at most equal to the sample spacing. Thus when the point-spread function (PSF) of the combined atmosphere and optics is similar in size to individual receptors of the CCD, care must be taken in interpreting resulting images.

The process of sampling is briefly summarized in Figure 1. A continuous scene might have an extensive spectrum F(s) as shown in Figure 1a. A CCD or other solid-state camera with a square receptor (pixel) of width d will have a transfer function specified by $\sin(\pi ds)/\pi ds$, or sinc(ds), where s is the spatial frequency. The energy spectrum available to the CCD for measurement is $F(s)\cdot\mathrm{sinc}(ds)$, a sinc function with attenuated side lobes shown in Figure 1b. When the scene is sampled at an interval p, the continuous spectrum is converted to a repetitive and overlapping array of discrete spectra with a periodicity of p^{-1} in the frequency domain. For CCD cameras, $d \leq p$ and the fundamental spectrum is significantly contaminated by

energy from adjacent spectra (Figure 1c). For example, at .25 cycles/pixel, half the energy may be aliased. If $d \geq 2p$, the only contamination would be from side lobes of neighboring spectra. Figure 1d illustrates such a spectrum, specified by:

$$I(s) = \sum_{n=-\infty}^{\infty} F(s + \frac{n}{p}) \cdot \mathrm{sinc}[d(s + \frac{n}{p})] \qquad (1)$$

Even so, as evaluation of the above series shows, roughly one quarter (the exact amount depends on F(s)) of the energy at .25 cycles/pixel is aliased.

In an image, aliased energy is manifest as fine detail which is distorted in two ways: small features may appear too large and their locations may be shifted by up to $d/\sqrt{2}$. As Root[3] has illustrated, for most scenes the distortions caused by CCD-style undersampling are not severe and indeed may improve the appearance of the image by increasing the "sharpness" and the amount of detail visible. This is because of the large amount of energy aliased when $d = p$. Most such images can be interpreted correctly, provided the scene is familiar to the analyst. Simple geometric targets may be grossly distorted, however, and scenes containing repetitive geologic structure or periodic atmospheric cloud features may likewise be misrepresented. It should be emphasized that low-pass filtering can do little to reduce these artifacts if any scene detail is to be preserved at all, because fundamental and aliased energy cannot be separated and are both reduced alike. Conversely, high-pass filtering emphasizes the most strongly aliased part of the spectrum, exacerbating interpretive difficulties.

Figures 2 and 3 illustrate the visual impact of CCD-type undersampling on a highly structured target and on an arbitrary real scene. In both figures the original "scene" is at the upper left. In Figure 2 this is a radial sinusoidal spoke pattern and in Figure 3 it is a part of a Viking image of the surface of Mars. In both figures the picture at the upper right represents the scene band-limited such that sampling at the CCD pixel spacing p produces isolated periodic image spectra, as in Figure 1d. For comparison, at center left are the images resulting from sampling the original "scene" with $d = p$. These simulate images from a CCD camera. At bottom are the log amplitude of the Fourier transforms of the top images. Finally, at right center are the CCD images after post-sampling low-pass filtering, illustrating the infeasibility of this approach as a remedy for undersampling.

Comparison of Figures 2 and 3 shows that artifacts are clearly evident in the geometric target, but are less so in the picture of Mars. The essential point is that in the presence of quasi-periodic structure, image information can be so fundamentally altered by CCD-type undersampling and aliasing as to preclude correct visual interpretation. In the properly sampled image (upper right) information in the overlapped spectral region is removed, protecting the analyst from false interpretations by rendering none possible. Further examples of sampling artifacts with various d and p may be found in Root.[3]

NASA is currently considering plans to launch an earth-orbiting satellite which would, as one possibility, carry a solid-state six-channel down-looking camera along with broad-band fore-and aft-looking cameras to provide stereoscopic images. These images would necessarily be aliased. Elevation data could be derived from these images using cross-correlation on a point-by-point basis to measure parallax. This must always be performed, whether by an observer using a stereoscope with a parallax bar or by digital computer. Aliasing probably will not effect this process except to slightly increase the r.m.s. error of the calculated parallax image, because most of the detail upon which registration is based is several pixels in size. The effect of aliasing is similar to the effect of reducing the signal to noise ratio, and can be controlled by varying the size of the correlation window.

Undersampling is not a problem of all electronic cameras. In vidicons or line scanners the sample dimension and the sweep rate and dwell time (which are factors of p^{-1}) can be chosen to satisfy the sampling theorem. In some instruments, notably the Viking Lander cameras and the Landsat MultiSpectral Scanner (MSS), aliasing is directionally dependent. The MSS sensor consists of a column array of six photoreceptors across which the image is swept by an oscillating mirror. Dwell time was selected so that in the scan direction the scene is correctly sampled, while in the orthogonal direction $d < p$ and the scene is undersampled. For the wide variety of linear array cameras now under development, such a hybrid sampling scheme is also a possibility. For these sensors, a two-dimensional image is constructed by utilizing a scan mirror or aircraft motion over the scene to sweep information past the sensors. Sampling can thus be controlled by varying scan-mirror or aircraft speed.

However, once a scene has been undersampled there is no simple procedure for resampling or filtering the data to produce a non-aliased image. This is because phase information has been irrevocably lost during integration for each pixel. Thus such intuitively simple and appealing suggestions as addition of neighboring pixels to simulate original sampling with $d = 2p$ will fail, at least in detail. Addition of adjacent pixels is a convolution operation specified in one dimension by the kernel (0 1 1), which has the Fourier transform

$$G(s) = e^{-i\pi s} \cos \pi s \qquad (2)$$

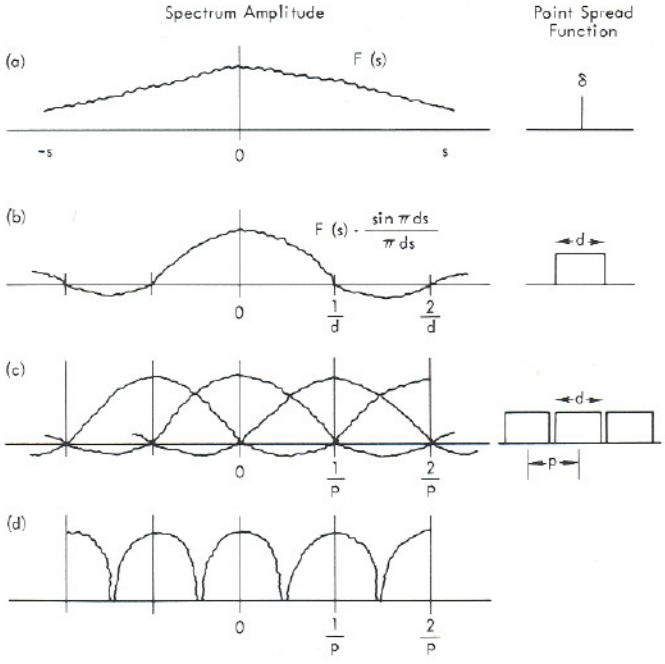

Figure 1.

Summary of sampling processes.

a) F(s) represents the continuous Fourier transform of a scene projected onto the image plane containing the sensor. Frequency is given by s. In this instance, the "lens" was assigned an impulse PSF; real optics might have a sinc or gaussian PSF, resulting in descreased amplitude of F(s) at large s.

b) Integration of the image photo receptors of width d modifies the image spectrum.

c) After sampling the image at intervals of p the discrete image spectrum is the sum of replicates of the continuous spectrum repeated at frequency intervals of p^{-1} cycles/pixel. For CCD cameras, d < p and severe overlapping will occur.

d) If D = 2p or if the continuous image spectrum is bind-limited, overlays of spectrum replicates can be minimized.

