J.M. Ogden | E.H. Adelson | J R. Bergen | P.J. Burt

Pyramid-based computer graphics

Pyramid-based graphics techniques can provide realistic computer graphics on small systems without the complexities of physics simulations.

Human beings have an intuitive feel for graphics. Graphics problems such as blending two images smoothly, interpolating to fill in missing image data, or creating realistic looking images are routinely solved by artists using traditional media. It is much more difficult to perform these tasks mathematically. A major step in solving a computer graphics problem is choosing an appropriate numerical representation for the image, one which allows us to use our visual and artistic intuition in a natural way.

Artists often tend to separate the spatial scales of an image when creating or altering a picture. When an artist paints a landscape, the coarse scale (low spatial frequency) information is filled in first, as a wash of color in a large region. Intermediate-size details can be added next with a medium-size brush. As a last step, the artist draws the fine details with a small brush. The image can be considered a sum of overlays of increasingly fine detail. When an artist touches up a damaged picture, both the large and small scale variations are considered in filling in the missing pieces. In synthesizing and combining images, we would like to imitate the artist's ability to see an image on both large and small scales by representing the image mathematically on many different spatial scales simultaneously.

The simplest way to represent an image is to set the numerical value of each pixel proportional to the image intensity. This representation is useful when we want to paint or draw directly one pixel at a time, or when simple manipulations of contrast or color are desired. But it becomes cumbersome when we want to look at an image at several spatial resolutions.

Abstract: *This paper describes pyramid solutions to graphics problems that have proven difficult in other image representations. The "physics simulation" approach grows more out of the physics and mathematical modelling traditions. Greater realism can be achieved by using the physics simulation approach but the complexity and computation time are vastly increased over the multiresolution pyramid approaches described here.*

Final manuscript received October 21, 1985
Reprint RE-30-5-1

A Fourier transform representation can be used to separate the various spatial scales of an image. Unfortunately, when we leave the familiar spatial domain for the spatial frequency domain our intuitive feel for the problem is lost. Operating on the Fourier transform of an image, we can no longer "see" local spatial features in a recognizable form. What is really needed is a representation that describes an image at multiple spatial resolutions, and also preserves the local spatial structure that allows us to "see" the picture at each scale. Pyramid representations are ideal for this class of problems.[1,2,3]

The pyramid representation and computer graphics

The pyramid representation expresses an image as a sum of spatially bandpassed images while retaining local spatial information in each band. A pyramid is created by lowpass-filtering an image G_0 with a compact two-dimensional filter. The filtered image is then subsampled by removing every other pixel and every other row to obtain a reduced image G_1. This process is repeated to form a Gaussian pyramid G_0, G_1, G_2, G_3 ... G_n (Fig. 1).

$$G_k\,(i,j)= \sum_m \sum_n G_{k-1}(2i + m,\, 2 + n), k = 1,N$$

Expanding G_1 to the same size as G_0 and subtracting yields the bandpassed image L_0. A Laplacian pyramid L_o, L_1, L_2, ... L_{n-1}, can be built containing bandpassed images of decreasing size and spatial frequency.

$$L_k = G_k - G_{k+1}, k = 0,N\text{-}1$$

where the expanded image $G_{k,1}$ is given by

$$G_{k,l}\,(i,j) = 4 \sum_m \sum_n G_{k,l-1}[(2i + m/2,\, 2j + n)/2)]f(m,n)$$

The original image can be reconstructed from the expanded bandpass images:

$$G_0 = L_0 + L_{1,1}, + L_{2,2} + ...L_{N-1,N-1} + G_{N,N}$$

Figure 2 shows an image represented as a sum of several spatial frequency bands.

The Gaussian pyramid contains lowpassed versions of the original G_0, at progressively lower spatial frequencies. This effect is clearly seen when the Gaussian pyramid "levels" are expanded to the same size as G_0 (Fig. 3a). The Laplacian pyramid consists of bandpassed copies of G_0. Each Laplacian level contains the "edges" of a certain size, and spans approximately an octave in spatial frequency (Fig. 3b).

This pyramid representation is useful for two important classes of computer graphics problems. First, tasks that involve analysis of existing images, such as merging images or interpolating to fill in missing data smoothly, become much more intuitive when we can manipulate easily visible local image features at several spatial resolutions. And second, when we are synthesizing images, the pyramid becomes a multiresolution sketch pad. We can fill in the local spatial information at increasingly fine detail (as an artist does when painting) by specifying successive levels of a pyramid.

In this paper, we will describe pyramid solutions to some graphics problems that have proven difficult in other image representations:

1. Image analysis problems
   (a) Interpolation to fill missing pieces of an image
   (b) Smooth merging of several images to form mosaics
2. Creation of realistic looking images
   (a) Shadours and shading
   (b) Fast generation of natural looking textures and scenes using fractals
   (c) Real-time animation of fractals.

Multiresolution interpolation and extrapolation

The problem

The need to interpolate missing image data in a smooth, natural way arises in a number of contexts. It can be used to remove spots and scratches from photographs, to fill in transmitted images that are incomplete, and to create interesting computer graphic effects.

A solution

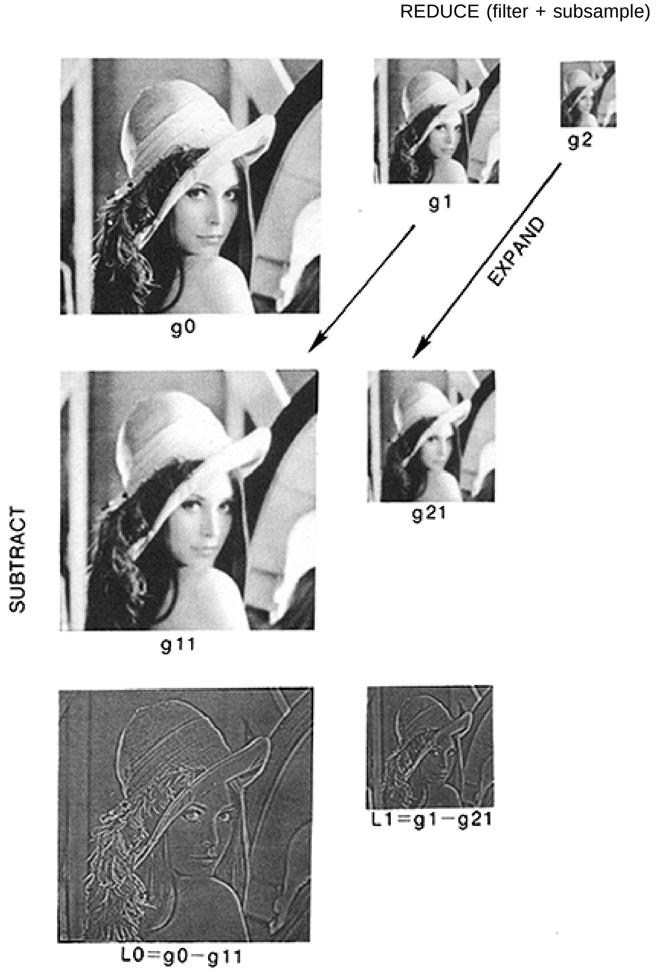

Fig. 1. *Building the pyramid.*

Insight into the problem of interpolation is gained by considering the image as a sum of patterns of many scales. A typical photograph includes small scale fluctuations due to surface texture, superimposed on more gradual changes due to surface curvature or illumination variations. Similarly, a painting is a composite made up of features of many scales, rendered with brushes of different sizes. In predicting the values of the missing pieces of an image, we need to consider intensity variations on both small and large scales.

Figure 4 shows a one-dimensional representation of an image G_0 that has some missing values. One method of interpolating to find the unknown region is to use a Taylor series expansion. If the size of the missing piece is comparable to the finest scale features of the image, a good estimate of the unknown value $x+dx$ can be obtained through a simple linear prediction based on the first derivative G_0' at the point x:

