Improved Reduced-Resolution Satellite Imagery

James Ellison and Jaime Milstein
The Aerospace Corporation
El Segundo, California 90245

513 - 61

47263

p. 16

Abstract

The resolution of satellite imagery is often traded-off to satisfy transmission time and bandwidth, memory, and display limitations. Although there are many ways to achieve the same reduction in resolution, algorithms vary in their ability to preserve the visual quality of the original imagery. These issues are investigated in the context of the Landsat browse system, which permits the user to preview a reduced resolution version of a Landsat image. Wavelets-based techniques for resolution reduction are proposed as alternatives to subsampling used in the current system. Experts judged imagery generated by the wavelets-based methods visually superior, confirming initial quantitative results. In particular, compared to subsampling, the wavelets-based techniques were much less likely to obscure roads, transmission lines, and other linear features present in the original image, introduce artifacts and noise, and otherwise reduce the usefulness of the image. The wavelets-based techniques afford multiple levels of resolution reduction and computational speed. This study is applicable to a wide range of reduced resolution applications in satellite imaging systems, including low resolution display, spaceborne browse, emergency image transmission, and real-time video downlinking.

1 Background

Satellite imaging systems like Landsat, collect and downlink large quantities of data. Associated ground systems may further process and store this data, as well as provide for its dissemination. Limitations on computer storage, transmission bandwidth, transmission time, and digital display resolution may restrict the amount of data used to represent an image. These issues affect image processing and storage on-board the satellite, preparation of the image for transmission, downlinking of image data, and reconstruction, storage and dissemination of the image to the end user. Such problems may be addressed by data compression techniques, by reducing image coverage, by reducing the number of gray levels (or colors), or by reducing

resolution. Some resolution-reducing techniques (for instance, edge-avoiding convolution) are scene-dependent. This paper considers only general resolution reduction algorithms. In particular, wavelets, a recently developed mathematical transform, is utilized as a resolution-reducing device and compared with some conventional algorithms for resolution reduction.

Section 2 discusses an example of a typical problem requiring resolution reduction. Some common methods for handling the problem are discussed, and the idea of wavelets is introduced. A quantitative measure is used for crude quality comparisons. Potential applications of a good solution to the resolution reduction problem are also suggested. In Section 3, resolution reduction algorithms are applied

to Landsat imagery, and numerical comparisons are given. Based on visual examination, experts[*] concluded that wavelets preserves image quality better than other methods tested. In Section 4, aerial images are used to illustrate the visual quality resulting from alternative methods. Conclusions are summarized and applications are suggested in Section 5.

2 Reducing Image Resolution

In this section we briefly discuss resolution-reducing algorithms based on subsampling, convolution and wavelets. We conclude by noting the applicability of a good resolution-reducing algorithm to other practical problems.

2.1 An Example of a Problem in Resolution Reduction

Suppose we wish to display a full-size M-by-N pixel image on a P-by-Q pixel screen, $P << M$, $Q << N$. This problem arises, for example, when the full 5984-by-6200 pixel scene presented by the Landsat Thematic Mapper (TM) is to be displayed on a conventional personal computer monitor, which may permit up to 512 rows and 650 columns. Under these circumstances, it is impossible to display the full scene on the pixel-limited display without sacrificing *resolution*, the minimal distance at which small adjacent objects can be distinguished [Rosenfeld and Kak, 1982, p. 215]. In this example, the original 5984-by-6200 pixel scene has 16 times the resolution of a 374-by-388 pixel rendition of it.

The visual degradation of a reduced-resolution image depends on the resolution reduction technique. Our goal is to reduce resolution in such a way that the eye's perception of the displayed scene is as close as possible to that of the full resolution scene. This is what we mean by the "display problem.".

The resolution-reducing algorithms discussed below have a common property which enables us to compare the reduced resolution imagery they produce: each algorithm can be represented as a series of applications of a 2-to-1 resolution-reducing technique, whether subsampling-by-2, wavelets, or some other methodology. If 2-to-1 resolution reduction is applied k times, then the algorithm produces a 2^k-to-1 resolution-reduced image, directly comparable to the image produced by applying any other 2^k-to-1 resolution-reducing algorithm. For example, since subsampling-by-16 amounts to 4 iterations of subsampling-by-2, it is reasonable to make quality comparisons between the results of subsampling-by-16 and that of applying 4 iterations of wavelets to the same original image: the resulting images have the same resolution and differ only in the algorithm applied.

We now focus on subsampling and wavelets. Each provides a practical, computationally efficient solution, independent of scene, subject matter, and prior degradation. Yet, subsampling and wavelets represent opposite extremes of mathematical soundness and visual appearance.

---

[*] The imagery discussed in this paper was presented to image exploitation professionals and other scientists of the EROS Data Center of the U. S. Geological Survey, together with a wide variety of scientists from the Landsat user community. "Expert conclusion" refers to the unanimous opinion of this population.

2.2 Subsampling

The most straightforward way to reduce the size of the array without losing coverage is by *subsampling*, that is, assembling an image from a regularly spaced subset of pixels in the original array. *Subsampling-by-n* involves the selection of pixels from every nth column along every nth row. Thus, subsampling-by-n results in an n^2-to-1 reduction in the number of pixels and an n-to-1 reduction in resolution.

Subsampling is widely employed as an efficient solution to the display problem. For example, as noted in Section 3, subsampling-by-16 is currently employed in preparing the Landsat browse product for the user community from the original TM scene. In principle, subsampling requires no computation and is therefore optimal in computational efficiency.

However, visual defects are introduced by subsampling-by-n. As n becomes a significant fraction of the width (in pixels) of any feature, these defects worsen. The following defects are typical of imagery produced by subsampling.

- Edges of solid bodies assume a staircase appearance.

- Even when a feature covers most of an n-by-n square, this information is lost if the sampled pixel happens not to fall within the feature.

- Separate, distinct features can merge.

- Linear features, i.e., long narrow features like roads, communication lines, and rivers, can disappear altogether.

- Small features can vanish.

- When the retained pixel is unrelated to its surroundings, this pixel shows up in the reduced resolution image as apparent noise.

- Artifacts can be introduced by random noise. As noise increases, larger artifacts become more common.

As resolution is reduced, some loss of image quality is unavoidable. However, much of the loss of visual quality just described is peculiar to the subsampling process itself. The obvious problem with subsampling is that the retained pixels provide no information about the discarded pixels. Generating the same amount of data, more effective resolution-reducing methods capture more representative visual data from the full resolution image than does subsampling. Instead of picking one pixel out of a fixed position in the n-by-n square, they define a value of the new pixel that better represents the pixel values in the n-by-n square it is replacing.

2.3 Convolution

Convolution, or *spatial filtering*, creates a new image by replacing each pixel value with a weighted average of its surrounding pixel values. As a resolution-reducing technique, convolution may be regarded as a generalization of subsampling, in which a convolution is performed at each subsampled point. The corresponding pixel in the new image is given the value of the convolution. When that convolution is the unit impulse function (1 surrounded with 0's), this process reduces to subsampling-by-n.