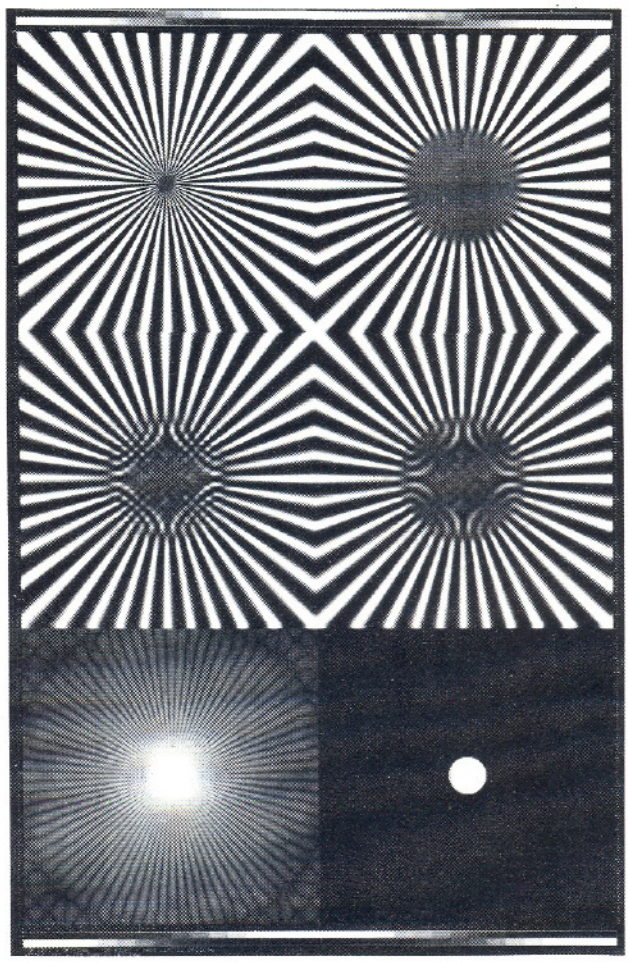

Figure 2.

Effects of CCD-style undersampling on a simple sinusoidal spoke target. Upper right: properly sampled image (d = 2p) of the band-limited target showing absence of information and artifacts at high frequencies. Center left: target sampled as seen by CCD (d =p). Artifacts caused by aliasing are seen instead of high-frequency data. Center right: CCD-sampled image after band-pass filtering showing impossibility of artifact removal. Lower left: spectrum of target; lower right: spectrum of correctly sampled image.

when d = p = 1. Applied to the image spectrum, this filter modifies the spectrum amplitude by cos πs and and also shifts the phase by $e^{-i\pi s}$, yielding a filtered image spectrum of

$$e^{-i\pi s} \cos\pi s \sum_{n=-\infty}^{\infty} F(s + n)\ \mathrm{sinc}(s + n) \tag{3}$$

which is not at all the same as the intended spectrum specified by

$$\sum_{n=-\infty}^{\infty} F(s + n)\ \mathrm{sinc}\ (2(s+n)) \tag{4}$$

This basically reflects the impossibility of separately filtering the individual terms of the series once summation has occured during sampling.

Aliasing in CCD images is most readily solved by degrading the quality of the image-forming optics. This has the effect of reducing the amplitude of F(s) at high frequencies so that overlap of the sprectra of the sampled scene is negligible. This requires, however, that the optical PSF be approximately gaussian in order to avoid introduction of optical artifacts. Simple defocussing is not a solution because of the onset of spurious resolution in the center of the modulation transfer function at defocusings greater than about $5.6\lambda\ f^2$, where λ is the wavelength of light and f is the optics f number[5]. For example, an optical system operating at f/15 and 5000 Å would begin to suffer phase reversals of π at a defocus greater than 0.6 mm.

Illusions in sampled images

Even if a scene is band-limited and properly sampled, there is no guarantee that it will be recognized from visual inspection of an arbitrary pictorial representation of the digital data. The sampling theorem only states that exact reconstruction of the function is possible. In order to do this correctly, it is necessary to interpolate using the infinite sinc function. This is not easy to do, and usually sampled data is displayed without any effort to reconstruct the original function.

One consequence of failure to reconstruct the continuous image before display is the visible presence, obvious in ordered images, of interference patterns resulting from sampling harmonic constituents of the scene near their nodes, where the amplitude is zero. The frequencies of these nodal zones can be specified in one dimension by observing that the sampling function is given by

$$\sum \delta\ (x+n) = |\ \cos\pi x\ |^{\infty} \tag{5}$$

Harmonic constituents of the spectrum are waves in space described by $A \cos 2\pi(\omega x+\phi)$, where A is amplitude, ω is frequency, x is spatial position, and φ is phase. At some x_o the wave will be sampled at its crest. If x_o is taken as the origin, the phase at which sampling occurs is seen to be the periodic function

$$\phi(x) = x(.5 - \omega) \tag{6}$$

To the eye, the characteristic frequency of φ(x) in an image is set by the repetition of the nodal zones, where $\phi = +\pi/2$. These occur twice per wavelength, so the apparent frequency of the nodal patters is 1-2ω. For $\omega < .\bar{2}5$ this nodal frequency exceeds the Nyquist (.5) and is aliased.

Figure 4 illustrates five waves of different ω. Each is oversampled and should resemble a simple sinusoid. Wavelengths from top to bottom are 40, 40/6, 40/10, 40/15, and 40/19 samples. A wave of the sampling frequency has a period of 40/20 samples and is represented as a binary sequence whose amplitude may range from ±1 to ±0 as a function of phase. It is evident that the simple nature of the sinusoid is masked for wavelengths shorter than about four samples, and that interference patterns may actually be the prominent feature of the image.

Visual inspection of images dominated by waves of frequencies greater than .25 cycles/pixel will result in inaccurate interpretation, even though the sampling theorem was observed. Because a discrete frequency spectrum of each of these sampled waves is simply $\mathrm{II}(s) = \delta(s + \omega) + \delta(s - \omega)$, where δ is the impulse function, s is frequency, and ω is the frequency of the wave, it follows that the nodal or interference patterns are completely illusory. In fact, they result from incorrect display of the discrete image.

