

Storing Art Images in Intelligent Computers

Any reproduction of a visual artwork creates a representation that fails to capture some properties of the original. Since most art is not sufficiently portable to make it available to all possible viewers, we accept the necessity of making reproductions with the consequent loss in properties, some of which may be vital to the proper appreciation of the artwork. Photography is the most well established of these reproduction methods. Critical viewers know that photographic reproductions fail to capture color, texture, three-dimensionality, surface texture, tonal gradations, fine detailed structure, and movement. Still, we make photographs and use them for teaching, scholarship, archiving, criticism and conservation of visual art materials.

Other reproduction methods are newer and less familiar, presenting new opportunities and demanding new compromises. Two notable such methods are the analog representations used on video discs and the digital represen-

Joan L. Kirsch and
Russell A. Kirsch

tations used on optical discs for storage of art images. These methods differ in important ways from photography, notably in their ability to capture movement and to be rapidly searched among large collections of such images; however, their spatial resolution is poor compared to photography.

We could continue to list other methods of representing artworks, but almost all common methods we might list share one important property—they ignore the art by being passive with respect to the content of the image. To suggest that a photograph or a video disc recording ignores the artworks recorded is to imply that there is an alternative. Can a storage medium look at a work of art in such a way as to be said to perceive it? We would hesitate to answer 'yes', unless some intelligent process had intervened in the recording so as to show some evidence of understanding. Ordinarily, we think of understanding associated only with human capabilities. However, great advances in the field of artificial intelligence over the past three decades of research have en-

ABSTRACT

Images of artworks can be stored in media that preserve different characteristics of the original. Differences exist in the extent to which we can preserve color, three-dimensionality, surface texture, fine structure, tonal gradations, temporal variations and other characteristics that lend uniqueness to individual artworks. Usually, we are willing to sacrifice some of these characteristics in exchange for the permanence and recoverability offered by storage media. Thus, a color slide (diapositive), which is a common medium for storing images of artworks, compromises all of the above properties to different extents but is nevertheless considered useful for the archival properties it offers for images of artworks. Digital storage media used in conjunction with computers offer new opportunities and demand new compromises in storing art images. An unusual challenge is offered by the possibility of providing intelligence to a computer. The authors make clear the sense in which we may ascribe intelligence to the computer and how this may be used to 'perceive' the image of an artwork. The computer then uses its knowledge of the artwork with respect to a large class of such works not only for archival storage but also to achieve economy in the use of the storage medium. The authors illustrate the achievement of storage economy as much as tens of thousands of times greater than storage without intelligence. The intelligence is provided to the computer as syntactic descriptions of classes of artworks. The syntactic descriptions incorporate insight from the art historian, critic or artist who uses innovative tools like shape grammars to provide the computer with a small part of the intelligence that the educated human viewer brings to the perception of the artwork.

Joan L. Kirsch, The Sturvil Corporation, Clarksburg, MD 20871-0157, U.S.A.
Russell A. Kirsch, The Sturvil Corporation, Clarksburg, MD 29871-0157, U.S.A.
Received 7 April 1988.

Fig. 1. Albrecht Dürer, *Melancholia I*, engraving, 240 × 187 mm, 1514.



Fig. 2. Computer scanned and reproduced detail with a scanning resolution of 50 micrometers. Eight levels of grey tone are shown in individual pixels.






Fig. 3. Einstein postage stamp, issued 4 March 1979, U.S. Treasury, Bureau of Printing and Engraving.



Fig. 4. Computer scanned and enlarged detail from below the eye in the Einstein engraving.



Fig. 5. Extreme enlargement of detail below the eye in the Einstein engraving.

couraged us to think of computers as having some degree of understanding in many disciplines. We see no reason why these insights from artificial intelligence should be denied to the visual arts. And we might expect that there would be different degrees of understanding to which computers might aspire, ranging from the superficial to that of an art historian, critic, educator or artist.