165

Convolutions have been tailored to widely varying purposes, including edge enhancement, smoothing, noise reduction, etc. Convolution has also been combined with other algorithms for selective application to scenes or parts of scenes. Depending on coefficients of the convolution, the pixels in the reduced scene may retain useful information from those discarded from the original scene. For this reason, the resulting image may be less subject to many of the defects characteristic of subsampling.

Computation required for any specific convolution is proportional to the number MN of pixels in the original scene: each of the PQ pixels in the n-to-1 reduced resolution image represents up to n^2 multiplications and additions, and $(n^2)PQ = MN$. Convolution offers a fast method for resolution reduction, though not as fast as subsampling. However, in our experience, for a specific convolution, apparent degradation typically varies greatly, depending on the nature of the scene, its texture, etc.

2.4 Wavelets

Wavelets may be regarded as a special kind of convolution. Wavelets developed rapidly from 1983 onward. There is now a large and rapidly growing literature on the subject [Meyer, 1986; Mallat, 1989; Chui, 1991; Press, 1991]. The present work uses coefficients defined by Daubechies [Daubechies, 1988]. Our purpose here is to discuss wavelets only to the extent necessary to provide a context for the present application.

As commonly employed, the term "wavelets" refers to a data compression technique with many elegant properties,

both theoretical and practical. When applied to an image represented by a $2P$-by-$2Q$ array, wavelets generates four P-by-Q arrays. One array, called the *smooth image*, is a reduced resolution version of the original image. The values in the present study were Daubechies's D_4 coefficients (or weights): $\frac{1}{4}(1+\sqrt{3})$, $\frac{1}{4}(3+\sqrt{3})$, $\frac{1}{4}(3 - \sqrt{3})$, $\frac{1}{4}(1 - \sqrt{3})$.

The computation time required for wavelets is, like convolution, proportional to the number of pixels in the original image. Used for resolution reduction, the number of pixels dealt with in each iteration of wavelets is $\frac{1}{4}$ that of the previous iteration. Thus, k iterations of wavelets, applied to an M-by-N pixel image, has a computation time proportional to $MN[1 + \frac{1}{4} +...+ (\frac{1}{4})^{k-1}]$. Since $[1 + \frac{1}{4} +...+ (\frac{1}{4})^{k-1}] < 1\frac{1}{2}$, for all positive k, the computation time for wavelets resolution reduction remains proportional to the number of pixels in the original image, independent of the size of the final reduced-resolution image. (In practice, clever implementation can significantly reduce the amount of computation.)

Only the smooth images are needed for the purpose of resolution reduction. Thus, k iterations of wavelets resolution reduction generate an image of the same 2^k-to-1 resolution reduction as k iterations of subsampling-by-2 (i.e., subsampling-by-2^k). The results of these algorithms are compared in Sections 3 and 4.

In a certain well-defined sense, for a given resolution reduction 2^k-to-1, k iterations of wavelets better preserve image quality and are not prone to pronounced artifacts such as those associated with subsampling.

2.5 Some Reduced Resolution Problems

Any technique that leads to better quality reduced-resolution imagery has many potential applications in Landsat and other satellite imaging systems. A few such applications are noted below.

The Landsat display problem

A full resolution image from the current Landsat Thematic Mapper typically requires a 5375-by-6468 pixel array. Any attempt to display such an image on a computer monitor, even one capable of a 1012-by-1012 display, requires a solution to the display problem. Moreover, a flexible solution would permit the individual user to tailor the final resolution to his or her display capability.

Landsat browse

The full resolution Landsat TM image consists of approximately 280 megabytes of data, about 40 megabytes per spectral band. The Landsat browse product is reconstituted for the user on location from data transmitted over phone lines or the Internet. Currently, three bands of data are reduced from 40 megabytes to 156 kilobytes per band using subsampling-by-16, to produce a false color, reduced-resolution version of the original image of about 335-by-404 pixels. This process avoids most of the data storage and transmission that would otherwise be required. Based on a "quick look" at the resulting image, the user can then request (and pay for) full-detail imagery of interest. A superior solution is one that gives better quality imagery of the same resolution than currently available. It would also be useful for the user to be able to select from a range of resolution-reductions. This would add to the current

full-resolution and 1/16th resolution alternatives a range of cost and bandwidth-intensive choices.

Downlink browse

This application postulates a high resolution satellite sensor with a downlink bandwidth constraint. The principle of operation is similar to that of the Landsat browse: the satellite downlinks a reduced resolution image for approval before transmitting (or even collecting) the full resolution image. This way, depending on the image and resolution desired, downlink bandwidth can be used or conserved.

Emergency spaceborne image communication

The downlink of a spaceborne remote sensing system could be jammed or otherwise dysfunctional. In this case, the satellite could be instructed to transmit a reduced resolution image to a communication satellite network for retransmission and downlinking.

Animation or real-time video downlinking

This scenario envisions the adventure movie scenario of an interactive capability enabling an imaging satellite to zoom in on a selected target area. Frequent images (animation) or real-time video would then be downlinked. Among the challenges in designing such a system is that of limited downlink bandwidth. However, the human eye is more forgiving of reduced resolution when viewing animation and video than when examining an individual image. This facilitates trade-offs of resolution reduction in favor of frame frequency. Suppose, for example, that the system has a 0.1 meter earth surface resolution and can downlink 24 megabytes of imaging data per second, with the ability to take and process up to 24 frames

per second. Such a system might be able to transmit a full-resolution, single-band Landsat-quality scene in 1.6 seconds. This system could instead be instructed to transmit 24 frames of a 100-by-100 meter square of the earth surface per second, at full one meter resolution. A good (and fast) 16-to-1 resolution-reducing algorithm might provide interactive real time video coverage of a 1.6 kilometer square.

3 The Landsat Browse Study[**]

3.1 The Current Browse Product

The current browse product provides users with an economical reduced resolution preview of Landsat imagery. The browse product is reconstituted on location from data transmitted over phone lines or the Internet. From this preview, the user decides whether or not to request full-resolution imagery.

Currently, subsampling-by-16 is applied to 3 bands of full resolution Landsat imagery in order to provide a single RGB reduced-resolution browse product of about 335-by-404 pixels. The user thus views 0.4% of the pixels from each of three bands of the original full resolution image. The subsampling deficiencies discussed in Section 2.2 are readily apparent in practice, as seen in imagery found in the next section.

The object of the study was to develop and investigate resolution-reducing algorithms that produce superior quality browse imagery over the full range of geographic scenes. In particular, such deficiencies in the current browse product as the potential loss of linear features should be overcome. The investigators established three ground-rules as fundamental to the study:

- Every candidate algorithm must produce browse images of the same resolution as those generated by the current system.

- In order to make the browse product available to the user in near-real time, the computer processing required to generate the browse image must not add more than 3 minutes to the total service delay.

- The browse product must be effective with the full range of geographic imagery.