It is standard practice when making a pictorial representation of a discrete image to let the average photographic density within a small square be proportional to the sampled brightness in the discrete image. This is equivalent to resampling the discrete image to create a continuous picture, using nearest-neighbor interpolation. This is done by convolving the discrete image with the function

$$\sqcap(x) = \begin{cases} 1: & |x| < 1/2 \\ 0: & |x| > 1/2 \end{cases} \tag{7}$$

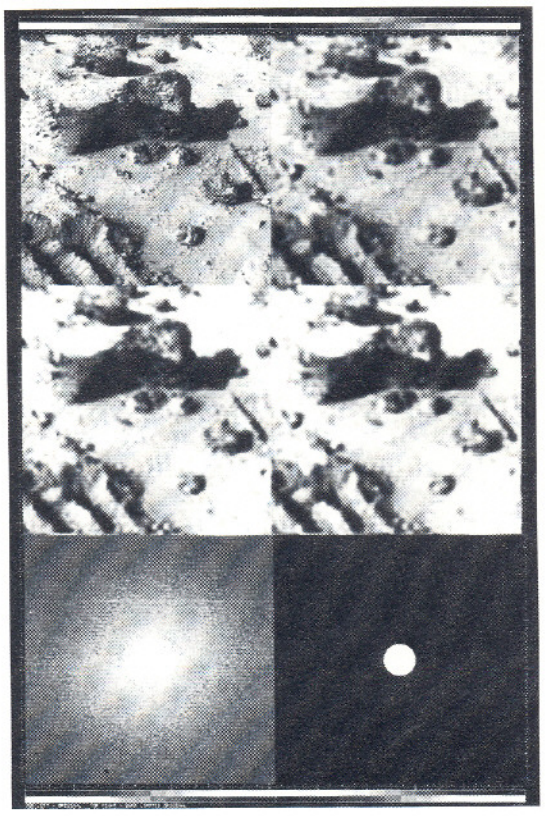

Figure 3.

Effects of CCD sampling on a typical scene. Upper left: continuous scene is simulated by Viking Lander image of Mars surface. Upper right: Properly sampled band-limited image with d = 2p (p = 11 pixels in the "scene"). Center row: scene sampled by a simulated CCD with d = p (left) and band-pass filtered after sampling (right). Bottom row: spectra of "scene" (left) and properly sampled image (right). Effects of aliasing in CCD image of scene are not as noticeable as they are for simple targets.

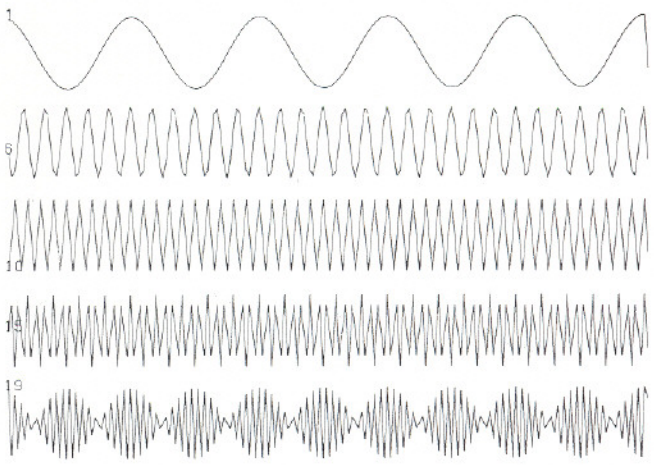

Figure 4.

Sampled sinusoids showing interference effects of the sampling function. Samples have been connected by line segments for clarity. From top to bottom, wavelengths are 40, 40/6, 40/10, 40/15, and 40/19 samples.

or multiplying the transform by sinc s. This approach, however, is incorrect. The proper way to reconstruct the continuous image requires convolving the image by sinc x and resampling at infinitesimally small intervals p. Clearly this cannot be done, but correct resampling to increase the pixel density in the output picture may be done by inserting the fundamental discrete spectrum into a field of zeroes and by capitalizing on a consequence of the similarity theorem to expand the image. This theorem states that a function g and its transform G are related to the function f and its transform F by the expression.

$$F\,(g(ax,by)) = \frac{1}{|ab|}\,G\,\left(\frac{fx}{a}\,,\,\frac{fy}{b}\right) \tag{8}$$

which says that reducing the dimensions of the transform will increase the dimensions of the image in the function domain. This is equivalent to increasing the spacing p^{-1} between replicates of the discrete transform, or decreasing the spacing between pixels. When this is done, the true nature of the scene becomes apparent in its pictorical representation.

To illustrate the resampling operation and to display the visual illusion in two dimensions, Figures 5-7 have been created. Figure 5 consists of rows of sinusoids extending from 0 cycles/pixel (cpp) at the top to 0.5 cpp (Nyquist frequency) in the center to 1.0 cpp at the bottom. Notice that the bottom half is the mirror image of the top half and that this is due to the aliasing of the entire bottom half of the image (the bottom half of the "scene" bears little resemblance to its image). Although the aliasing phenomenon is not the subject of this section, it is instructive to observe it in an image. Interference artifacts due to sampling of waves having frequencies greater than 0.25 cpp can be observed in the form of hyperbolae in the the center and the four sides of the figure. These structures do not actually exist in the target even though they are readily observed in the image. An expansion of the image performed by simply replicating pixels produces no changes because the high frequencies are retained.

In Figure 6 the picture was expanded by a factor of two by using bilinear interpolation. The central hyperbolae are suppressed although the border ones remain. Contrast and resolution have been lost especially near the center of the picture because of the low response of the interpolating function to high frequency waves.

In Figure 7 the image was expanded by a factor of two by separating replicates in the discrete transform as described above. The image contains no waves with frequencies greater than .25 cpp, and the first-order interference artifacts have vanished. Furthermore, the image quality has not been degraded.

As a rule, to insure that digital images will be properly visually interpreted, it is necessary that the band-limited scene be sampled at a rate more than four or five times the band limit frequency. If this condition is not met during data acquisition, it may be achieved by properly resampling the image.

Some interference artifacts will remain until the pixel spacing is reduced to zero, which is the case for a continuous image. However, the strongest of these are removed by an image expansion of two times. Even this may be infeasible with the large images produced by current sensors because Fourier transforms must be computed. A compromise would seem to be interpolation using a cubic spline[5] or other function of finite size but performance superior to bilinear interpolation.

It is evident that in an arbitrary image of a real scene or of random noise interference artifacts are not obvious (Figure 8). This is probably because they are masked by strong non-artifactual harmonic constituents of the image. However, if these waves are suppressed by filtering, then the only information left will be artifact. Thus we might predict that in high-pass filtered images interference artifacts will be apparent, but inspection of Figure 8 indicates that this does not seem to be the case. The directional artifacts first noticed by Goetz et. al.,[6], in filtered Landsat images and described in detail by Gillespie[7] are evident in Figure 8. These artifacts take the appearance of a directional fabric or pattern of oriented quasi-linear features. They are strongest in images filtered by functions having high contrast and sharp features in the spectral domain. It has been argued that they are real features of the scene; however, their artifactual nature is emphasized by their presence in filtered random noise. They do not appear to be related to the interference phenomenon because, unlike interference artifacts, they are not affected by resampling.