The question we raise about computer perception is not moot because it has been demonstrated recently that storage media using computers in novel ways can exhibit some limited understanding of works such as the architectural drawings of Frank Lloyd Wright [1] and the fine art images of artists like Joan Miró [2] and Wassily Kandinsky [3]. Computer understanding has also been demonstrated for paintings by Richard Diebenkorn [4], for a comparison of paintings by Georges Vantongerloo and by Fritz Glarner [5], for woodcuts by the German artist Jacob Fauser [6] and for contemporary computer art by Harold Cohen [7]. For each of these artists, it has been shown that a computer can be said to 'understand' the artist's work in one of several senses. In the simplest case, the computer can classify an image upon viewing it. In a deeper sense, the computer can be said to understand the image by being able to place it in chronological sequence with images of other works by the same artist. In a still deeper sense, the computer understands the artwork because it has produced the work by itself. And finally, in the deepest

sense, the computer can be said to understand an art image by being able to recognize its formal properties and by being able to synthesize other works that are stylistically similar.

All these capabilities have been demonstrated and are the subject of ongoing research. We therefore feel justified in saying that a computer need not ignore an artwork that it sees. If the matter ended there, we might precipitate discussions in aesthetics, or even in epistemology, without any necessary practical consequences. But there are many practical consequences that follow from computer understanding of art images. In this paper we discuss one of these: the ability of computers to reduce greatly the storage requirements needed for archiving such images.

THE STORAGE PROBLEM

Digital storage of images is more expensive than photographic storage. This expensiveness is apparent if we consider only the cost of producing a stored record. The reason for storing images is usually to perform subsequent operations on the stored images, ranging from the common case of retrieval of the image to the more esoteric cases of modification of the image and answering questions about the image. All of these operations can be done on photographic images by varying degrees of manual manipulation, but when speed is important, digital storage is more economical for retrieval of images. When automated

operations of the more esoteric kind are desired, only digital storage allows us to perform such operations as image enhancement, pattern recognition, image comparison and searching of images directly without the intervention of a text-based description. We are therefore led to consider the storage of images of artworks in the form of digital records accessible to conventional computers.

The digital medium of choice has become the optical digital disc. Capacities for such storage range from 600 megabytes on 'compact disc read only memories' (CD-ROM), which can be read only by the user, to 2000 megabytes on the much more expensive 'write once read many' (WORM) optical discs, which can be recorded once by the user and read thereafter as often as desired, as can the CD-ROMs. Complementing the storage media is the technology for scanning images and for producing the digital signal to be recorded on the optical discs. This technology is much more mature, having been used with computers for over three decades.

It is useful to understand the relations among artworks, their scanned images and the storage requirements needed for recording those images. To make this relation clear, consider the engraving by Dürer (1514) shown in Fig. 1. A print of this engraving was scanned with a scanner having a resolution of 50 micrometers per picture resolution element (pixel). We can see the extent to which this black-and-white print is adequately resolved with such a comparatively fine scan by

studying the eye detail shown in Fig. 2. Here we can see the individual pixels. We notice that the line work in Dürer's engraving, which resolves quite well on the original print, is poorly resolved at this scanner resolution. Nevertheless, the whole scanned image requires a storage capacity of 18 megabytes at this (inadequate) resolution. Furthermore, this image has been scanned to produce an image with about 65 times as much information as a conventional black-and-white TV image contains. It follows, therefore, that a scanned image of the Dürer print, using only the resolution common with TV cameras and video-discs, will be quite inadequate if the fine structure of the engraved lines is to be resolved. However, if no detailed structure is needed in the stored image, the coarse resolution may be adequate.

Occasionally it is suggested that, for so-called 'line work', the use of computer graphics technology is more appropriate than scanned images. Certainly, there are images, of at least peripheral interest in the fine arts, for which line drawings can create adequate representations of the original works. The most obvious examples occur in architectural drawing. Here, the use of computer graphics line drawings can provide greater storage economy than can scanned images. There are commercial machines in the com-

puter-aided design (CAD) field that scan architectural drawings and immediately convert them to line drawings for many purposes, one of which is storage economy.