In the course of the study, several algorithms were investigated: subsampling, wavelets, 3-by-3 convolution, and various hybrid algorithms. These algorithms varied in the quality of the resulting browse product and in processing time.

In the final study phase, experts visually compared the various browse images and products both to the full resolution image and to one another. In the earlier study phases, resolution-reduced images were compared in terms of an objective measure we now describe.

168

3.2 A Measure of Image Degradation

This study used a quantitative measure we refer to as the *sequential correlation coefficient (SCC)*, defined as the average correlation between the intensity at a pixel and that of its immediate neighbor on the right. (This measure is not to be confused with more sophisticated imagery measures involving two dimensional statistical correlation.) The SCC can be used as a crude measure of image degradation. In principle, the SCC can assume any value between -1 and 1, the more positive the value, the less the average change. For example, the SCC of an image is 1 if all pixels in the same row have the same intensity. An SCC of 0 suggests a completely random "snow-like" image (e.g., the TV screen when a channel is not broadcasting). For practical purposes, the SCC of recognizable imagery is generally well above .60.

For insight into the significance of the sequential correlation coefficient, compare almost any scene or picture of interest to "snow". A "real" scene tends to be a patchwork of regions and well-defined objects or features. Two adjacent pixels are more likely than not to fall within the same feature or region, have similar intensity, coloring, etc. In the "snow" scene, however, even adjacent pixels are likely to be dissimilar. For any "real" scene the greater the distance between the pixels, the more likely they are to fall into different regions or features, having unrelated, widely varying colors (intensities in various bands). Thus, for any real scene, as the resolution is reduced, the SCC can be expected to decrease. This is clearly true of subsampling-by-n, as n increases.

The SCC is not a completely reliable measure of image quality as interpreted by the eye. For example, "turning down the contrast" of an otherwise good quality image can reduce the eye's perception of quality while increasing the SCC. As a practical matter, a one or two-percent difference between SCCs is unpredictive of comparative visual quality.

In the early phases of the study, the SCC proved a useful heuristic for comparing image degradations caused by alternative reduction algorithms. Final conclusions were based on the judgment of expert viewers representing the user community and were consistent with the SCC-based findings.

3.3 The Three-Phase Browse Study

Phase 1 of the study was an assessment of a wide variety of candidate algorithms and an initial proof-of-concept of iterated wavelets as a resolution-reducing methodology. Phase 1 used *Landsat P data* - Landsat full resolution imagery after radiometric and geometric correction. Phase 2 investigated two additional algorithms, checked processing speeds, and extended the investigation to *Landsat raw data* (i.e., full resolution images not radiometrically or geometrically corrected). Phase 3 investigated two additional algorithms, each computationally faster than wavelets and more effective than subsampling. Table I surveys the algorithms tested in the course of the study. In addition to a wide variety of Thematic Mapper images, Phase 3 included digitized aerial imagery with resolutions higher than that of the current Landsat. Examples of these reduced resolution images are found in Section 4.

TABLE I
Reduced Resolution Algorithms
Investigated in Each Study Phase

Algorithm	Phase 1	Phase 2	Phase 3
Subsampling			
Subsampling-by-2	x		
Subsampling-by-4	x		
Subsampling-by-8	x	x	
Subsampling-by-16[+]	x	x	x
Subsampling-by-32	x	x	
Wavelets			
One iteration	x		
Two iterations	x		
Three iterations	x	x	
Four iterations[++]	x	x	x
Five iterations	x	x	
3 x 3 Convolution		x	
Hybrid			
SS-by-2, 3 wavelets iterations[+++]	x	x	
SS-by-4, 2 wavelets iterations[+++]		x	
SS-by-8, 1 wavelets iteration			x

+The current algorithm for browse
++The candidate wavelets algorithm for browse
+++The major candidate hybrid algorithms for browse

Landsat Browse Study - Phase 1

The Phase 1 study used a 5965-by-6967 pixel scene[***] that included an urban setting having many linear features. RGB false color images were generated using Bands 5, 4, and 3 of the seven spectral bands obtained from the Landsat Thematic Mapper. Ten RGB images were generated, five by iterated subsampling-by-2 and five by iterated wavelets. The SCC was evaluated for each color (band) of each image.

As expected, the SCC tended to decrease with each application of subsampling-by-2 and with each application of wavelets to the Urban P-data scene. Figure 1 illustrates dramatically different behavior of the SCC

[***] All Landsat P data and raw data used in this study were supplied by Stuart Doescher of the U. S. Geological Survey (USGS) EROS Data Center in Sioux Falls, SD.

when iterated wavelets is compared to iterated subsampling-by-2. Data is shown for only one spectral band (Band 3) because band-to-band variation in the SCC was negligible in every case. With each iteration of subsampling-by-2, the SCC drops about 0.08 until, with the fifth iteration, the SCC falls below 0.60, the "threshold of intelligibility". By comparison, a single application of wavelets induces a loss of about 0.04. The next four applications of wavelets together result in an additional loss of about 0.03. (The slight increase in the SCC for lower resolution wavelets, though negligible, is an artifact of the crudeness of the SCC as a measure of image quality.) Consequently, after 5 iterations of wavelets the SCC is approximately that of one iteration of subsampling-by-2, while the SCC of the image resulting from sub-sampling-by-32 (5 iterations of subsampling-by-2) suggests a severely degraded image.
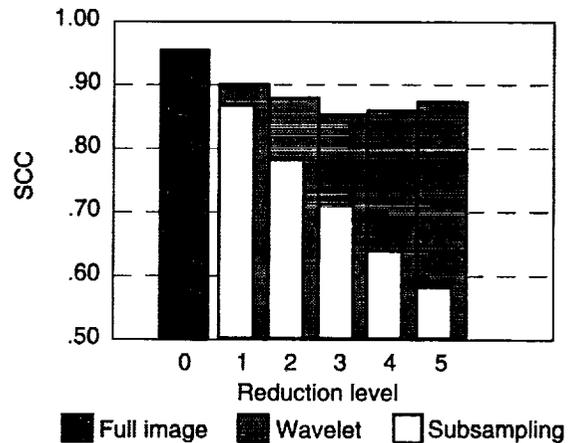

Figure 1. Comparison of Sequential Correlation Coefficients For Subsampling Band 3

The Phase 1 results showed that the wavelets approach is a good alternative to the present subsampling technique. Wavelets-generated imagery retained more

features at reduced resolution and had fewer artifacts: in particular, linear features were never obliterated.

Landsat Browse Study - Phase 2

Phase 2 of the study used Landsat raw data to examine the robustness of wavelets in conserving image quality. This phase also addressed computation time issues. A conventional 3-by-3 convolution [Pratt, 1991, p. 303] was tested as a foundation for a browse capability (see Section 2.3). Milstein's *hybrid-1* technique was also investigated. This method, consisting of subsampling-2 followed by iterated wavelets, was expected to reduce processing time by 75%, compared to iterated wavelets alone.