In Figure 8 directional fabrics are introduced by convolution of a Landsat scene and also random noise (shown in the top row) with a one-dimensional uniform weight or "box" filter having the Fourier transform

$$A_L(s) = \frac{1}{2n+1} + \frac{2}{2n+1}\sum_{j=1}^{n}\cos 2\pi sj \tag{9}$$

For large n, A_L approximates sinc ((2n + 1)s). For figure 8, a one line (vertical) by eleven sample (horizontal) box filter (2n + 1 = 11) was chosen. The high-pass filter is given by

$$A_H(s) = 1 - A_L(s) \tag{10}$$

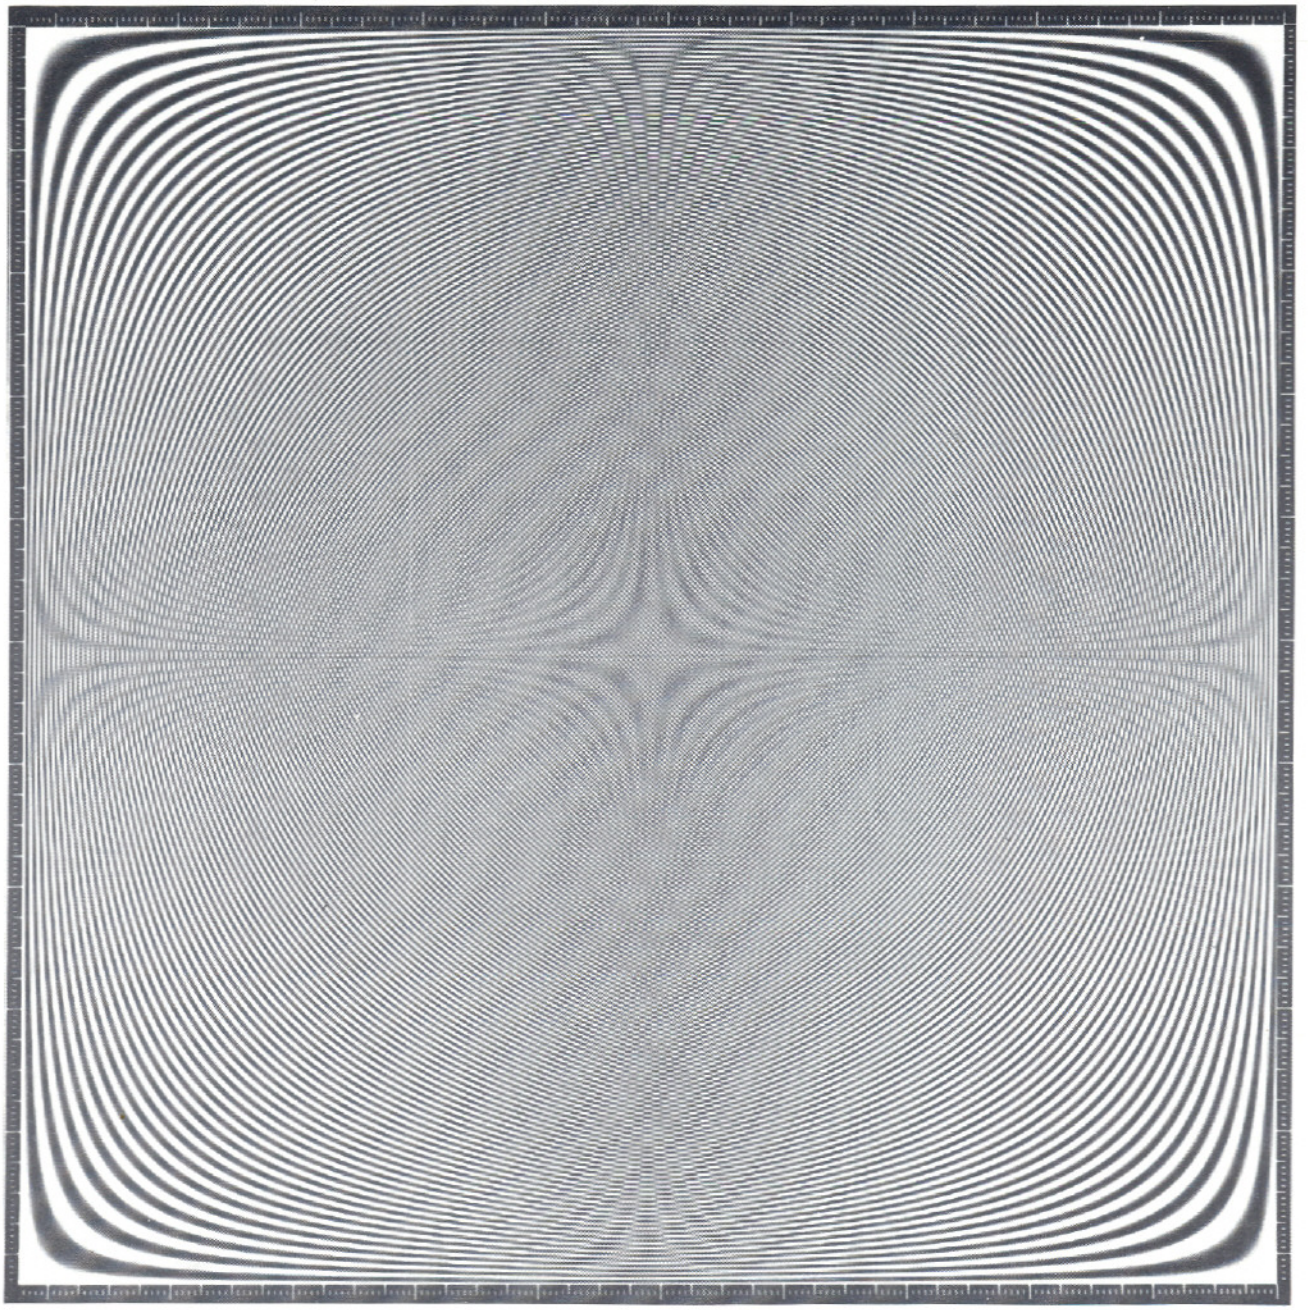

Figure 5.

Sinusoid target showing interference patterns as hyperbolae at center and in central regions of the four edges. Each line in target is a sinusoid in the horizontal direction. Frequency changes in vertical direction from 0 at top line to 1 at bottom. Note mirror image caused by aliasing around the middle line (at the Nyquist frequency). Image size has been expanded by pixel replication.