We must be careful not to infer that other kinds of drawings can be susceptible to the same treatment. Again, an illustration will make the issue clear if we consider line engravings. Figure 3 shows a print from a contemporary fine engraving made for a postage stamp by a master engraver at the United States Treasury. Upon superficial inspection, it appears to be a line drawing and therefore capable of being represented by a sequence of line segments commonly used in computer graphics. But if we look at an enlarged detail of the image (Fig. 4), we begin to see, and can see more clearly in a further enlargement (Fig. 5), that the image is composed of complex shapes and not of lines at all. The engraver has control over line quality that goes beyond the control that can be exercised, for example, by an etcher.

To test whether the detailed structure of Fig. 5 is an artifact of high magnification or an essential part of the composition, we performed a simple test. We showed the detailed structure of Fig. 5 to the engraver who had made the original engraving a few years previously. We asked him to locate the detailed image of Fig. 5 in the whole

composition of Fig. 3 without the helpful intermediate illustration of Fig. 4. He took about 2 minutes to describe the structure of Fig. 5, immediately concluding that it must occur at the edge of Einstein's eye, and then, in another minute, located the exact spot where the detail could be found. The conclusion from this simple experiment is something known to most artists: everything counts! And so the gratuitous proposal that this engraving be treated as 'nothing more' than a line drawing fails a simple test.

We expect that most true drawings in which line structure is important would similarly fail to be adequately represented by computer graphics line drawings. And therefore, for drawings, only the crudest representations, such as TV scans for images, can be achieved with the obvious use of the most readily available technology.

RECOVERING IMAGES FROM STORAGE

When an image is scanned with a TV camera, there is a direct correspondence between the original image, the scanned version, the version stored in computer memory and, finally, the image that is recovered for display. For every pixel extracted by the scanner from the original image, a corre-

Fig. 6. Richard Diebenkorn, *Ocean Park #111*, oil and charcoal on canvas, 336.2 x 336.7 cm, 1978. (Hirshhorn Museum and Sculpture Garden, Smithsonian Institution, Museum Purchase, 1979)





Fig. 7. A linear representation of the structure of Fig. 6.






Fig. 8a. A grammar for the linear structure of Diebenkorn's *Ocean Park* paintings.

sponding one is stored in memory and a corresponding one is displayed. But this correspondence need not be so direct. It is acceptable for the stored image not to correspond to the scanned image, so long as the recovered image still has the direct correspondence with the scanned image. This is the approach taken with common code compression techniques. In code compression, the redundancy of the scanned image is exploited to save on storage requirements. What is stored is an encoded version of the scanned image in which certain commonly occurring arrays of adjacent pixels (such as all white or all black ones) are represented with a more economical code than would be used with more direct storage. Then, when the image is recovered for display, it is

decoded to produce the correspondence between the displayed image and the original scanned image. The resulting storage economy can range up to a factor of about tenfold.

With encoded images, it is proper to speak of the displayed image as having been reconstructed from the encoded representation in storage. There are two kinds of such reconstructions, unique and approximate. The first kind reconstructs the scanned image identically; the second does not. The reason for using these two types of encodings is that unique reconstruction achieves absolute fidelity to the scanned image, but approximate reconstruction can achieve greater storage economy.

The cost of storage is determined by the kind of approximation one uses. A

simple kind of approximation that is widely used depends on the degree of the fine structure in the image. This approximation starts by transforming the image into an entirely new image not apparently like the original. A useful property of this new image is that all the fine, detailed structure is located in one part of the image and the gross structure is located in another part. It is possible to recover the original image by a certain reverse transformation. But, first, the part of the transformed image containing the fine structure is removed. This yields a smaller image that can be stored more economically. To recover the original image, a reverse transformation is made, which yields an approximation to the original image. But the fine structure is gone. This so-called




Fig. 8b. A grammar for the linear structure of Diebenkorn's *Ocean Park* paintings (continued).

'spatial frequency filtering' does not depend on any interpretation of the image. All fine structure is treated equally, whether it be a pimple on a nose or a pebble on the ground. And it can yield storage economy by as much as a factor of 10 to 100.