The data used for this study consisted of Band 5, 4, and 3 raw data for two 5984-by-6400 scenes: a forested mountain scene and a scene consisting of clouds, water, and vegetation. As in Phase 1, five levels of reduction were applied to each scene, using each of the four algorithms. The resulting SCC values are shown in Figure 2 (the forested mountain scene) and Figure 3 (the clouds, water, and vegetation scene).

Compared to Phase 1 results, the degradation represented by the decline in the SCC for the two raw images is slightly greater for wavelets and significantly greater for subsampling, and there is noticeable variation from band to band. This is seen in Figures 4 and 5, which compare SCCs of the 16-to-1 reduced resolution images generated by the four algorithms. Otherwise, SCC findings for subsampling and wavelets do not differ very much from those of the Phase 1: the rapid degradation that occurs for subsampling greatly exceeds that of wavelets.

As suggested by Figures 2 through 5, the performance of the 3-by-3 convolution as a resolution-reducing technique was only marginally better than subsampling. However, the SCCs for hybrid-1 resolution reduction were nearly identical to their pure wavelets counterparts. This unexpected finding suggested hybrid-1 as a viable, high speed alternative to wavelets.

Phase 2 analysis also addressed the question of the relative sensitivity of the browse image to the uncorrected distortions in the raw image under the various algorithms. It was found that neither the wavelets algorithms nor the hybrid algorithms propagated the geometric or radiometric errors for any level of resolution. Both wavelets and hybrid methods proved robust, in particular, when applied to raw image data or to uncorrelated data. This finding dispelled concern for possible error propagation.

These algorithms were implemented by approximately 160 lines of C code. The runs on a Sun SPARC 10 Workstation showed that the run-time performance of all the algorithms meets Landsat 3-minute time constraint. For 16-to-1 resolution reduction, subsampling was by far the fastest algorithm (½ second for non-computational processing), followed by convolution and hybrid-1 (30 seconds), and wavelets (180 seconds).

Landsat Browse Study - Phase 3

Phase 3 investigated two additional algorithms, each computationally faster than wavelets and more effective than subsampling. Phase 3 used a wide variety of full resolution Landsat imagery, in addition to still higher resolution aerial imagery. The
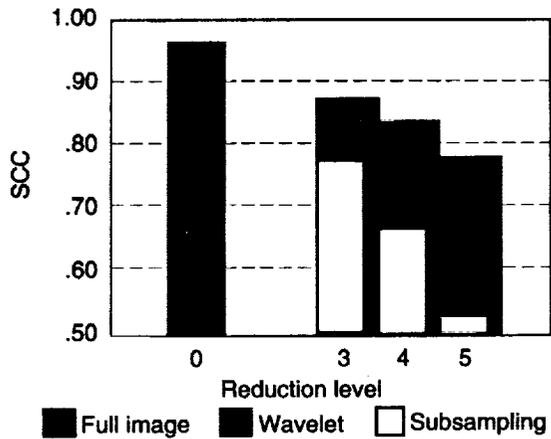
Figure 2. Comparison of SCCs for Subsampling and
Wavelets Generated Imagery Band 3,
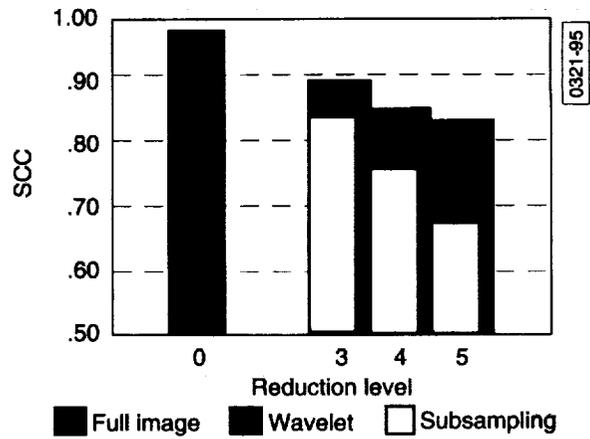Forested Mountain Scene, Raw Data

Figure 3. Comparison of SCCs for Subsampling and
Wavelets Generated Imagery Band 3,
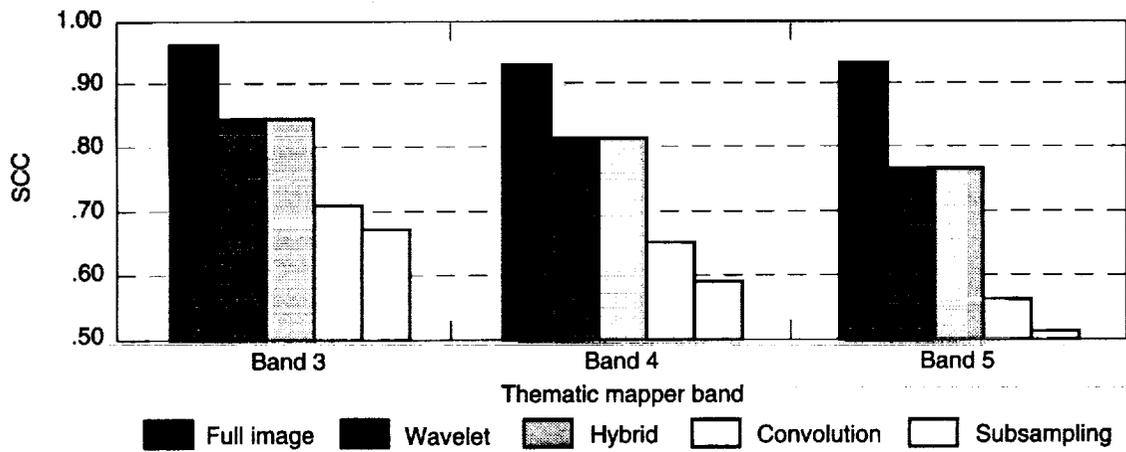Clouds, Water & Vegetation Scene, Raw Data

Figure 4. Comparison of SCCs for Four 16-to-1 Resolution Reduction Algorithms: Band 3,
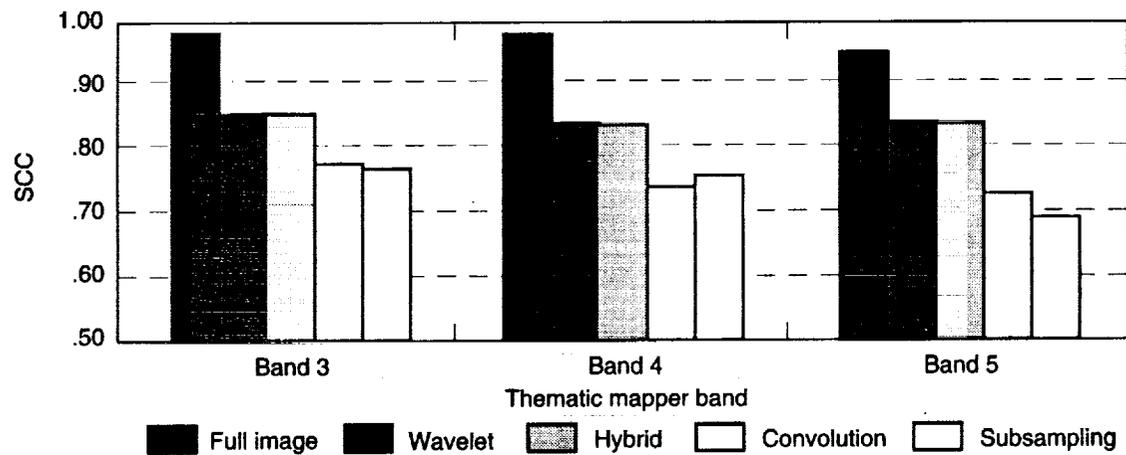Forested Mountain Scene, Raw Data

Figure 5. Comparison of SCCs for Four 16-to-1 Resolution Reduction Algorithms: Band 3,
Clouds, Water & Vegetation Scene, Raw Data

aerial imagery is shown in reduced resolution and discussed in Section 4.

