


Packing It In

Fractals play an important role in image compression

By IVARS PETERSON



The flickering flames of a campfire highlight the jagged forms of nearby rocks. Gnarled pines crouch on the stony shore. Patches of golden wildflowers, scattered across a field, glow in the last light of a weary sun. In the distance, a saw-toothed range of mountain peaks, wreathed in clouds, tears into the sky.

Although the details in this mountain landscape are easily captured in a photograph, a computer recreation would normally require a complicated program or the storage of millions of bits of data. Researchers are now looking for ways to cut down the amount of information needed to reproduce such a scene in all its detail. Their success could result in efficient means for storing data in a computer's memory, for transmitting photographs over telephone lines, for recognizing specific objects in a landscape and for simulating natural scenery on a computer. Someday, it may even be possible to convey a movie from one computer to another simply by sending a chain of formulas down a telephone line.

The idea is to start with a digitized picture. Such a picture may consist of a 1,000-by-1,000 grid of dots, or pixels. Each pixel is assigned, say, eight bits of data to represent 256 different shades of gray or an equal number of different colors. Thus, the entire picture can be thought of as a string of 8 million ones and zeros, one digit for each bit. If this string of digits can be encoded in some way to produce a new, shorter string of digits, then the image is said to have been compressed. Of course, the compressed string should be able to reproduce, pixel for pixel, the original picture.

Mathematician Michael F. Barnsley and his colleagues at the Georgia Institute of Technology in Atlanta believe that the key to image compression is in the redundancy found in natural forms. For instance, one pine needle is more or less like any other pine needle. "So you don't need to describe each one over and over again," says Barnsley. "You need to describe only

one."

Furthermore, nature is full of shapes that repeat themselves on different scales within the same object. A fragment of rock looks like the mountain from which it was fractured. Clouds keep their distinctive appearance whether viewed from the ground or from an airplane window. A tree's twigs often have the same branching pattern seen at the tree's trunk.

In all these examples, zooming in for a closer view doesn't smooth out the irregularities. Instead, the objects tend to show the same degree of roughness or branching at different levels of magnification. In 1975, Benoit B. Mandelbrot, now at Harvard University, coined the word "fractal" to describe such irregular and fragmented shapes, which can be magnified endlessly and still retain their complicated structure (SN: 3/21/87, p.184; 1/21/84, p.42).

The mathematics of fractals has already been used to create images that look a lot like clouds, mountains and other forms (SN: 11/20/82, p.328). Usually the process involves feeding into a computer a small set of numbers that generates a basic shape, such as a triangle. That shape is then recreated many times on smaller and smaller scales within the original figure. Random variations are thrown in to make the image look a little rougher and, as a result, more realistic. Such artificial landscapes can be mathematically magnified to reveal more detail, just as a close-up lens probes deeply into a natural scene.


However, these fractal images have been created largely by trial and error. They're the result of computer doodlers stumbling upon mathematical procedures (or algorithms) that happen to lead to drawings that look like natural objects. In contrast, Barnsley has tackled the problem of starting with a natural object and finding a specific fractal to fit it. He and his co-workers have been studying how a scene's geometry can be analyzed to generate an appropriate set of rules that can then be used to recreate the scene. Because fractal mathematics is

a compact way to store the characteristics of an object, this approach would compress the content of an image into just a few equations. Barnsley described his scheme for image compression recently in Chicago at the annual meeting of the American Association for the Advancement of Science. Mathematical details appear in the *PROCEEDINGS OF THE NATIONAL ACADEMY OF SCIENCES* (Vol.83, No.7).


Several important mathematical concepts lie at the heart of Barnsley's scheme. The procedure relies on mathematical operations called affine transformations. An affine transformation behaves somewhat like a drafting machine that takes in a drawing (or the coordinates of all the points making up the lines in a drawing), then shrinks, enlarges, shifts or skews the picture and, finally, spews out a distorted version of the original.