$$G_0 (x + dx) = G_0 (x) + dx\, G_0'(x).$$

If the missing piece is large compared to the fine scale features of

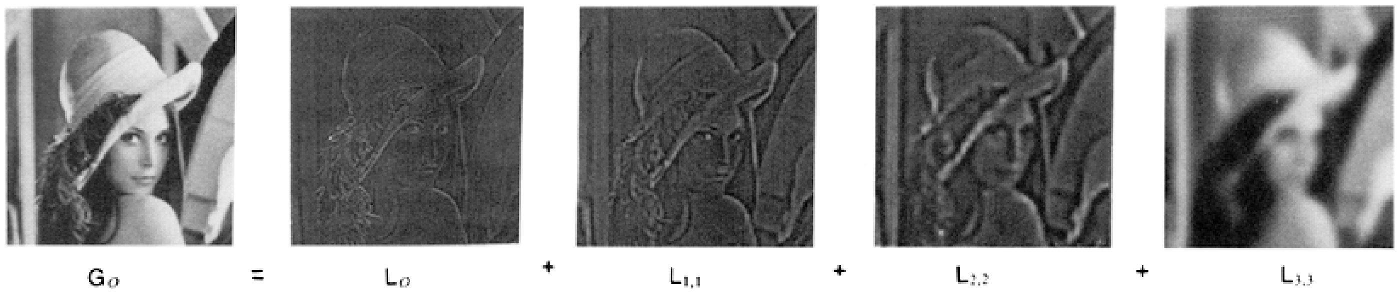

Fig. 2. *The pyramid as a sum of spatial frequency bands.*

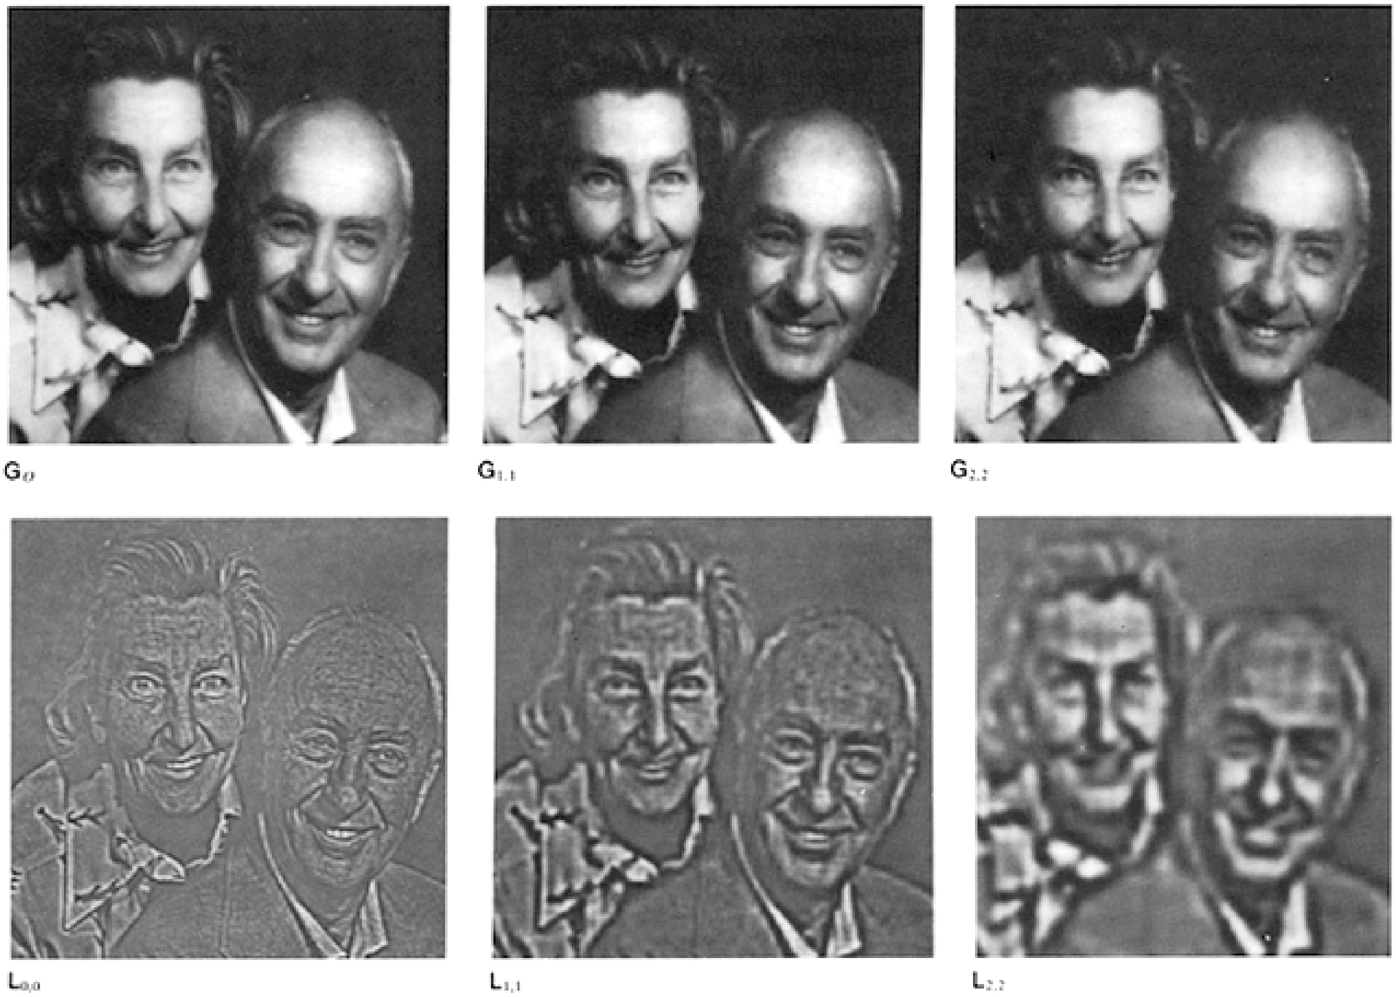

Fig. 3. *Expanded Gaussian and Laplacian pyramids.*

the image, we need to examine variations on larger scales as well. We can compute such an estimate by fitting the function $G_0(x)$ to a Taylor polynominal of higher degree. This involves taking higher order derivatives that represent the image variation over a larger number of pixels. One disadvantage of this approach for computer graphics is that it is computationally expensive. It is also difficult to adjust the degree of the interpolating polynomial to account for missing regions of different sizes.

An alternative way to look at the missing information on many scales is to build a Gaussian pyramid. This represents the image G_0 at different spatial resolutions ranging from fine (G_0) to coarse (G_n). The unknown piece of G_0 is also missing from G_1, G_2, ... G_n. Note that the size of the missing region is reduced in the reduced pyramid levels. Now, instead of fitting a Taylor polynomial of high degree at the finest spatial scale, we use linear interpolation at multiple spatial resolutions to fill in the missing information.

In the example shown, the size of the missing piece is large compared to the fine scale variations contained in G_0 and G_1, but is comparable to the feature size in G_{n-1} and small compared to the coarse features of G_n. Starting with G_0, we use linear extrapolation to predict the values of unknown points with two known neighbors.

$$G_0\,(i) = 2\,G_0\,(i+1) - G_0\,(i+2)$$

For this example, most of the unknown values of G_0 are "left blank." A small border one pixel wide is extrapolated into the unknown region. Now we build G_1 and extrapolate again. At this reduced resolution, the extrapolated border corresponds to a larger proportion of the unknown region. When the extrapolated G_1 is expanded to full size, the border also expands in size from one pixel to several pixels. Continuing to lower spatial resolutions, we eventually reach a reduced pyramid level G_n, where the unknown region has shrunk to only one pixel. Extrapolation at G_n gives a pyramid level with all the values filled in.

An extrapolated image is built by reconstructing the extrapolated pyramid. Starting with G_n, we expand to form $G_{n,1}$. Where a pixel in the next highest frequency band G_{n-1} is missing, the value from $G_{n,1}$ is used. Continuing this process, we form an extrapolated image G_0, with all unknown points filled in.

Examples

Figure 5a shows a portrait that has had ink spilled on it. The locations of the ink spots are indicated in a mask image, Fig. 5b. When a two-dimensional multiresolution interpolation procedure is applied, the missing image points are filled in smoothly, as shown in Fig. 5c. In many cases, the result is so natural looking that the flaw would not be detected except on close