The major issues treated in Phase 3 were the investigation of two more hybrid algorithms and the comparison and evaluation by experts from the scientific community of 16-to-1 reduced imagery generated by alternative algorithms: subsampling-by-16, wavelets, hybrid-1, *hybrid-2* (subsampling-by-4, followed by two iterations of wavelets), and *hybrid-3* (subsampling-by-8, followed by one iteration of wavelets). Compared to iterated wavelets, hybrid-2 and hybrid-3 reduce the number of computations by factors of 16 and 64, respectively.

Experts found that 16-to-1 reduced resolution imagery produced by wavelets, hybrid-1, and hybrid-2 were virtually indistinguishable from one another, though slightly superior to hybrid-3 imagery. All were found far superior to imagery produced by subsampling-by-16. Experts considered imagery produced by wavelets and the three hybrid techniques useful for various purposes, but agreed that imagery produced by subsampling-by-16 had little value except for cloud determination.

This three-phase study established that a Landsat browse product based on either wavelets or a hybrid methodology offers a significantly better quality browse product within the Landsat processing time requirements than the current subsampling-based system. The new techniques produce more trustworthy imagery which can be stored and transmitted efficiently. Roads, communication lines, power lines, rivers, and other linear features are much better preserved by wavelets and the hybrid algorithms, and there are seldom artifacts.

Furthermore, these new methodologies provide greater flexibility, with the potential to meet future image reduction requirements arising from higher resolution imagery created by new sensor technologies.

4 Examination of Gray Scale Images

We now discuss a few reduced resolution aerial images used in the final phase of browse study. Figures 6, 7 and 8 show 16-to-1 reduced resolution versions of an aerial scene. This scene of an Air Force base, includes many roads and paths, a small runway, hills, and so forth.

Figure 6 shows the image after applying four iterations of wavelets to the full resolution image. All roads are clearly discernible, although there is some fade in-and-out or striation along the principal roads. Detailed hillside contour and erosion patterns are visible. It seems possible to make out much of the detail within the base itself. The SCC of this image is 0.90, compared to the full resolution image SCC of 0.98.