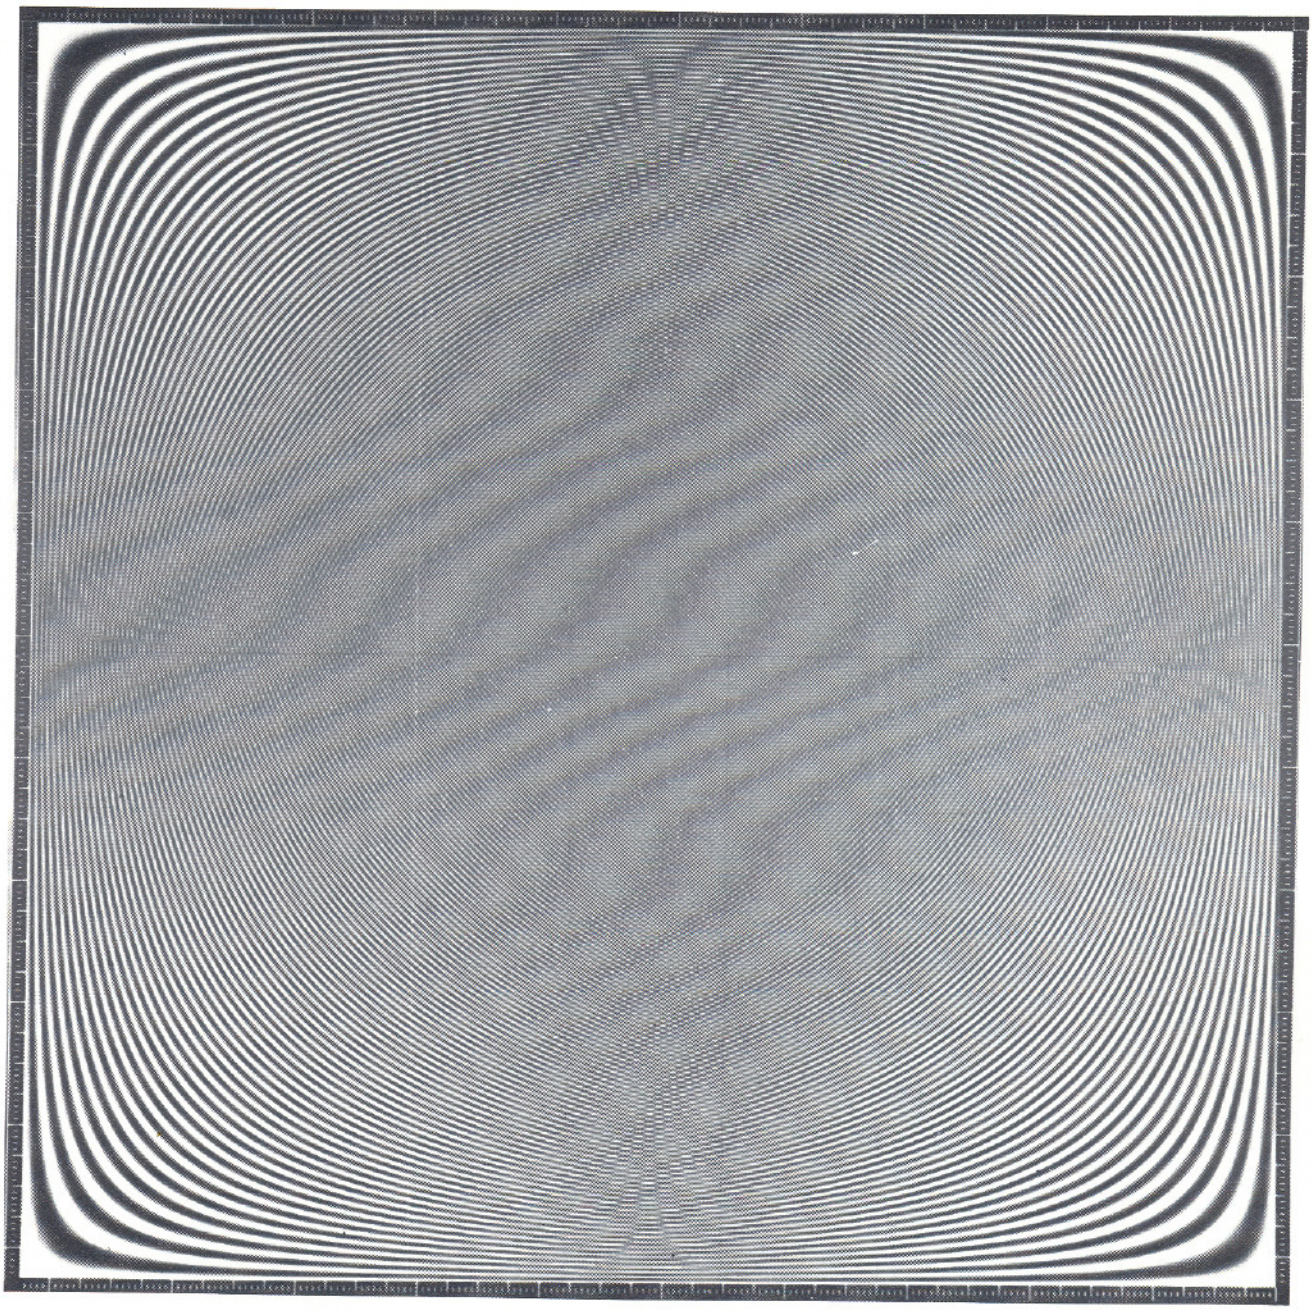

Figure 6.

Sinusoidal target showing artifact suppression after two-fold size expansion on using bilinear interpolation. Contrast at high frequencies is also lost.