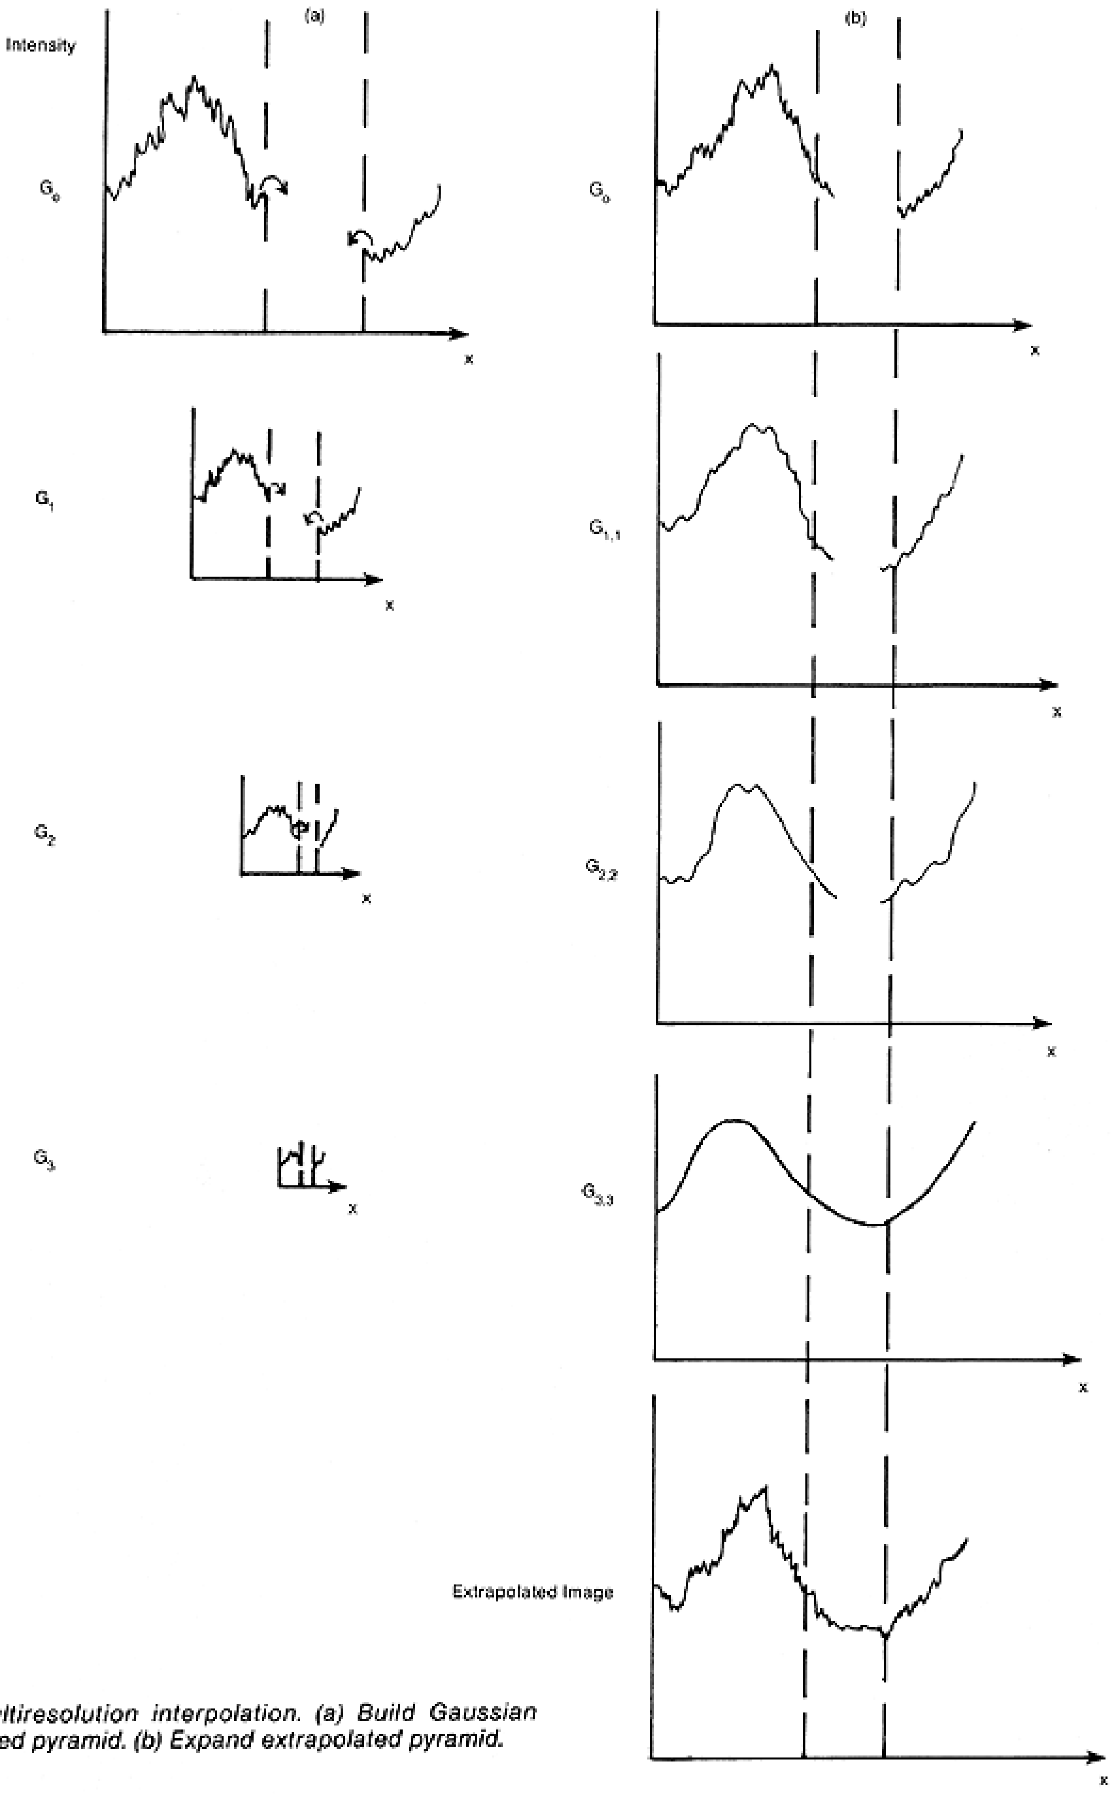

Fig. 4. *Multiresolution interpolation. (a) Build Gaussian extrapolated pyramid. (b) Expand extrapolated pyramid.*

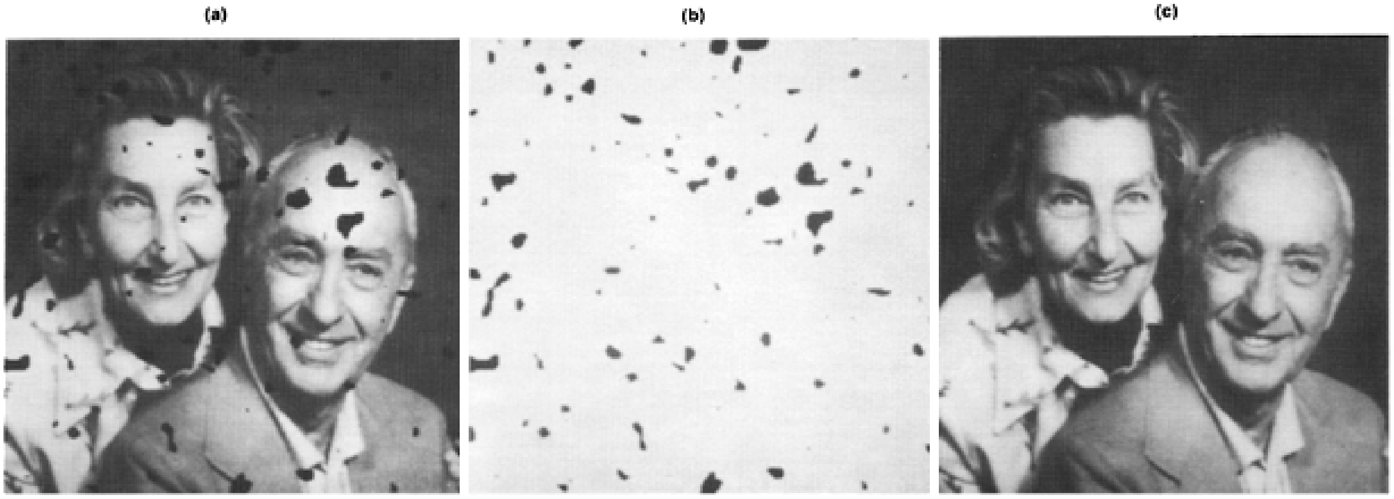

Fig. 5. *Interpolation to fill missing points in an image.*

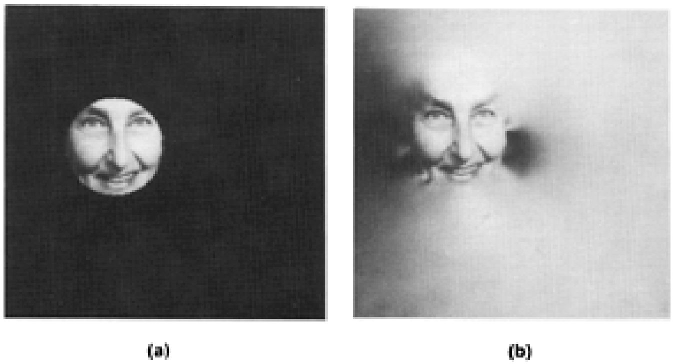

Fig. 6. *Extrapolation example.*

examination. Figure 6 illustrates extrapolation when known image points represent only a small island within the image domain.

Image merging

The problem

It is frequently desirable to combine several source images into a larger composite. Collages made up of multiple images are often found in art and advertising, as well as in science (for example, NASA's mosaic images of the planets). images can even be combined to extend such properties as depth-of-field and dynamic range.

The essential problem in image merging may be stated as one of "pattern conservation." Important details of the component images must be preserved in the composite, while no spurious pattern elements are introduced by the merging process. Simple approaches to merging often create visible edge artifacts between regions taken from different source images.

To illustrate the problems encountered in image merging, suppose we wish to construct a mosaic consisting of the left half of an apple image, Fig. 7a, and the right half of an orange, Fig. 7b. The most direct procedure is to simply join these images along their center lines. However this results in a clearly visible step edge (Fig. 7c).

An alternative approach is to join image components smoothly by averaging pixel values within a transition zone centered on the join line.[3,4] The width of the transition zone is then a critical parameter of the merging process. If it is too narrow, the transition will still be visible as a somewhat blurred step. If it is too wide, features from both images will be visible within the transition zone as in a photographic double exposure. The blurred-edge effect is due to a mismatch of low frequencies along the mosaic boundary, while the double-exposure effect is due to a mismatch in high frequencies. In general there is no choice of transition zone width that can avoid both artifacts.