The appearance of a reconstructed image that has been spatial frequency filtered is easy to compare with the unfiltered image. The filtered one appears to be a blurred version of the original. Other kinds of approximate reconstruction are less easy to compare with their original images. One such reconstruction method, which claims to achieve large storage economy, uses fractals [8]. A fractal curve has a complex structure that is suggestive of, but different from, the structure of natural objects. To the super-

ficial observer, images constructed from fractal curves often appear realistic. This property has been exploited in computer graphics research to create artificial images that lack the obvious geometric shapes often associated with computer-generated images. For example, pictures of flowers or mountains have been constructed with fractal curves to create a pleasing illusion that, nevertheless, is unconvincing to a botanist or geologist. Furthermore, even scanned photos that have been encoded for storage and reconstructed with fractal curves can create the illusion of being 'realistic' reconstructions. We would expect the pimples and pebbles would look realistic, but not like their originals. So, if we are willing to accept that once one has seen one pebble one has seen all

pebbles, this method of achieving storage economy has some attractiveness—but only if we are tolerant of deviant pebbles! For those who would 'see the world in a grain of sand', such economy is less acceptable.

PROVIDING IMAGE INTELLIGENCE TO THE COMPUTER

The storage and recovery schemes thus far discussed exhibit a benign indifference to the content of the images dealt with. Even for images that have been approximately reconstructed, the approximation is not based on any understanding of the image content. Rather, the image's statistics constitute the basis for accepting and re-

jecting image properties during approximate reconstruction. Surely we can do better than this!

Indeed we can in several ways. We consider first a method whereby the computer clearly has some knowledge of the image content, although the knowledge is insufficient to reconstruct an image. Rorvig has shown how primitive-feature extraction procedures can be used by the computer to classify a set of woodcuts [9]. The primitive features consist of lines, angles and density distributions in the image. These were correlated with aesthetic judgments made by human viewers of the same set of images. Close correspondence between the human judgment and the machine ranking was seen to be possible. Clearly, the simple primitives provided to the computer were insufficient to characterize the images completely or to reconstruct them. This procedure, nevertheless, could provide the computer with an elementary

component of understanding of an image similar to the understanding by a human viewer.

A much larger step in the direction of computer understanding of art images occurs in the work of Harold Cohen [10]. The images he provides to the computer do not come from any external source but, in fact, are generated by the computer itself. A program written by Cohen uses algorithms that represent principles of his composition. These are used to generate drawings both abstract and figurative. The drawings all have a readily identifiable style and can be produced in unlimited quantity. It is clear that the computer has sufficient intelligence to produce images that are both aesthetically interesting and consistent in style.

Since Cohen's images are generated *ab initio*, the storage problem for these images has a transparently simple solution. To reconstruct any image it is only necessary to know the