Now we examine Figure 7, the same resolution image, produced via subsampling-by-16. The road along the left edge of the military base has become a series of short, disjoint vertical segments, not much different in shape or intensity from horizontal segments just to their right. The same problem exists to varying degrees along most roads. Although the original image was virtually free of noise, the subsampled version has taken on a very noisy appearance, especially within the base area, where small features could assume the greatest importance to the user. This same "pseudo-noise" has washed out much of the
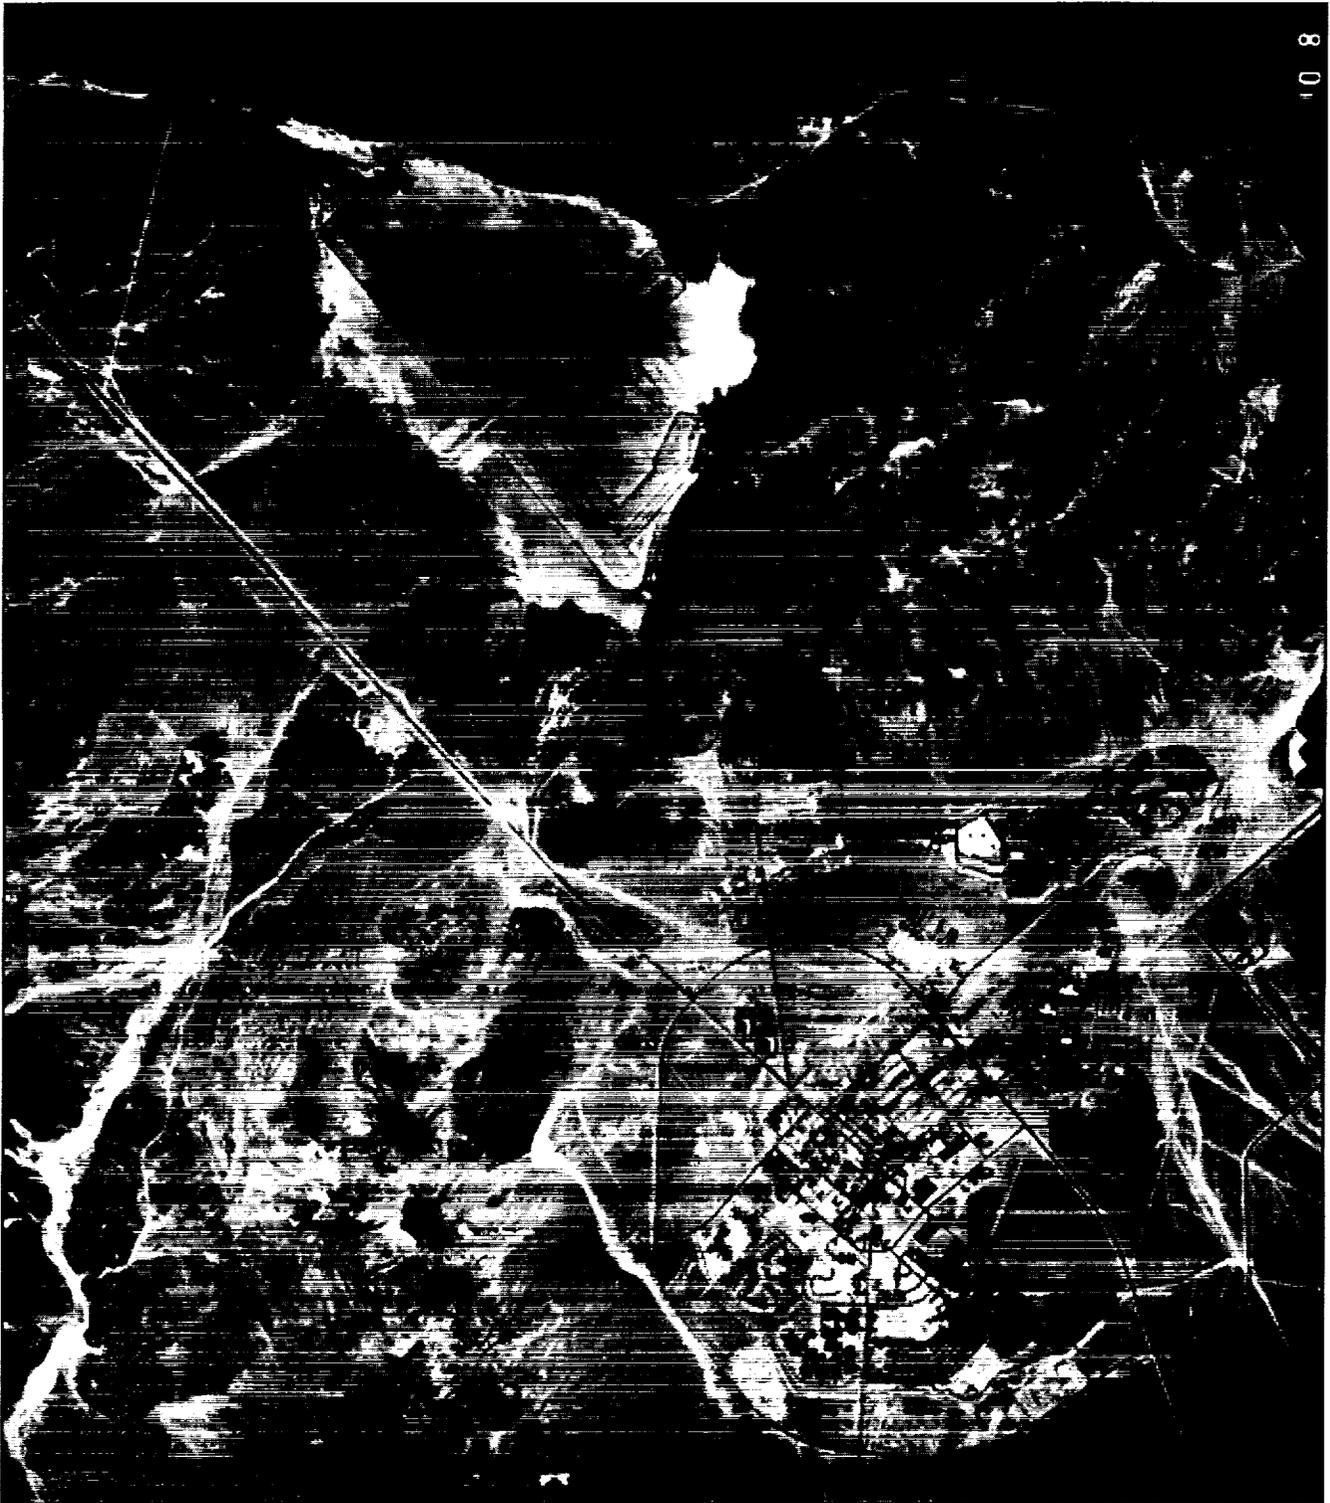
Figure 6. Reduced Resolution Airbase Image After Four
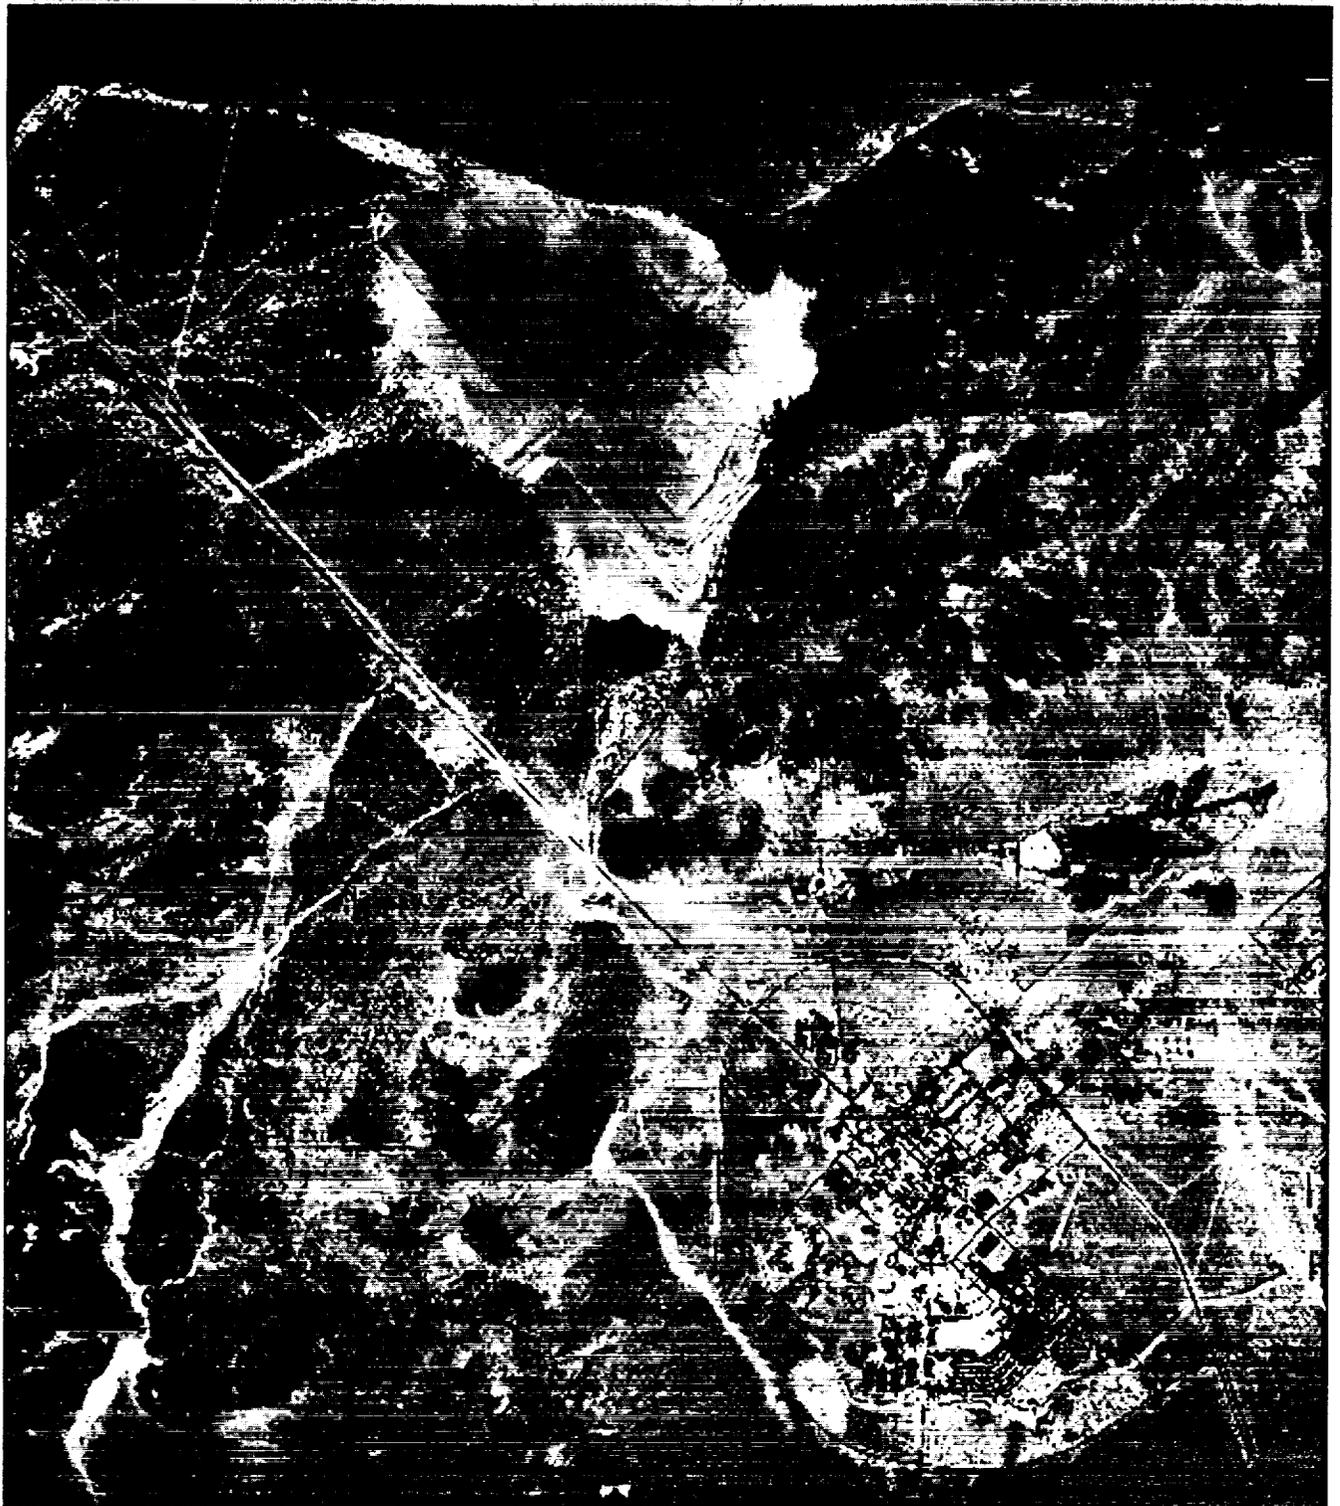Iterations of Wavelets