Affine transformations can be applied to any object — triangles, leaves, mountains, ferns, chimneys, clouds — or even the space in which an object sits. In the case of a leaf, the idea is to find smaller, distorted copies of the leaf that, when fitted together and piled up so that they partially overlap, form a "collage," which approximately adds up to the original, full leaf. Each distorted, shrunken copy is defined by a particular affine transformation — a "contractive map," as it's called — of the whole leaf. If it takes four miniature copies of the leaf to approximate the whole leaf, then there will be four such transformations.

Now the original image or "target," whether leaf or cloud, can be thrown



Barnsley et al./PNAS



away, leaving only the corresponding collection of affine transformations. These can be used to recreate the original image by, in a sense, mathematically molding a chunk of space. That's done by starting with a point somewhere on a computer screen. Applying one of the available transformations to the point shifts it to a new spot. That spot is marked. Again, randomly applying one of the transformations shifts the point to another location. The new spot is colored in, and the process is repeated again and again.


Amazingly, the point doesn't hop about aimlessly, first pulled one way, then another. Instead, a pattern gradually emerges. The point's colored tracks add up to an image called an attractor. In the case of the four leaf transformations, the attractor is an object that looks very much like the original leaf.

An approximate collage of a leaf can be constructed from smaller, distorted copies of the whole leaf (left). Randomly iterating the corresponding set of four affine transformations generates an attractor that looks like the original leaf (right).

How this part of the process works can best be seen in a simple example that can be worked out on a sheet of squared graph paper. Imagine a rectangle with three of its corners labeled 1, 2 and 3. Suppose there are also three transformations. The first shrinks everything toward corner 1, the second shrinks everything toward corner 2, while the third shrinks everything toward corner 3, always by a factor of one-half.


Select a starting point somewhere on the grid. Randomly apply one of the three transformations. That transformation will designate a new point halfway between the original point and one of the corners, depending on the choice of transformation. Randomly applying a second transformation (it may be any one of the three available) locates another point halfway between the previous point and the appropriate corner. Chasing the point around the sheet of graph paper, marking each landing spot, produces a remarkable fractal object known as a Sierpinski gasket. This object consists of a complicated array of triangles nested within triangles, and so on to smaller and smaller scales.

It turns out that any particular collection of affine transformations, when iterated randomly, produces a unique fractal figure. The trick is to find the right group of transformations to use for generating a particular image. That's done using the "collage" process (for example, a leaf covered by little copies of itself) described in an earlier paragraph. Furthermore, the probability of using a certain transformation need not be the same as the probability of applying any other transformation in the set. And because some grid squares are likely to be visited



Barnsley et al./PNAS

Successive close-ups of the fronds of a fractal fern reveal new features at each level of magnification.



Barnsley, Sloan, Reuter/© 1987 Georgia Tech Research Corp.

more often than others, keeping track of the relative number of visits to each square provides a way to specify color brightness and intensity or to define a gray scale. In this way, a lot of information is packed into a few formulas.

“We can compress images by encoding them as the collection of rules for which the random iteration generates the image,” says Barnsley. “When people see it on the screen, when they see us do it on a computer, it blows them away. It’s an extraordinary thing to watch.”

Using this technique, Barnsley and his group have been able to create remarkable three-dimensional renderings of

natural objects such as ferns. He has also been able to come up with reasonable fractal reproductions of photographs taken from magazines like NATIONAL GEOGRAPHIC.

In one instance, Barnsley used 57 affine transformations (or maps, as they are often called) and four colors – a total of 2,000 bytes of information – to model three chimneys set in a landscape against a cloudy sky. “The idea is that we can fly into this picture,” he says. “You can pan across the image, you can zoom into it, and you can make predictions about what’s hidden in the picture.”

As the picture is blown up to show more and more detail, parts of it degenerate into nonsense, but some features, such as

This scene, showing smoking chimneys in a landscape, was computed using 57 affine transformations. Zooming in on the smoke shows that it remains reasonably realistic even to great magnifications, yet the same data base is used throughout.

the chimneys, the smoke and the horizon, remain reasonably realistic, even when the image compression ratio is better than 10,000 to 1.

“It’s only a matter of how far we’re willing to agree the picture makes sense,” he says. “Some of it is quite nonsensical at this level, but nonetheless, it shows the potential for this method of image compression. The prospects look very, very good for the future.” □