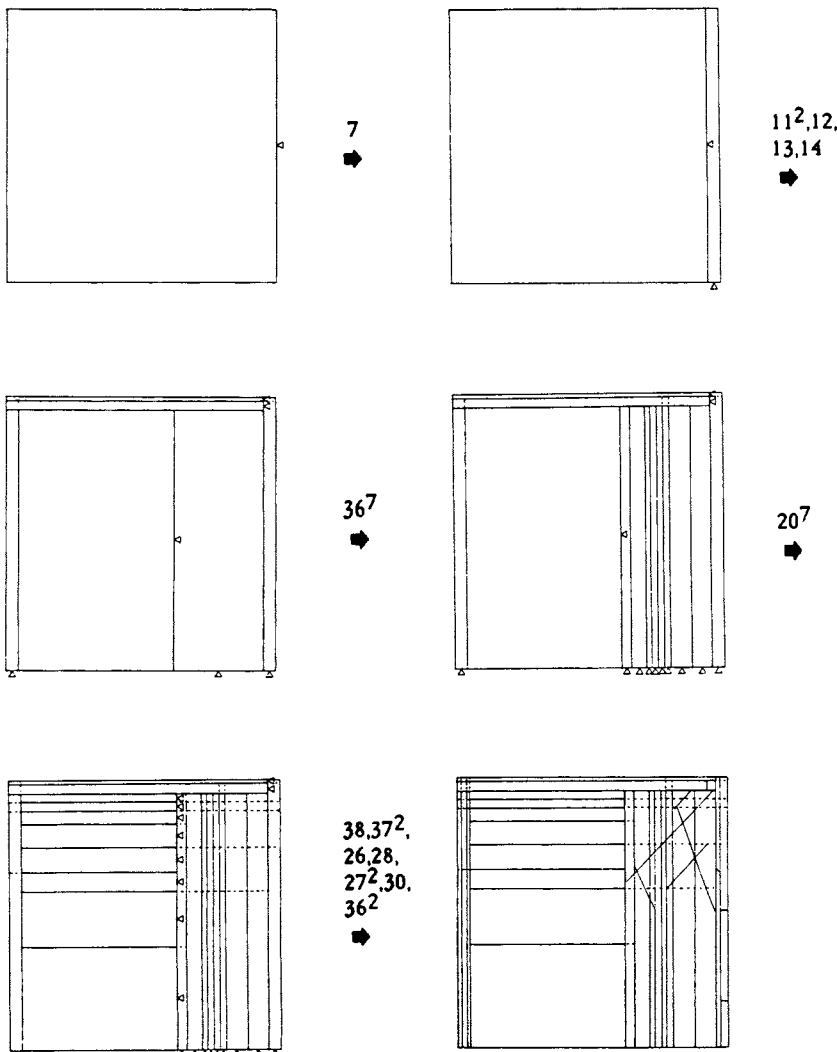
algorithms for image generation and the specific options exercised by the computer in the production of the particular image. The storage requirement for this information, which is peculiar to any single image, is modest. So, in a sense, we have a case of a computer that knows almost everything about such an image and can convert that knowledge into economical storage of the image. Of course, Cohen's computer knows all about its own images and nothing about ones generated by any other artist. But it is interesting to speculate whether such a machine might be able better to understand drawings of other artists by using its demonstrated ability to generate drawings of its own! At present that question is quite open.

If the computer's knowing a class of images (in the sense that the whole class of Cohen's drawings can be generated) leads to great storage economy, it is important to know whether that approach can be achieved when we are dealing with pre-existing images. The question was answered affirmatively in 1964 [11]. The technique introduced was drawn from computational linguistics but modified to deal with images: the use of picture grammars. It was shown that a large class of images could be described succinctly by a grammar that successively transforms two-dimensional shapes into final forms that correspond to recognizable images. The resulting field of syntactic analysis of images developed many tools, the most productive of which was the shape grammar [12].

Although shape grammars are important conceptual devices for describing images, they have made more tangible contributions. Beginning in 1977, architects created shape grammars for a wide range of disciplines, including buildings [13,14], landscapes [15], furniture designs [16] and window patterns [17].

The schematization required by shape grammars poses little limitation in design areas such as architecture, where schematic drawing is already widely accepted. Consequently, grammars have been readily accepted in these cases. As we might predict, great economy of representation can be achieved because each grammar implies a large, or infinite, set of designs. To store the description of a single design, it is necessary only to specify, within the grammar, the options exercised that distinguish that design. As

Fig. 9. Six stages in the generation of Fig. 7 according to the grammar of Fig. 8. The notation 11^2 denotes rule 11 applied 2 times.



before, such a specification is orders of magnitude less costly in storage than would be a scan of the drawing for a design.

These architectural examples were the first practical use of grammars to describe a very large class of images in such a way that the images could be regenerated from the class description. But the question still remained whether images such as paintings could be treated this same way. Again, we have answered this question affirmatively by constructing a grammar for Diebenkorn's *Ocean Park* series of paintings [18]; the question has subsequently been answered also by Knight, who constructed a sequence of grammars for the successive stylistic periods of Vantongerloo and Glarner [19]; and by Lauzzana and Pocock-Williams, who constructed a rule system for the skeletal organization of Kandinsky's paintings [20]. All these examples provided the computer with sufficient understanding of the target class of images for us to credit the computer with substantial understanding of the style of the artists. And it is this understanding that can be exploited for many purposes, among them the

economical storage of the approximate representations of the images described by the grammars.