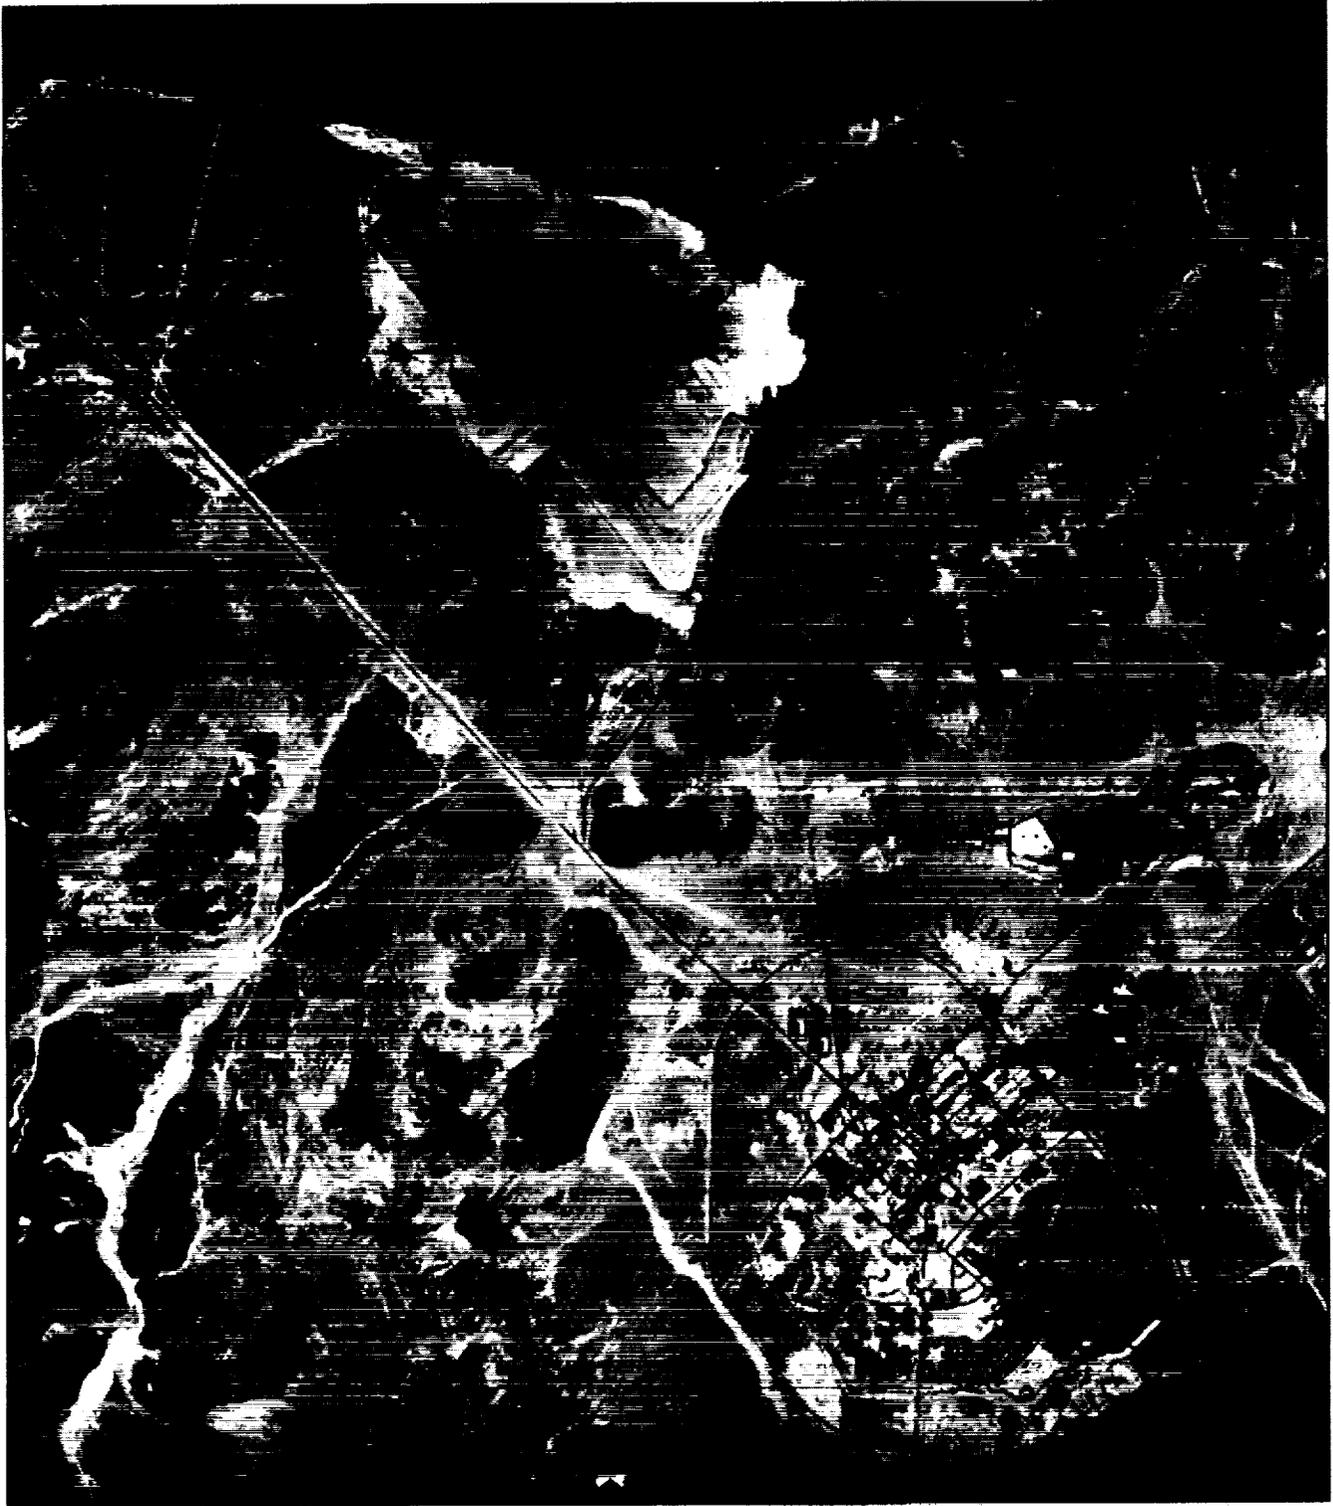Figure 7. Reduced Resolution Airbase Image After Subsampling-by-16