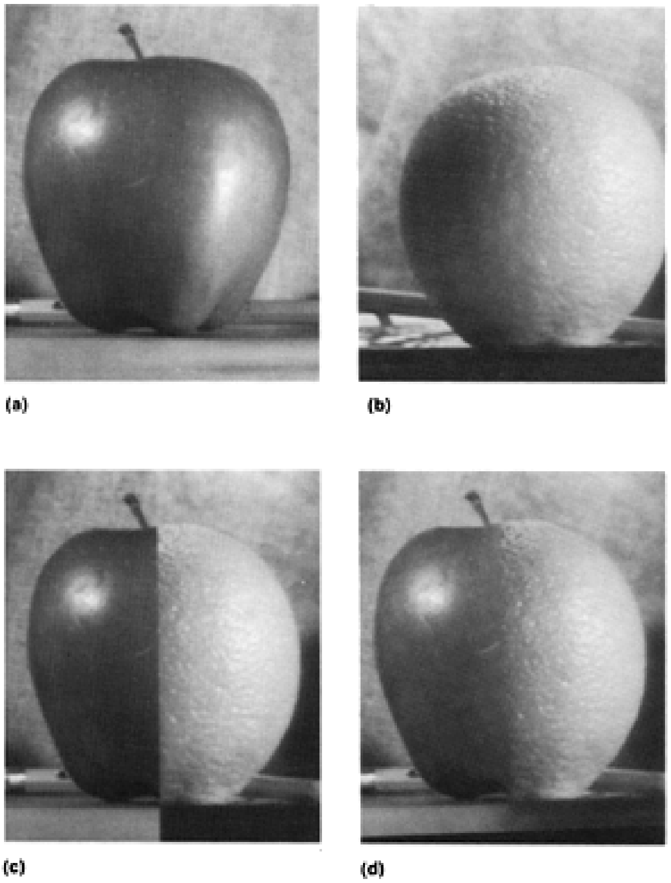

Fig. 7. *Multi-resolution spline of apple an orange. (a) Apple, (b) orange, (c) cut and paste composite (d) multi-resolution pyramid mosaic.*

Multiresolution spline

We can resolve the transition zone dilemma if the images are decomposed into a set of bandpass components before they are merged. A wide transition zone can then be used for the low frequency components, while a narrow zone is used for the high frequency components. In order to have smooth blending, the width of the transition zone in a given band should be about one wavelength of the band's central frequency. The merged bandpass components are then recombined to obtain the final image mosaic.

Let S_0, S_1, S_2, ... S_k be a set of K source images. A set of binary mask images M_o, M_1, M_2, ... M_k determine how the source images should be combined. M_k is "1" where source image S_k is valid, and "0" elsewhere. Simply multiplying S_k 5 M_k and summing over k would give a "cut and paste" composite with step edges. Instead, we build a Laplacian pyramid L_{k1} for each source image, and a Gaussian pyramid M_{k1} for each mask image. A composite Laplacian pyramid L_{k1} is formed by "cutting and pasting" at each spatial scale by weighting each source pyramid level by its corresponding mask:

$$L_{cl}\,(i,j)= M_k\,(i,j)\; L_{kl}\,(i,j)$$

The final image is reconstructed from L_C by expanding each level and summing. Smooth blending is achieved because the transition zone in each pyramid level is comparable to a wavelength of the central frequency of that level. When this procedure is applied to the apple and orange images of Fig. 7, an "orple" is obtained with no visible seam (Fig. 7d).

Multifocus

When assembling information from multiple source images we need not always proceed region by region guided by mask images. Some types of information can be merged in the pyramid node by node and be guided by the node's own value. Here we show how this type of merging can be used to extend the depth-of-field of an image or increase its dynamic range.[3]

Figures 8a and 8b show two exposures of a circuit board taken with the camera focused at different depth planes. We wish to construct a composite image in which all the components and the board surface are in focus. Let L_A and L_B be Laplacian pyramids for the two source images. The low frequency levels of these pyramids should have nearly identical values, since changes in focus have little effect on the low frequency components of the image. On the other hand, changes in focus will affect node values in the pyramid levels where high spatial frequency information is encoded. However, corresponding nodes in the two images will generally represent the same feature of the scene, and will differ primarily in attenuation due to blur. The node with the largest amplitude will be in the image that is most nearly in focus. Thus, "in focus" image components can be selected node-by-node in the pyramids rather than region-by-region in the original images. A pyramid L_C is constructed for the composite image by setting each node equal to the corresponding node in L_A or L_B that has the larger absolute value:

If $\qquad L_{Al}\,(i,j) > L_{Bl}\,(i,j)$

then $\qquad L_{Cl}\,(i,j) = L_{Al}\,(i,j)$

else $\qquad L_{Cl}\,(i,j) = L_{Bl}\,(i,j)$

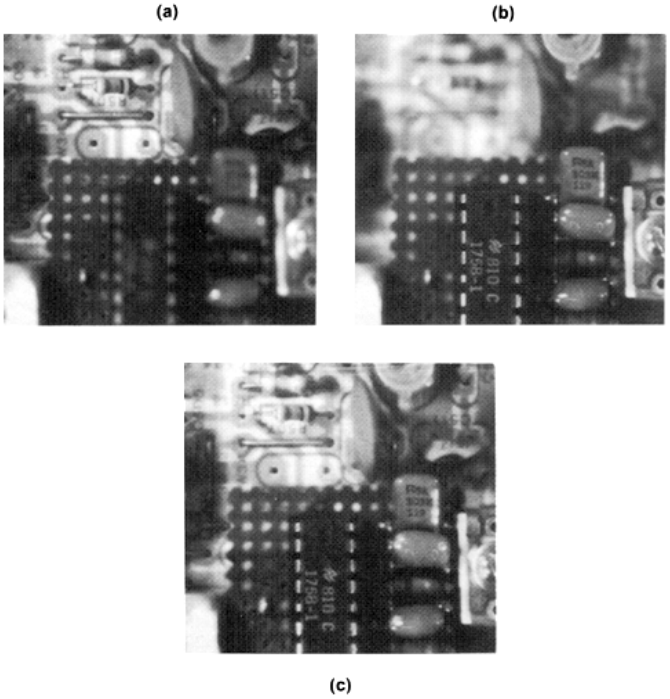

Fig. 8. *Multi-focus composite. (a, b) Two images of the same scene taken with different focuses, and (c) composite with extended depth-of-field.*

The composite image is then obtained simply by expanding and adding the levels of L_C. Figure 8c shows an extended depth-of-field image obtained in this way.

Creating realistic looking images: shadows and shading

The problem

We will consider three different approaches to the problem of creating a realistic looking image. The first approach is to use the computer as a paint box. No mathematical description of the scene is given. All the renderings, shading, shadows, and highlights are done "by hand," as though the artist were using a canvas. The advantage of "paint box" approach is complete artistic control. The disadvantage is that it is time consuming to create an image and difficult to make major changes without redrawing completely.

A second approach is to create a three-dimensional mathematical "universe." The artist specifies the location of objects in this new world, their shapes and physical properties, and the location of light sources. A two-dimensional, photograph-like image of the three-dimensional world is made by tracing a large number of light rays as they are reflected, refracted and absorbed. The advantage of this "physics simulation" approach is that very realistic looking images can be created. This approach is also flexible in that the viewing angle and properties of the component parts can be changed as input parameters. The disadvantage is that it requires a complex physical model and a lot of computation time.

There is a third, "multiresolution," approach that lies somewhere between the first two and combines some of the advantages