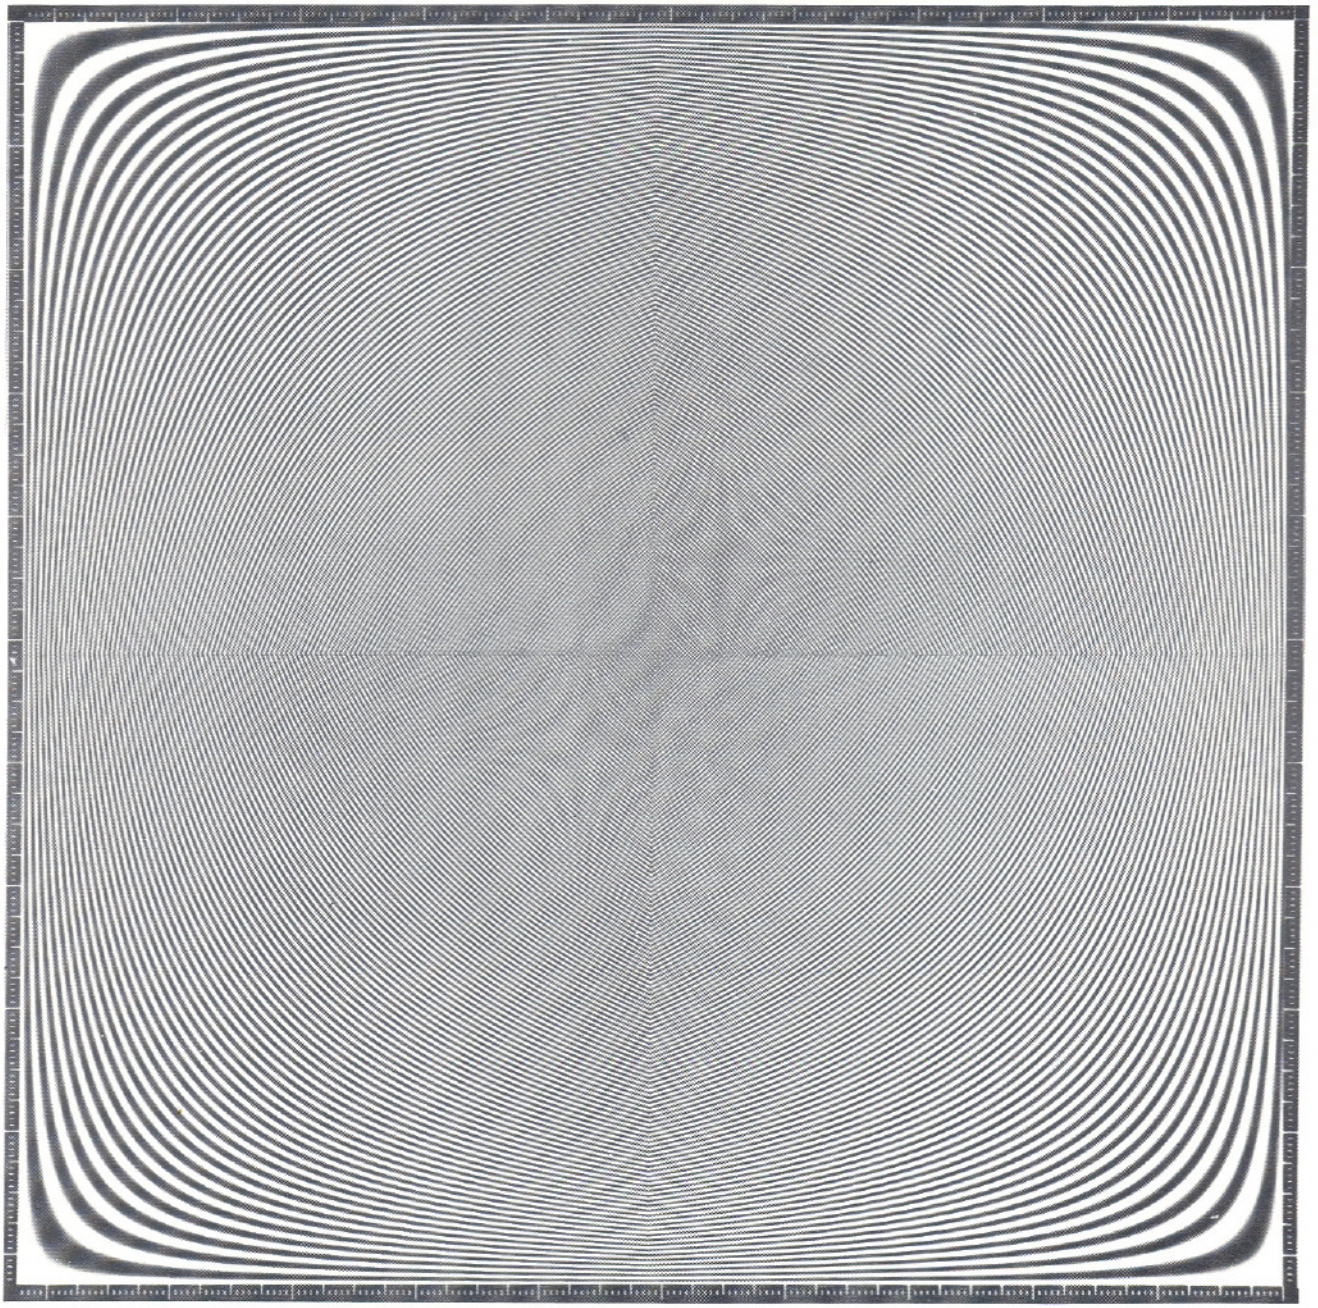

Figure 7.

Sinusoidal target showing artifact suppression after two-fold size expansion using the similarity theorem (sinc interpolation). The few artifacts remaining would be removed by further interpolation to create a more nearly continuous image. Contrast at high frequencies is preserved.

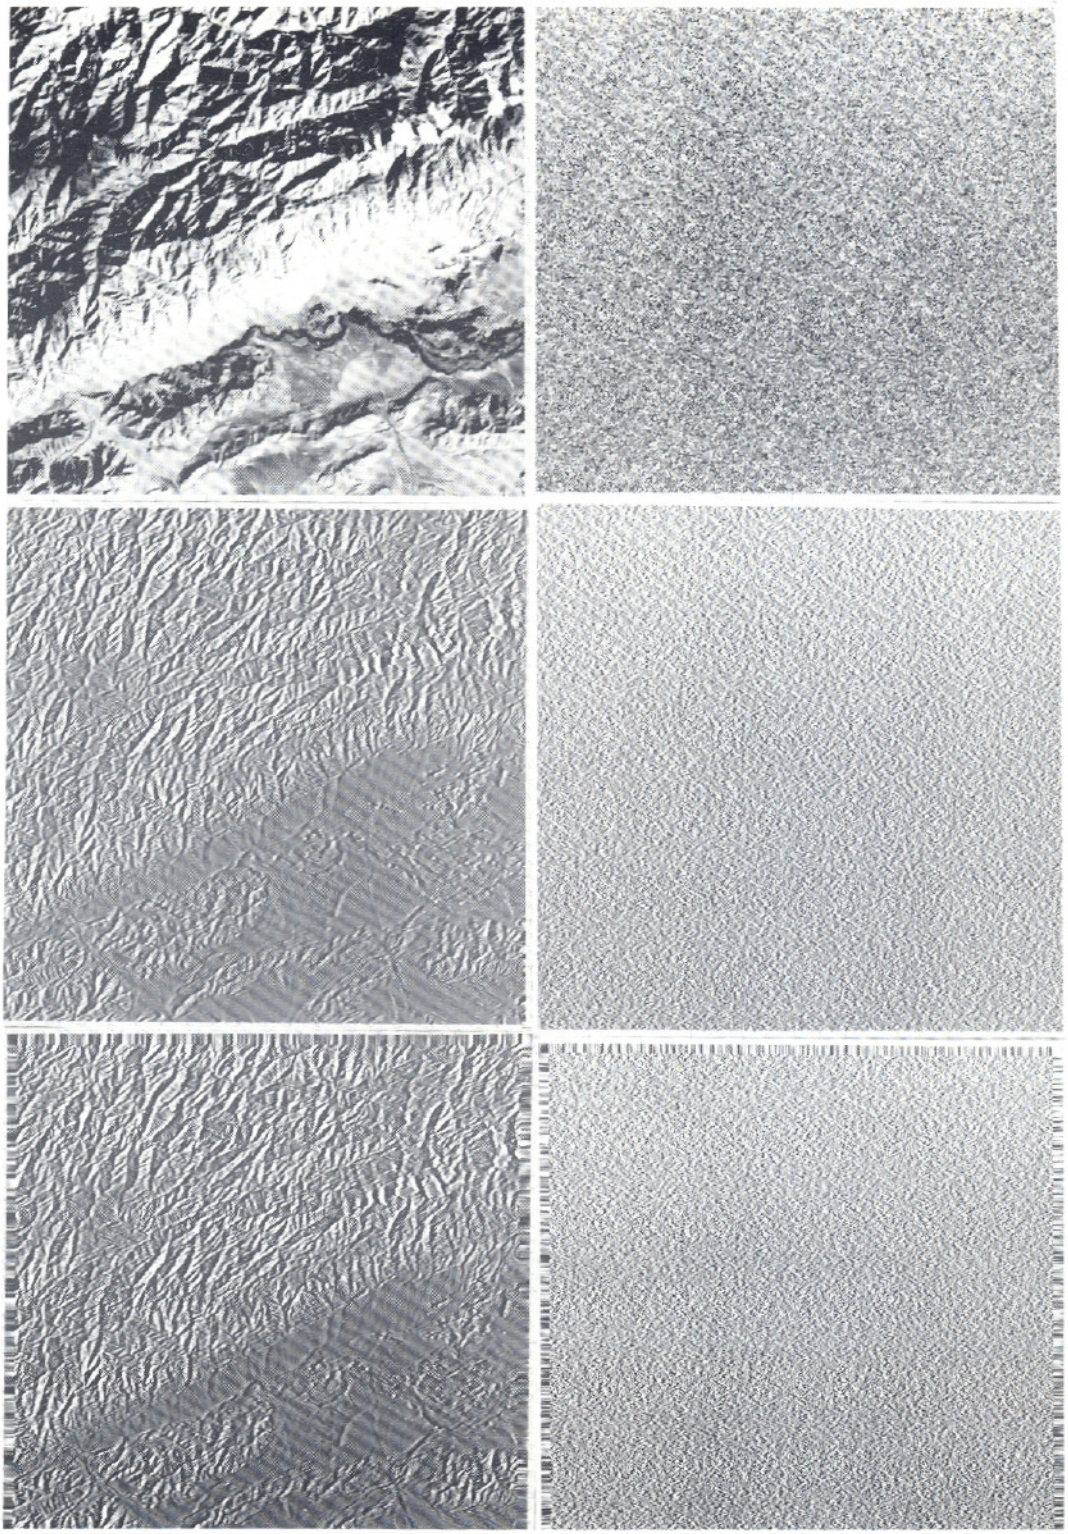

Figure 8.