ECONOMY RESULTING FROM INTELLIGENT IMAGE STORAGE

We wish, now, more specifically to calculate the economy that results from storage of images in an intelligent computer. We will start with Richard Diebenkorn's paintings and determine how much storage is needed for an image described by the grammar [21]. Because the grammar contains recursion, there are an infinite number of paintings described (all of them purportedly in the style of Diebenkorn). For any such painting, there is a sequence of rule applications that, when applied in the proper order, will result in a representation of the target painting. At certain points in the process, the grammar provides a set of alternatives among which a choice may be made in producing a picture. We must use information to specify which choice is to be made. If there are N alternatives offered by the grammar at

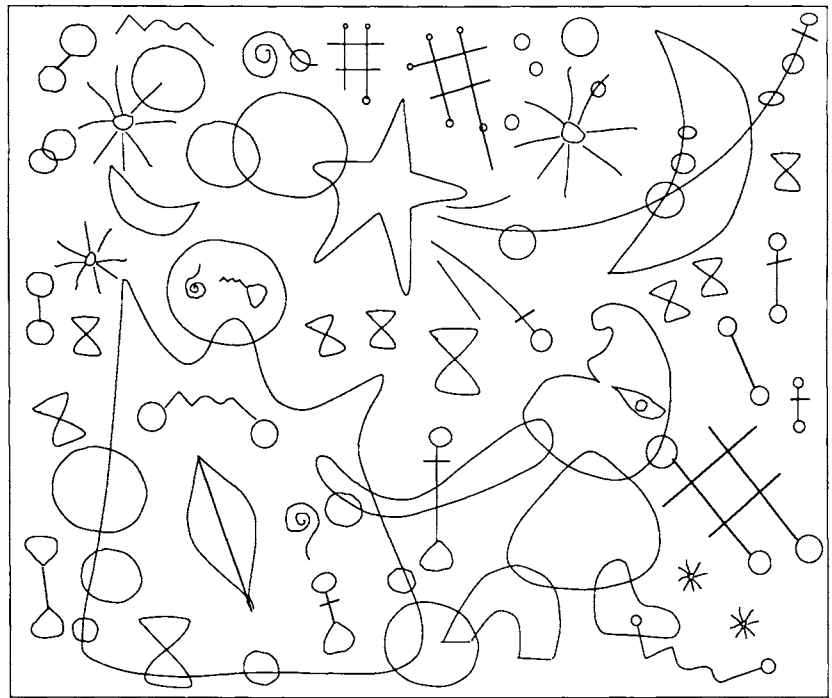


Fig. 10. An artificial Miró composition based on a catalog of prototype shapes.

that point, making such a choice always can be done by specifying no more than the logarithm (to the base 2) of N bits of information. Thus, if there were eight choices, any one of these could be specified with no more than 3 bits of information, four choices with 2 bits, two choices with 1 bit, etc. Then, by summing the information associated with each of the choices made in the production of the picture, we determine how much information is needed to specify the whole picture, with respect to the grammar.

For an example, we will choose the picture of Diebenkorn's *Ocean Park No. 111* shown in Fig. 6. The representation to be used is shown in Fig. 7. Such representations of the paintings are provided by the grammar taken from Kirsch [22] as shown in Fig. 8. The process of generating Fig. 7 from the grammar is shown in Fig. 9. Here we see six of the 32 stages in the generation process.