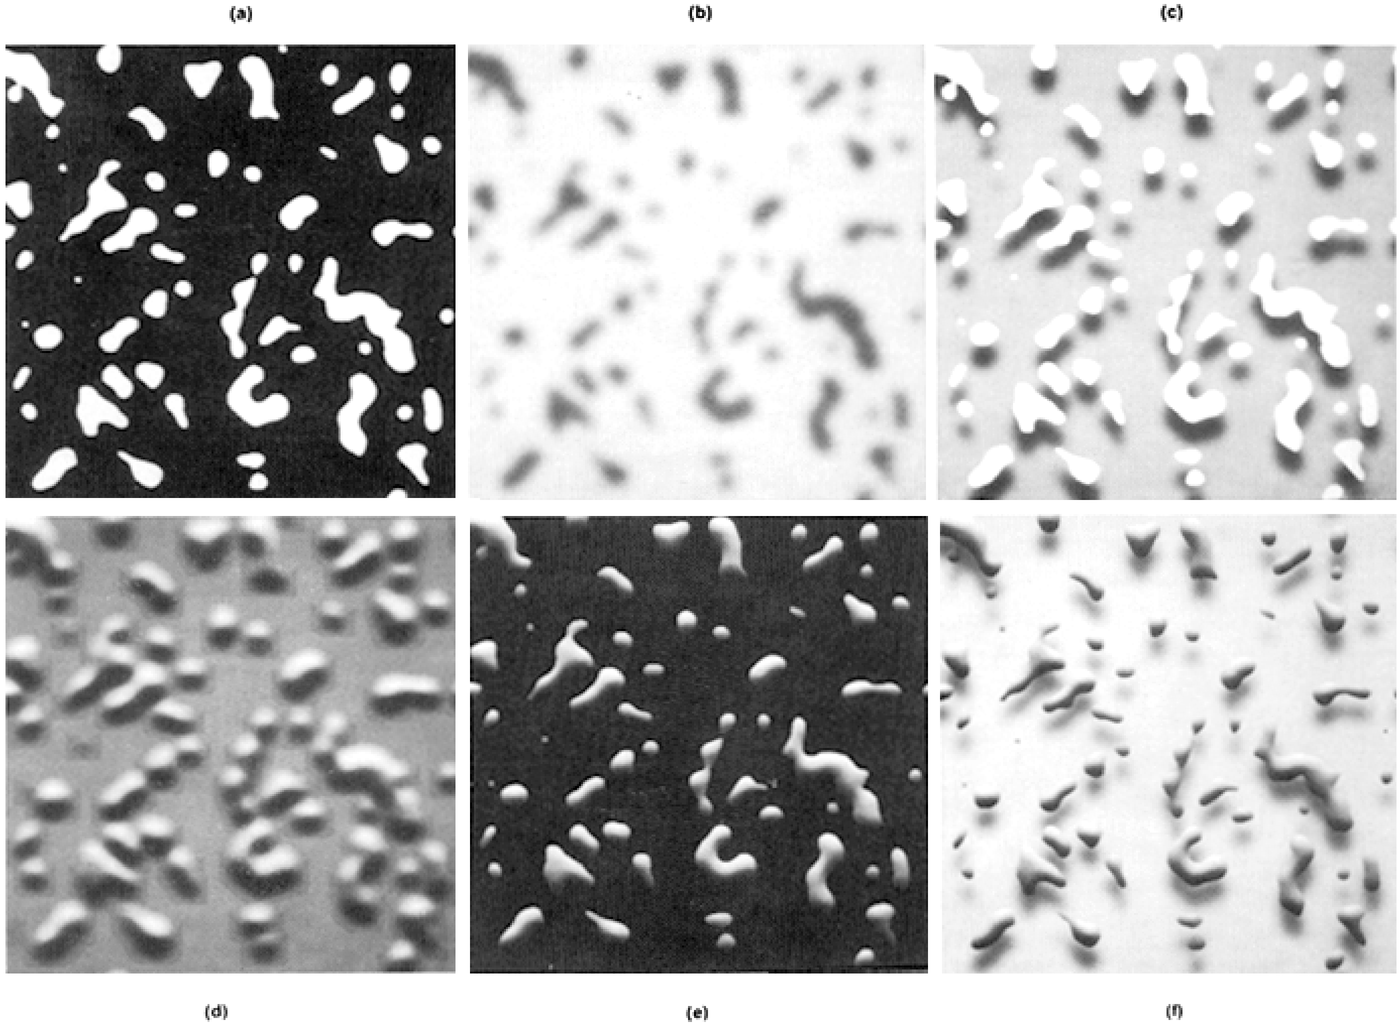

Fig.9. *Multiresolution shadowing and shading of flat shapes.*

of each. As in the paint box approach, we are concerned primarily with painting the two-dimensional image. The difference lies in the type of information on the artist's palette. We usually think of a palette as an array of colors that can be blended and applied to an image. Using the pyramid we can extend the definition of palette to include multiresolution shape and edge information as well as color and intensity. Multiresolution lowpass and bandpass copies of image features are extremely useful in creating special effects and in adding realism to an artifically generated shape. The advantage of this approach is that natural looking images can be generated quickly without resorting to an elaborate physics simulation. Artistic decisions are made by viewing the two-dimensional image and the pyramid levels and combining desired elements of each.

An example

As an example of how an artist might use the pyramid as a spatial frequency palette, consider the problem of making flat shapes (Fig. 9a) appear three-dimensional by adding realistic looking shadows and shading. Both the shadows and shading resemble blurred copies of the original shape. Building a Gaussian pyramid of Fig. 9a, we select a lowpass copy that resembles soft shadows (Fig. 9b). Comparing Fig. 9a and a slightly displaced 9b pixel by pixel, and taking the maximum value at each point, gives an image of shadowed paper cutouts floating or sitting on a glass-topped table (Fig. 9c). Now we need to add dimension to the white cutouts. It has long been known that filtering an image with a "gradient" filter (1, -1) gives an effect of side illumination. When gradient filtering is done at multiple resolutions, there is a great improvement in this bas-relief effect. Performing a combination of gradient and lowpass filtering on Fig. 9a gives us the low-frequency relief 9d. If the minimum of Figs. 9d and 9a is taken, the result is an image of three-dimensional-looking droplets (Fig. 9e). Finally, we add Fig. 9e to 9c to form shadowed droplets in Fig. 9f.

Clearly, many other interesting graphic effects can be generated by imaginative use of the pyramid. It provides the artist with a convenient and efficient way of accessing certain important features of an image—shapes and edges—at multiple resolutions. For certain problems, especially when computation facilities are limited, the multiresolution approach offers considerable realism for little computation time.

Pyramid generation of fractals

The problem

The computer graphics community has adopted fractals as a remarkably effective way of synthesizing natural looking tex-