Directional artifacts introduced in filtered images are not suppressed by interpolation. Left column shows processing on part of NASA Landsat image 1074-04253 of the Astin Tagh mountains of western China. Image width is 40km, north is 15° counter-clockwise from up. Right column shows identical processing on normally distributed random noise. Top row contains original images. Center row shows same images after high-pass filtering ("edge-enhancement") by convolution with a 1 x 11 element uniform weight "box" filter. Artifacts in both images are obvious as a directional fabric comprised of short linear features oriented about + 20° from horizontal. Bottom row illustrates doubling of image size. Artifacts are unaffected by sinc interpolation.

so frequencies in the sample direction less than about .09 cpp were suppressed by the filter. The filtered images are seen in the middle row of Figure 8. This particular filter was chosen because it is commonly used to "enhance edges", particularly in Landsat images.

That the directional fabric in the filtered images is different from the interference phenomenon shown in Figures 4-7 can be demonstrated by doubling the size of the image, resampling using the similarity theorem as before. The directional fabric remains in the sampled pictures in the bottom row of Figure 8, although resampling suppresses interference artifacts.

Effects of quantization in spatially filtered images

Quantization of real and continuous scene brightness to integer values is a characteristic of all digital images. In the illustrations above, the quantization interval was 2^7 so that a sinusoid ranging in amplitude from -1 to +1 was encoded in 256 gray levels. In real scenes contrast may be reduced so that only a few quanta describe the range of amplitudes of a harmonic component of the image spectrum. This is commonly true for high frequency waves in unfiltered images, but may be true for any component in a filtered image. A problem can occur if the dynamic range of the sampled data is sufficiently narrow that integer truncation seriously modifies the true intensity values in the scene. In this event the quantized sinusoid is represented by series of stacked square functions, or after sampling by the sum of truncated sequences of δ functions. For instance, a sine wave of frequency .125 cpp which is quantized in two levels may be represented by the sequence (...-1, -1, -1, 0, 1, 1, 1, 0, -1, -1, -1, 0...). The Fourier transform of this sequence is not simply $\delta(s + .125) + \delta(s - .125)$, the transform of the sinusoid, but is instead the sum of the transforms of the three elementary sequences (0, 0, 0, -1, 0, 1, 0, 0...), (0, 0, -1, 0, 0, 0, 1, 0...), and (0, -1, 0, 0, 0, 0, 0, 1...), all of which are periodic with a wavelength of 8 samples and which add to produce the original sequence. The transform of these impulse pairs are themselves sinusoids having frequencies of .125, .25, and .375 cpp and the spectrum of the quantized sinusoid is

$$\sum_{n=1}^{3} \sin(2\pi n.125s) \tag{11}$$

In this example, if the spectrum were quantized at the same level as the image, each sinusoid in the spectrum would itself be a series of alternating impulses

$$\sum_{m=1}^{3} \sum_{n=-\infty}^{\infty} (-1)^n \ \delta(s - .125mn) \tag{12}$$

Thus harmonic overtones are created by the act of quantization, and these overtones may have frequencies exceeding the Nyquist frequency. In this case, the overtones will be aliased.

Figure 9 illustrates this phenomenon on a synthetic sinusoid image. The figure consists of four rows, each showing a sinusoid at the left, the log amplitude of the Fourier transform of this sinusoid to its right, and a high-pass filter applied to the sinusoid at far right (after contrast enhancement). The difference between the four rows is in the dynamic range of the sinusoid. The frequency of the wave was controlled by populating a single element in an otherwise empty spectrum and inverse transforming into the function domain. The sinusoid thus created was then scaled to the different amplitude ranges shown in the left column of Figure 9. The Fourier transforms of the quantized sinusoids, displayed as negatives for visibility, are reproduced in the center column and clearly show the presence of harmonic overtones.

Filtering of an image affected by this phenomenon may modify the amplitudes at the fundamental and overtone frequencies differently. To illustrate this, frequencies less than or equal to the fundamental frequency of the quantized sinusoid in Figure 9 were suppressed. The filtered images, shown in the right column of Figure 9, clearly show the presence of energy at overtone frequencies. In images of real scenes, this phenomenon must result most noticeably in an increase of high-frequency noise. However, all frequencies will be affected in an image dependent fashion, producing artifacts which are only indirectly related to the continuous scene.

Quantization artifacts cannot be removed from an image by filtering, except by also removing all fundamental energy at the same frequency. Resampling is likewise of no use. They are best minimized by selecting appropriate signal amplification prior to digital encoding. However, in any digital image some quantification artifacts will be found, and the analyst must anticipate encountering them during careful inspection of digital images.

Conclusions

Digital images exhibit three kinds of artifacts. These arise from violation of the sampling theorem, improper display, and quantization or integer representation of scene radiance.

Solid-state electronic cameras such as CCD's are prevented from sampling in accordance with the sampling theorem by their physical design. The principal consequence of this is undersampling and aliasing resulting in visual disruption of pictures of simple geometric targets or highly structured scenes. However, image