We start with a blank canvas as shown in the first image of Fig. 9. Corresponding to this image is the starting symbol of the grammar, OPP. We notice that there are three rules (nos. 1, 2, 3) that provide options for developing this blank canvas. On the average, with $\log(3) = 1.58$ bits of information, we can specify our choice, which is, for this picture, the second rule. Now, the grammar confronts us with a single rule for further developing the

resulting OP/S that resulted from rule 2. This single rule (no. 4) creates a blank canvas labeled with a Q. Since there are no choices provided, no information need be specified, and the cumulative total information remains at 1.58 bits. At this point we are confronted with a set of four rules (nos. 5, 6, 7, 8), any one of which may be invoked. For the intended picture, we choose rule 7 by specifying $\log(4) = 2.0$ bits of information. This raises our cumulative total to $1.58 + 2.0 = 3.58$ bits of information. We thus far have produced the second image in Fig. 9.

To produce the third image in Fig. 9, we first invoke rule 11 twice to produce the two top horizontal bands ($2 \times \log(4)$), rule 12 once to produce the vertical band on the left ($1 \times \log(4)$), rule 13 to produce the medial vertical band ($1 \times \log(4)$), and finally rule 14, which removes the /S subscript that allows the previous rules to be applied ($1 \times \log(4)$). This adds 10.0 bits to yield a cumulative total of 13.58 bits.

The fourth image in Fig. 9 is produced by seven applications of rule 36. Each such application represents a choice among five alternatives, yielding a contribution to the total of $7 \times \log(5) = 16.25$ bits.

Then the fifth image is produced by seven applications of rule 20, which is chosen among five alternatives, yielding $7 \times \log(5) = 16.25$ bits.

Finally, the target image is produced from rule 38 (2.32 bits), rule 37 twice (4.64 bits), rule 26 (2.0 bits), rule 28 (2.0 bits), rule 27 twice (4.0 bits), rule 30 (2.32 bits) and rule 36 twice (4.64 bits). The cumulative total is thus 68.0 bits.

What the above calculation shows is that, if we provide the computer with the intelligence to understand the whole class of Diebenkorn compositions represented by the grammar of Fig. 8, any picture thereafter can be specified very economically. The example requires only 68 bits, or less than 9 bytes, of data to store it. This number should be compared with the storage requirements for an ordinary TV scan of the same painting, which would require about 0.25 million bytes. The dramatic difference in storage requirements is accounted for not only by machine intelligence but also by the fact that two different kinds of representation are being compared.

It is important to point out that these two kinds of representation also elicit different levels of understanding from people and machines. For the computer, a TV scan carries no meaning whatsoever. However, a grammatical representation does furnish it with a degree of understanding. Of course, for people, the TV scan is immediately comprehensible. The grammar is a human artifice created with great effort and insight, but it can produce a schematic representation that can be understood by both people and machines. If we wish to use the computer for operations upon images such as automatic searching of a large collection of images rather than textual descriptions of these images, then the grammatical approach is the only one of the two storage methods that makes such a search possible.

The extreme storage economy available for the Diebenkorn paintings, in an intelligent computer possessed of a grammar, results from our generosity in accepting the linear schematization of his works that we see in examples like Fig. 7. Like any representation, such linear schematization results in a loss of essential information from the original painting. But the loss achieves the dual gains of storage economy and intelligent understanding by the computer.

How might we preserve more information in a representation? An obvious approach is to help the computer to understand shape. Since shape does not appear essential in the

Diebenkorn paintings, we can consider an artist like Joan Miró, for whom shape is essential. In a recent article, we discussed how artificial Miró compositions can be constructed based on a catalog of Miró shapes [23]. A typical such example is shown in Fig. 10, which may be compared with another example shown in Fig. 8 of the aforementioned article and with a photograph of a Miró in Fig. 9 of the same article. These compositions in the style of Miró can be represented very economically since each shape is drawn from a small catalog of prototypes, each of which can be specified with about 5 bits for identification, another 20 bits for size and another 20 bits for location. With those 45 bits representing each shape, a composition containing N shapes can be represented with $45 \times N$ bits.

While this method results in much greater economy of representation than TV scanning, we have arrived only at an ad hoc solution. Leyton [24] and Jakubowski [25] have suggested better ways to represent shape. Leyton's scheme uses a grammar for the shapes themselves just as we have used one for the compositional arrangements in Diebenkorn. He has devised grammar rules that correspond to the successive deformations that transform a simple circle into elaborate shapes with invaginations and evaginations that occur at maxima and minima of curvature. We currently are investigating a first such grammar for Miró shapes. This grammar allows us a wide choice of shapes as well as a natural characterization of how the artist draws and how the viewer sees.