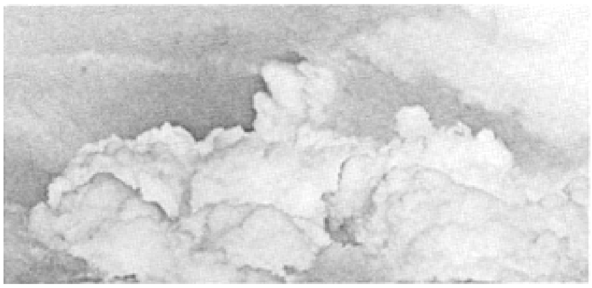

Fig. 10. *Cloud as a sum of random circles.*

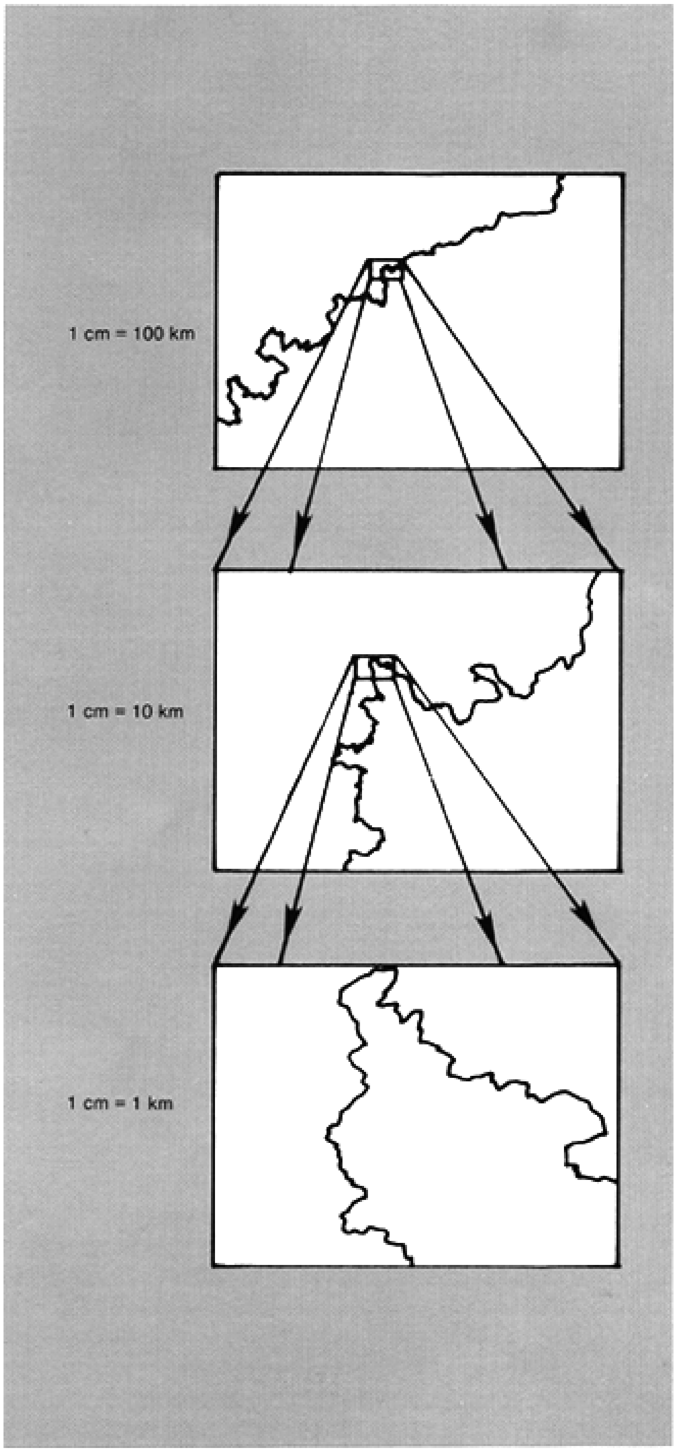

Fig. 11. *Coastline on three different spatial scales.*

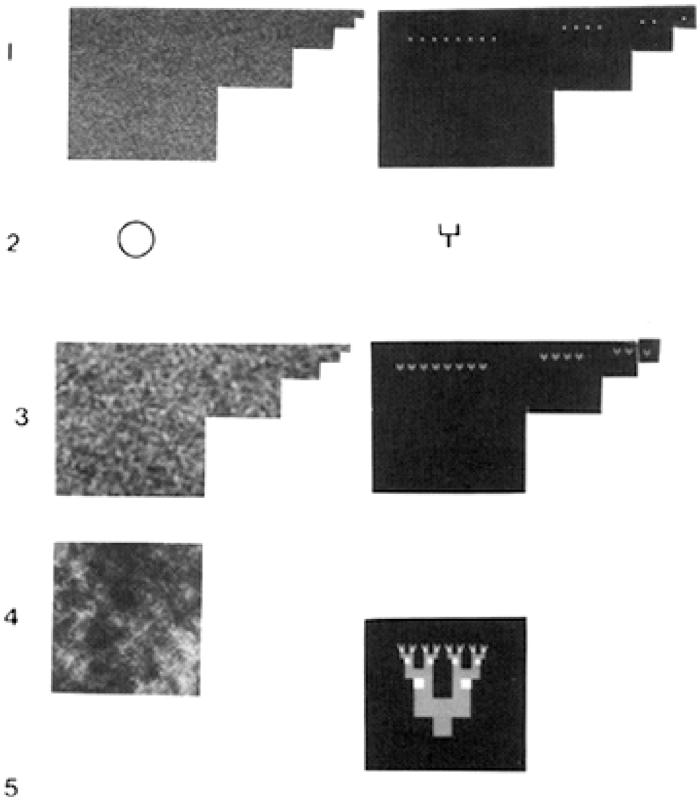

Fig. 12. *Pyramid generation of fractals.*

tures.[5,6,7,8] The problem is to generate these textures quickly using multiresolution techniques.

Fractals and the pyramid

Traditional mathematics has relied on idealized models of the complicated and irregular forms of nature. Structures such as clouds, mountains, and coastlines are difficult to describe in terms of continuous, differentiable functions. Recently, Mandelbrot devised a new class of functions called "fractals" to express these complex natural forms.[5]

A fractal function includes both the basic form inherent in the object and its statistical or random properties. For example, a cloud can be visualized as a sum of random circles. The basic circle shape is seen at randomly distributed sizes and positions (Fig. 10).

Fractals have the property of self-similarity over many different geometric scales. A fractal appears similar as the spatial scale is changed over many orders of magnitude. As an example, consider a very jagged, rocky coastline. The coastline looks qualitatively similar when plotted on different special scales (see Fig. 11).

We have devised a fast fractal generation technique based on the pyramid algorithm, which takes advantage of the self-similarity of fractals. The pyramid breaks an image up into a sum of bandpassed images plus a lowpass filtered image. If an inherently self-similar fractal image is decomposed into pyramid form, one would expect the bandpassed images to look similar at each spatial frequency scale. Conversely, if similar patterns were entered into each spatial band of a pyramid, the reconstructed image should look like a fractal.

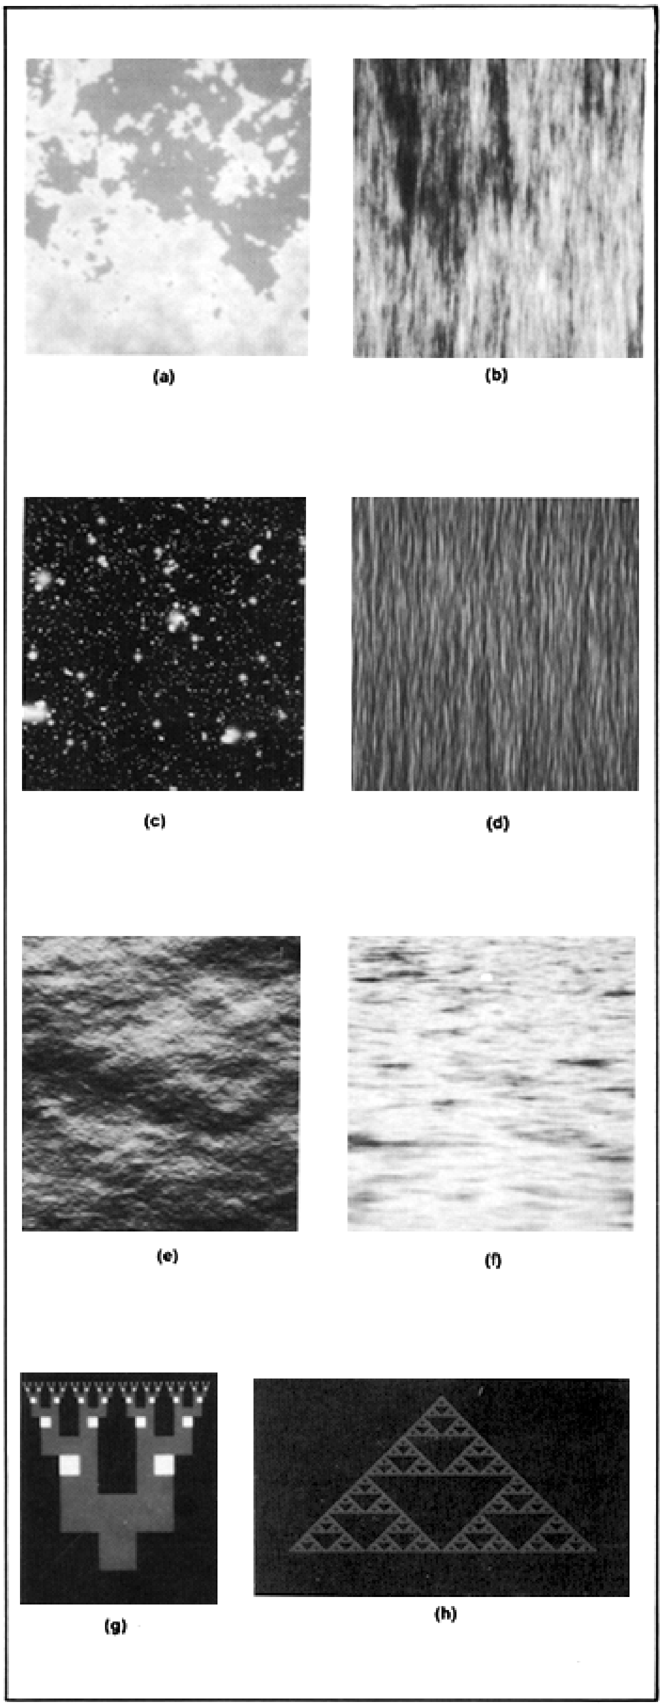

Fig. 13. *Examples of fractals. (a) Clouds, (b) woodgrain, (c) galaxy, (d) papyrus, (e) granite, (f) waves, (g) fractal "tree," and (h) fractal gasket*

Fractal generation method

The method used involves replicating a basic form, or "generator," at a succession of spatial scales, then summing to produce a fractal image (Fig. 12):

1. A "seed image" is set. This can be a regular array or a random pattern of dots.
2. A basic form, or "generator, " is selected. This shape will appear at many spatial scales.
3. The seed image is convolved with a filter having the same shape as the generator. Each dot in the seed image will reproduce as an image of the generator. The result is a sum of generators with amplitude and position determined by the seed image. The convolved image is entered into the various levels of the pyramid.
4. The pyramid is reconstructed by expanding the various levels and adding. Because we have included generators over a wide range of spatial frequencies, the reconstructed image is a self-similar fractal.
5. The fractal image can be threshholded, colored, shaded, filtered, or otherwise manipulated to give pleasing effects.

Some examples of pyramid fractuals are given in Table I and in Fig. 13.

Realtime fractal animation method

A moving fractal image (for example, clouds drifting by, or rippling waves) can be decomposed into a sum of spatial frequency bands. The different spatial frequency components move at different speeds. In an ocean wave, for example, low-frequency swells move more quickly than high-frequency ripples on the water's surface.

An initial fractal image is expressed as a sum of several spatial frequency components. The image is advanced one time frame by moving each spatial frequency component a specified amount and summing again. This creates a new fractal, which is slightly changed from the previous frame. Continuing this process gives an animated sequence of continuously evolving fractal images. Each spatial frequency component can be periodic with a fairly short repeat distance (say, half the image width), but the sum has a much longer period if the different components move at various speeds that are not multiples of each other. We have demonstrated this technique using our DeAnza 1P8500 image processor. Each new frame takes about 1/15th of a second to compute, and the overall repeat time was several minutes. Figure 14 shows an example of fractal animation for a beach scene (also see front cover).