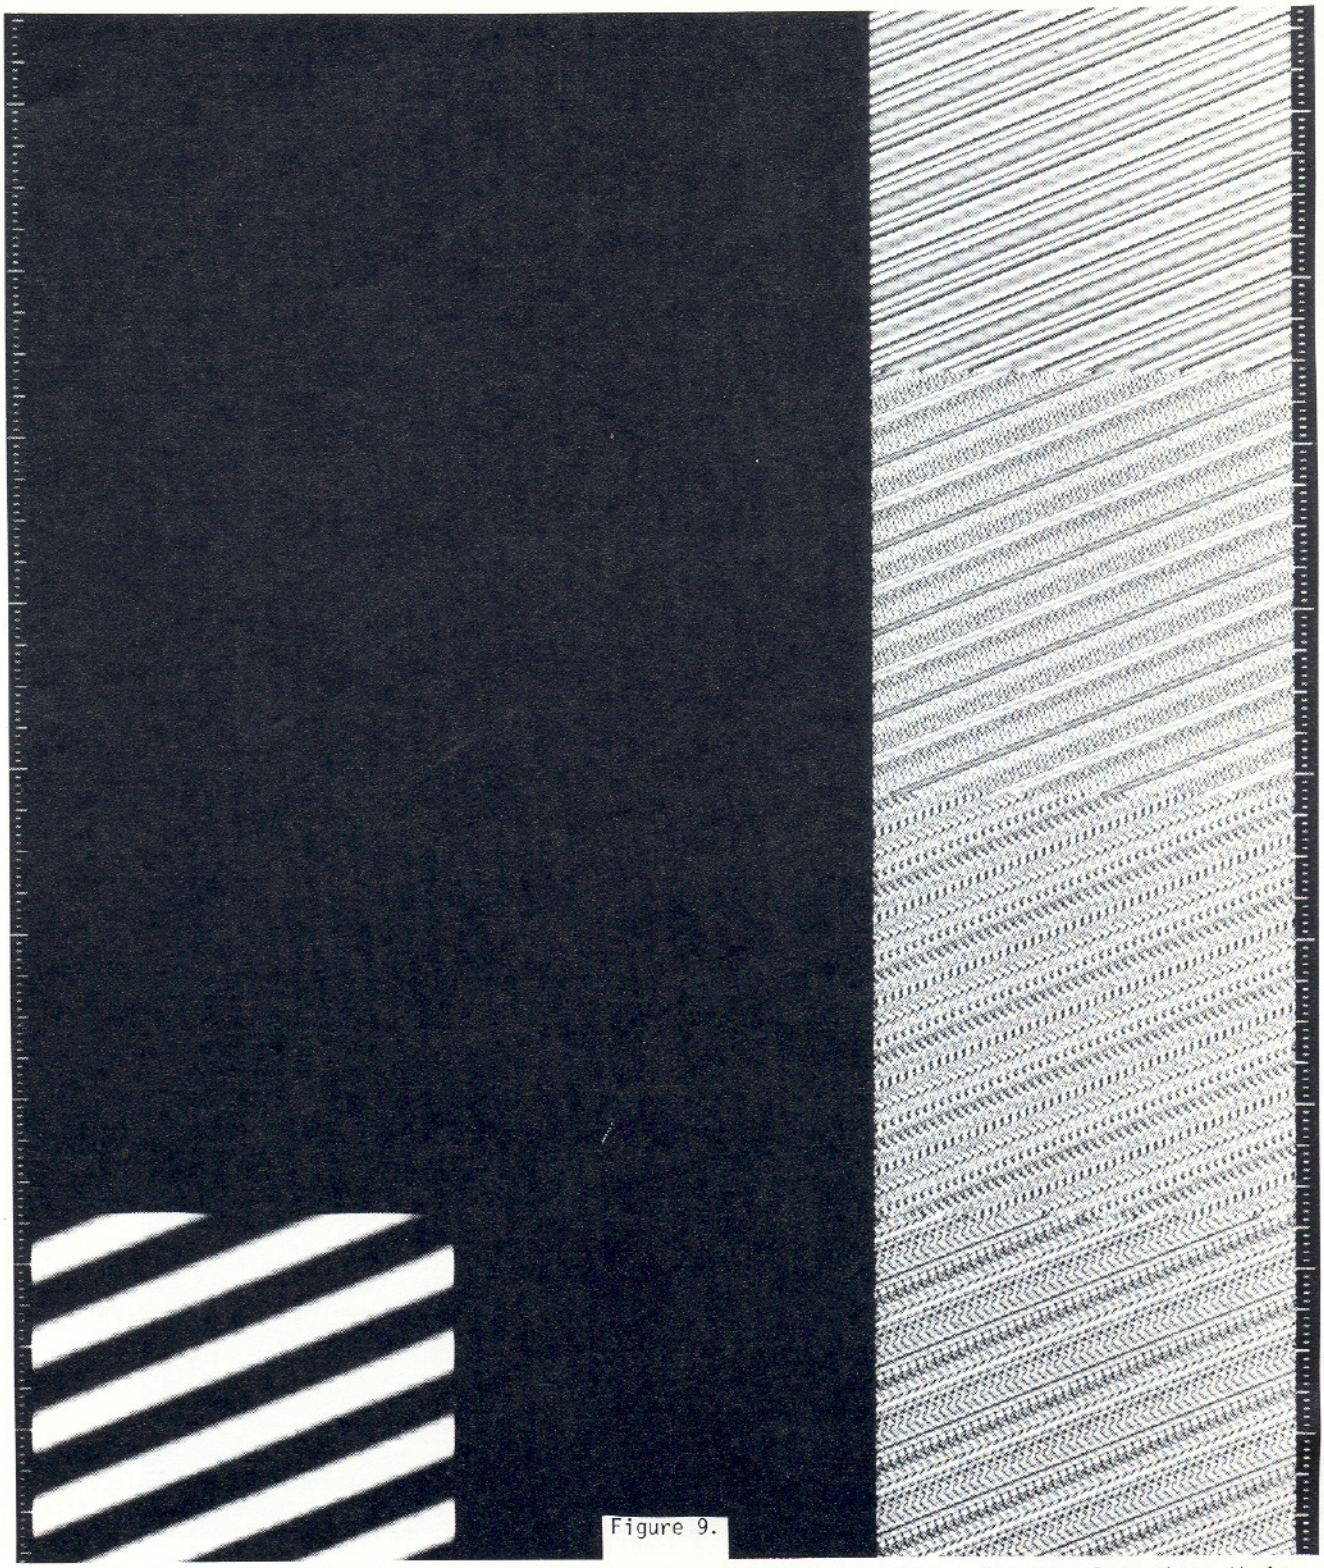

Figure 9.

Quantization artifacts in sinusoids. Left column shows sinusoidal targets. Center column shows their spectra. Although the target is a simple sine wave, overtones are present in the transforms. Right column shows filtered target with power at and below the fundamental frequency of the sine wave suppressed. In row 1 the sinusoid has a brightness range of 2^3 gray levels; in row 2 the target occupies 2^4 gray levels; in row 3 it occupies 2^6 gray levels; and in row 4 the brightness range is 2^8. As the target contrast increases more overtones appear but the power at the fundamental frequency increases. For a continuous target an infinite number of overtones of zero amplitude would be present. Contrast in filtered images has been adjusted for display.

filtering may reveal these effects even in images of ordinary scenes. Critical sampling or oversampling may be achieved by degrading the optics to produce a broad pointspread function, which band-limits the image projected onto the sensors.

Properly sampled scenes may be reconstructed by interpolation of the discrete image. This costly process is widely over-looked, resulting in the introduction of interference patterns in images. Again, simple geometric targets are most obviously affected but filtered images of most scenes are adequately displayed by conventional means.

Quantization of image data necessarily introduces energy at harmonic overtones of the image spectrum. This energy is aliased if the frequency of the overtones exceeds .5 cycles/pixel. It can not be selectively removed from the image by filtering, and the best way to suppress it is to maximize the amplification of the sensor before digital encoding. The effect is reduced as the dynamic range is increased.

Acknowledgments

This paper presents one phase of research conducted at the Jet Propulsion Laboratory, California Institute of Technology, under NAS 7-100, sponsored by the National Aeronautics and Space Administration. Support was also given through California Institute of Technology President's Fund Grant number PF-116.

References

1. Goodman, J. W., "Introduction to Fourier Optics," McGraw-Hill, P. 25, 1968.
2. Bracewell, R., "Fourier Transform and Its Applications," McGraw-Hill, p. 191, 1965.
3 Root, G., "A Qualitative Study of the Trade-off Between Sample Spacing and MTF for Discrete Image Sensors," J.R. 900-631, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, 91103, 1973.
4. Linfoot, E. H., "Fourier Methods in Optical Image Evaluation," Focal Press, p. 22, 1964.
5. Rifman, S. S., "Digital Rectification of ERTS Multispectral Imagery," Symposium on Significant Results Obtained From Earth Resources Technology Satellite-1, NASA SP-327, pp. 1131-1142, National Aeronautics and Space Administration, Washington, D.C., 1973.
6. Goetz, A. F. H., Billingsley, F. C., Gillespie, A. R., Abrams, M. J., Squires, R. L., Shoemaker, E. M., Lucchitta, L., and Elston, D. P., 1975, Application of ERTS Images and Image Processing to Regional Geologic Problems and Geologic Mapping in Northern Arizona: Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, Rpt. 32-1597, 181 p.
7. Gillespie, A. R., "Directional Fabrics Introduced by Digital Filtering of Images," Proceedings, 2nd International Conference Basement Tectonics, University of Delaware, Newark, Delaware, July 1976, Contribution No. 58, pp. 500-507.