CONCLUSION

We have seen several examples of how a computer may be given the intelligence to perceive representations of artworks. We have also seen how expensive it is to provide raw scanned images to a computer storage system. These scanned images are effectively invisible to the computer, a fact deceptively easy to forget since those same images are readily visible to the human viewer.

Once the large investment in providing appropriate intelligence to the computer has been incurred, many rewards accrue to the art historian, critic, educator, archivist and artist. We have directed our attention only to the question of economy of storage.

But it is reasonable to expect that a computer that can view images intelligently in one way can do so in other ways, too.

Acknowledgement

We wish to thank Harry Rand and Sanford Resler for pushing us toward the asymptotes of English and truth.

References and Notes

1. H. Koning and J. Eizenburg, "The Language of the Prairie: Frank Lloyd Wright's Prairie Houses", *Environ. and Planning B: Planning and Design* 8, 295-323 (1981).
2. R.A. Kirsch and J.L. Kirsch, "Describing Painting Styles to a Computer", *Leonardo* 21, No. 4 (1988) (forthcoming).
3. R.G. Lauzzana and L. Pocock-Williams, "A Rule System for Aesthetic Research in the Visual Arts", *Leonardo* 21, No. 4 (1988) (forthcoming).
4. J.L. Kirsch and R.A. Kirsch, "The Structure of Paintings: Formal Grammar and Design", *Environ. and Planning B: Planning and Design* 13, 163-176 (1986).
5. T.W. Knight, "Transformations of De Stijl Art: The Paintings of Georges Vantongerloo and Fritz Glarner", *Environ. and Planning B: Planning and Design* (forthcoming).
6. M.E. Rorvig, R. Helfer and S. Fitzpatrick, "Automatic Image Classification by Psychometric Mapping", (Working Paper 87-2, Graduate School of Library and Information Science, University of Texas at Austin, 13 November 1987).
7. Harold Cohen, Becky Cohen and Penny Ni, *The First Artificial Intelligence Coloring Book* (Los Altos, CA: William Kaufman, Inc., 1984).
8. M.F. Barnsley and A.D. Sloan, "A Better Way to Compress Images", *Byte* (January 1988) pp. 215-223.
9. Rorvig, Helfer and Fitzpatrick [6].
10. Cohen, Cohen and Ni [7].
11. R.A. Kirsch, "Computer Interpretation of English Text and Picture Patterns", *IEEE Trans. Electronic Computers* 13, 363-376 (1964).
12. George Stiny and James Gips, *Algorithmic Aesthetics* (Berkeley, CA: Univ. of California Press, 1978).
13. Koning and Eizenburg [1].
14. F. Downing and U. Flemming, "The Bungalows of Buffalo", *Environ. and Planning B: Planning and Design* 8, 269-293 (1981).
15. G. Stiny and W.J. Mitchell, "The Grammar of Paradise: On the Generation of Mughul Gardens", *Environ. and Planning B* 7, 209-226 (1980).
16. T.W. Knight, "The Generation of Hepplewhite-Style Chair Back Designs", *Environ. and Design B* 7, 227-238 (1980).
17. G. Stiny, "Ice-Ray: A Note on the Generation of Chinese Lattice Designs", *Environ. and Planning B* 4, 89-98 (1977).
18. Kirsch and Kirsch [4].
19. Knight [5].
20. Lauzzana and Pocock-Williams [3].
21. Kirsch and Kirsch [2].
22. Kirsch and Kirsch [2], [4].
23. Kirsch and Kirsch [2].
24. M. Leyton, "A Process Grammar for Shape", *Artificial Intelligence Journal* (1988) (forthcoming).
25. R. Jakubowski, "A Structural Representation of Shape and Its Features", *Information Sciences* 39, 129-151 (1986).