Conclusions

Pyramid representations have much in common with the way people see the world. Human beings like to look at things on many spatial scales simultaneously. A strong analogy exists between an artist who paints images by adding progressively finer details, and a computer scientist who constructs images by adding together the spatial frequency bands of a pyramid. Representing an image on multiple spatial scales allows us to do things like blending apples and oranges naturally. When we express a picture as a sum of pyramid bands and operate on each band separately, we are taking an artistically familiar "painter's" approach. This contrasts with the "physics simulation"

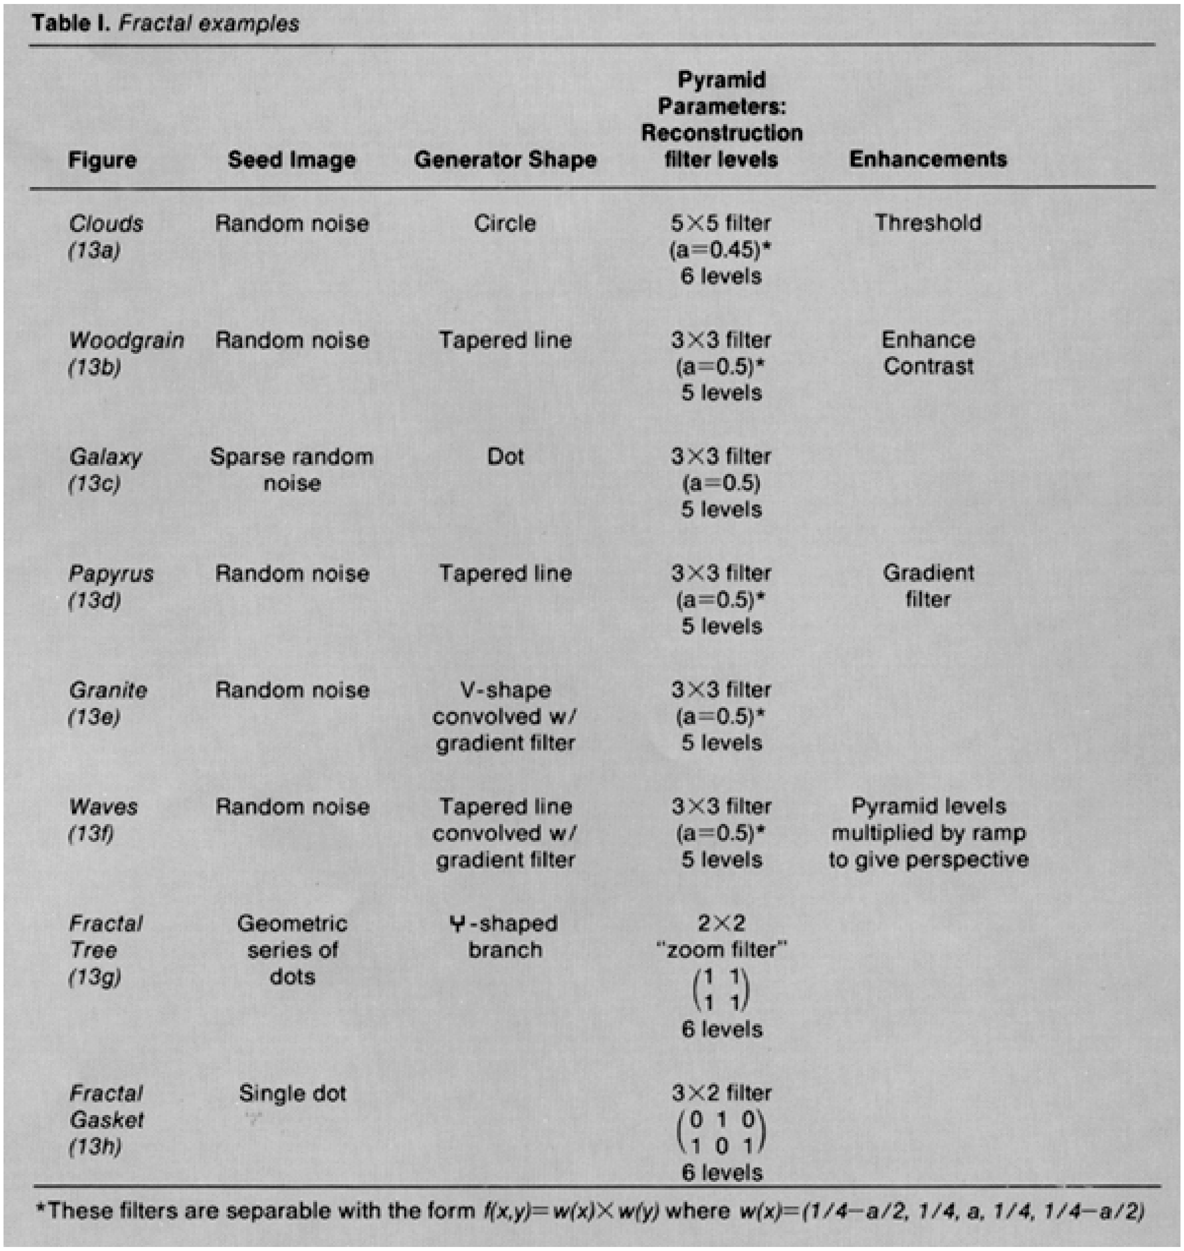

Table I. *Fractal examples*

Figure	Seed Image	Generator Shape	Pyramid Parameters: Reconstruction filter levels	Enhancements
Clouds (13a)	Random noise	Circle	5×5 filter (a=0.45)* 6 levels	Threshold
Woodgrain (13b)	Random noise	Tapered line	3×3 filter (a=0.5)* 5 levels	Enhance Contrast
Galaxy (13c)	Sparse random noise	Dot	3×3 filter (a=0.5) 5 levels	
Papyrus (13d)	Random noise	Tapered line	3×3 filter (a=0.5)* 5 levels	Gradient filter
Granite (13e)	Random noise	V-shape convolved w/ gradient filter	3×3 filter (a=0.5)* 5 levels	
Waves (13f)	Random noise	Tapered line convolved w/ gradient filter	3×3 filter (a=0.5)* 5 levels	Pyramid levels multiplied by ramp to give perspective
Fractal Tree (13g)	Geometric series of dots	Ψ-shaped branch	2×2 "zoom filter" $\begin{pmatrix}1 & 1\\ 1 & 1\end{pmatrix}$ 6 levels	
Fractal Gasket (13h)	Single dot		3×2 filter $\begin{pmatrix}0 & 1 & 0\\ 1 & 0 & 1\end{pmatrix}$ 6 levels	

*These filters are separable with the form *f(x,y)= w(x)×w(y)* where *w(x)=(1/4−a/2, 1/4, a, 1/4, 1/4−a/2)*

approach, which grows more out of the physiw and mathematical modelling traditions. Greater realism can be achieved by using the physics simulation approach, but the complexity and computation time are vastly increased over the multiresolution pyramid approaches described. Considering how well we can fake reality with the pyramid, it seems particularly well suited for making "realistic looking" computer graphic images on small systems.

References

1. P.J. Burt, "Fast Filter Transforms for Image Processing", *Computer Graphics and Image Processing*, 16, pp. 20-51 (1981).
2. P.J. Burt and E.H. Adelson, "The Laplacian Pyramid as a Compact Image Code", *IEEE Transactions on Communication* COM-31, pp. 532-540 (1983).
3. E.H. Adelson, C.H. Anderson, J.R. Bergen, P.J. Burt and J.M. Ogden, "Pyramid methods in image processing", *RCA Engineer*, Vol. 29, No. 6, pp. 33-41 (Nov./Dec. 1984).
4. P.J. Burt and E.H. Adelson, "A Multiresolution Spline with Application to Image Mosaics," *ACM Transactions on Graphics*, Vol. 2, pp. 217-236 (1983).
5. B. Mandelbrot, *The Fractal Geometry of Nature*, Freeman, San Francisco (1982).
6. A. Fournier, D. Fussell, and L. Carpenter, "Computer Rendering of Stochastic Models," *ACM Graphics and Image Processing*, Vol. 25, No. 6, p. 371 (1982).
7. A. Pentland, "Fractal-Based Description of Natural Scenes," *IEEE Transactions*, SMC-13 (1983).
8. J.M. Ogden, "Generation of Fractals Using the Burt Pyramid," presented at the Optical Society of America meeting (October 1985).