Figure 8. Hybrid-3 Reduced Resolution Airbase Image

topographical hillside detail found in Figure 6. If there had been significant random noise in the original image, the subsampled image would have been much more seriously degraded. The SCC of this image is 0.77.

In Figure 6, on the periphery of the base, about 3 inches from the left and 2 inches from the bottom of the image, is a small array of white objects. Even if we cannot identify this feature, we can use it as an aid in comparing the images. In Figure 7, we see that this feature is distorted beyond recognition (not surprising in view of the discussion in Section 2.2).

Figure 8 shows the effect of the hybrid-3 algorithm: subsampling-by-8, followed by wavelets. Under close scrutiny, we see slight but definite degradation, compared with the iterated wavelets image (Figure 6). For example, the small array is still visible, but the viewer is less certain as to its boundary. Yet, overall image quality seems much closer to pure wavelets than to pure subsampling. In fact, the SCC of this image is 0.89, compared to 0.90 SCC value for the wavelets image. In view of the visual quality of the hybrid-3 image, and the processing speed of the hybrid-3 algorithm (64 times that of wavelets) this algorithm could be an attractive alternative to iterated wavelets when computational speed is important.

The hybrid-1 and hybrid-2 images are not reproduced here. The hybrid-1 image appears visually indistinguishable from the pure wavelets image. The hybrid-2 image is distinguishable from the pure wavelets image but only in the finest of visible detail. There is no significant difference in the SCC values for wavelets, hybrid-1, and hybrid-2.

The quality of the hybrid-1 and hybrid-2 products, together with their processing speed-ups (respectively 4-to-1 and 16-to-1) compared to that of iterated wavelets, again make them serious alternatives to iterated wavelets in many applications. As a group, iterated wavelets, hybrid-1, hybrid-2, hybrid-3 constitute a prepackaged trade-off set of algorithms, which could give the user the luxury of choosing his or her own speed-quality trade-off.

5 Summary and Applications

The resolution of an image is the distance required between small objects in order to distinguish them from one another. In satellite imaging systems it is often desirable to generate reduced resolution versions of satellite imagery. Some deterioration in the visual quality of the imagery inevitably results from this process. However, some resolution-reducing algorithms are more effective than others in preserving the visual quality of the original image. We noted that a resolution reducing algorithm that does a good job in retaining visual quality has many potential applications to satellite imaging systems.

We recounted a study in which a variety of resolution-reducing algorithms were investigated in an effort to provide a superior browse product for Landsat imagery. Using a crude quantitative measure, we compared the current technique, subsampling-by-16, to a resolution-reducing technique based on a conventional convolution, an iterated wavelets-based algorithm, and several hybrid algorithms involving subsampling followed by iterated wavelets. Comparing images of the same resolution, those produced by iterated wavelets had quantitative measures

superior to those resulting from convolution and still more so from subsampling. The hybrid algorithms ranged from faster, with imagery visually indistinguishable from that of iterated wavelets, to much faster, with imagery of slightly lower quality than that produced by iterated wavelets. Imagery produced by pure subsampling was distinctly inferior compared to that of wavelets or any of the hybrid algorithms. Visual inspection by experts confirmed the findings suggested by the quantitative measure. Each of these algorithms can support resolution reductions of 2^k-to-1, $k \geq j$ ($j = 0$ for iterated wavelets, i for hybrid-i, $i = 1, 2$ or 3). The new algorithms were validated using the full variety of Landsat TM data, both P data and raw data, as well as higher resolution aerial imagery. All ran fast enough to satisfy browse requirements.

The wavelets-hybrid set of algorithms provide a speed-selectable set of 2^k-to-1 resolution reduction algorithms ($k = 0,1,...$) applicable to a variety of imaging satellite system problems, including the Landsat display problem, the downlink browse problem, emergency spaceborne image communication, and real-time video downlinking, in addition to the Landsat browse problem.

Dedication

The authors dedicate this paper in honor of U. C. Berkeley Professor Hans J. Bremmermann, whose work on Haar functions pioneered early developments in the theory and application of wavelets.

References

Daubechies, I. (1988). Othonormal Bases of Compactly Supported Wavelets. *Comm. Pure Appl. Math., Vol. 41, pp* 909-996.

Mallat, S. G. (1989). Multiresolution Approximations and Wavelet Orthonormal Bases of L^2(R). *Trans. Amer. Math. Soc., Vol. 315, pp* 69-87.

Meyer, Y. (1986). Ondelettes et Fonctions Splines. Presented at the *Seminare Equations aux Derivees Partielles*, Ecole Polytechnique, Paris.

Muller, J-P, ed. (1988). *Digital Image Processing in Remote Sensing.* Taylor and Francis, Inc., Philadelphia.

Pratt, W. K. (1991). *Digital Image Processing, Second Edition, p* 301. John Wiley and Sons, NY.

Rosenfeld, A. and Kak, A. C. (1982). *Digital Picture Processing, Vol. 1, Second Edition.* Academic Press, Inc., NY.

Strang, G. (1989). Wavelets and Dilation Equations: a Brief Introduction. *SIAM Review, Vol. 31, No. 4, pp* 614-627.