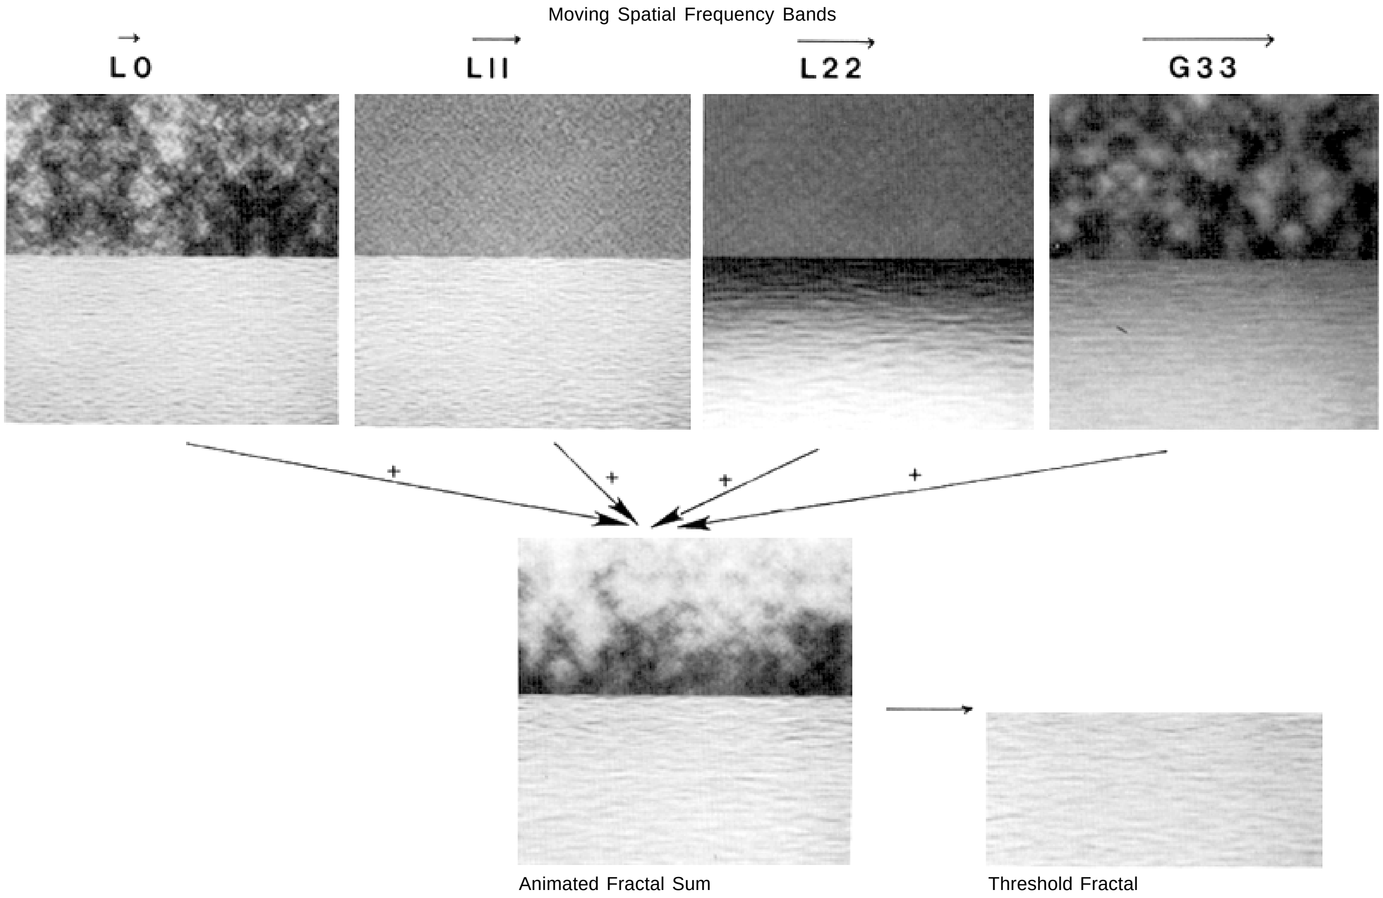

Fig. 14. *Fractal animation method*

Joan M. Ogden received a BS in Mathematics from the University of Illinois in 1970, and a PhD in Physics from the University of Maryland in 1977. After two years as a Postdoctoral research associate at Princeton Plasma Physics Laboratory, she started her own consulting company, working on a variety of applied physics problems.

In 1982 Dr. Ogden joined the Advanced Image Processing Research group at RCA Laboratories, first as a consultant and, in 1984, as a part-time Member of the Technical Staff. Her research interests at RCA include applications of the pyramid algorithm to noise reduction, data compression, and texture generation. She is the author or coauthor of more than 20 articles, has one patent pending, and is coinventor of an issued patent.

Dr. Ogden was awarded a National Science Foundation Visiting Professorship at Princeton University in 1985. She is on leave from RCA to teach at Princeton during the 1985 academic year.

Contact her at:
Princeton University
Princeton, N.J.
(609) 452-4966

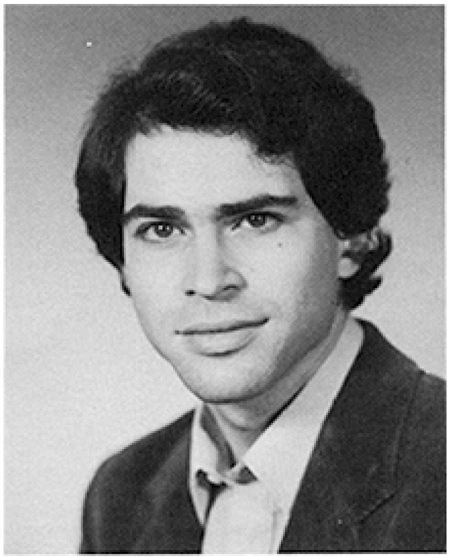

Edward H. Adelson received the BA degree, summa cum laude, in Physics and Philosophy from Yale University in 1974, and a PhD degree in Experimental Psychology from the University of Michigan in 1979. Dr. Adelson joined RCA Laboratories in 1981 as a Member of the Technical Staff. As part of the Advanced Image Processing group, in the Advanced Video Systems Research Laboratory, he has been involved in developing models of the human visual system, as well as image-processing algorithms for image enhancement and data compression.

Dr. Adelson holds two U.S. patents and has published a dozen papers on vision and image processing. His awards include the Optical Society of America's Adolph Lomb Medal (1984) and an RCA Laboratories Outstanding Achievement Award (1983). He is a member of the Association for Research in Vision and Ophthalmology, the Optical Society of America, and Phi Beta Kappa.

Contact him at:
RCA Laboratories
Princeton, N.J.
Tacnet: 226-3036

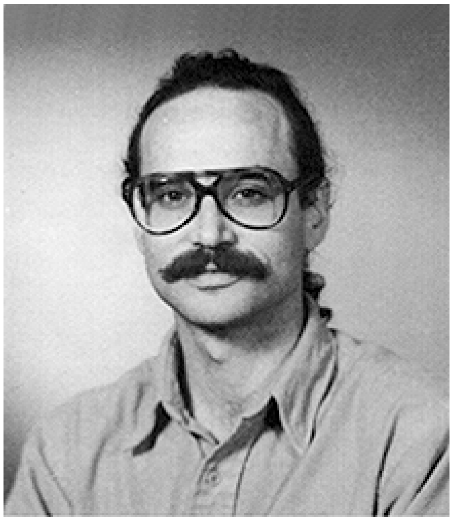

James R. Bergen received the BA degree in Mathematics and Psychology from the University of California, at Berkeley, in 1975, and the PhD degree in Biophysics and Theoretical Biology from the University of Chicago in 1981. His interests concern the quantitative analysis of information processing in the human visual systems, and the application of this analysis to problems in imaging system design and evaluation.

During 1981-82, Dr. Bergen was a Postdoctoral Member of Technical Staff at Bell Laboratories. In 1982 he joined RCA Laboratories as a Member, Technical Staff. His research activities at RCA include theoretical and experimental work on human vision, as well as perceptual aspects of HDTV systems and image processing techniques. In 1983, he shared an RCA Laboratories Outstanding Achievement Award.

Dr. Bergen has published many papers on visual psychophysics and perception, and is a frequent speaker at conferences on these subjects. He is a member of the Optical Society of America and the Association for Research in Vision and Ophthalmology.

Contact him at:
RCA Laboratories
Princeton, N.J.
Tacnet: 226-3003

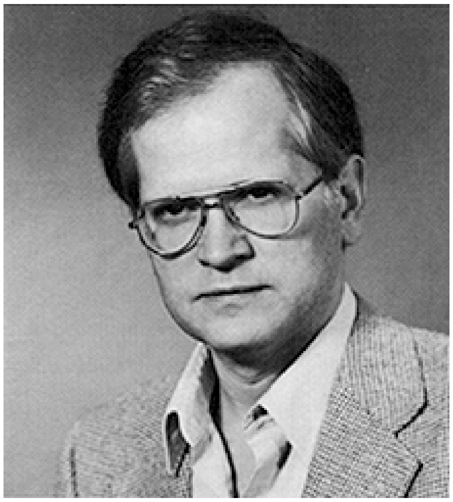

Peter J. Burt received the BA degree in Physics from Harvard University in 1968, and the MS and PhD degrees, both in Computer Science, from the University of Massachusetts, Amherst, in 1974 and 1976, respectively. From 1968 to 1972 he conducted research in sonar, particularly in acoustic imaging devices at the U.S. Navy Underwater Systems Center, New London Connecticut, and in London, England. As a Postdoctoral Fellow, he has studied both natural vision and computer image understanding at New York University (1976-78), Bell Laboratories (1978-79), and the University of Maryland (1979-80). He was a member of the engineering faculty at Rensselaer Polytechnic Institute from 1980 to 1983.

In 1983 Dr. Burt joined RCA Laboratories as a Member of the Technical Staff. He worked in the areas of computer vision, image data compression, and human perception, and in 1984 he became Head of the Advanced Image Processing group.

Dr. Burt has published more than 20 technical papers in his fields of expertise and is a member of the IEEE.

Contact him at:
RCA Laboratories
Princeton, N.J.
Tacnet: 226-2451