The Anatomy of Painting Style: Description with Computer Rules

Joan L. Kirsch and Russell A. Kirsch

Abstract—Experts recognize artistic style by invoking knowledge of plastic elements and their juxtaposition. Algorithmic methods now make it possible for this knowledge to be expressed to a computer. Compared to the alternative of describing style in natural language, the computer methods offer the advantage that they can be tested for validity. The tests are accomplished by using the computer to generate new compositions in the style described. The authors demonstrate these methods with algorithmic descriptions of the styles of Richard Diebenkorn and Joan Miró and the generation of new compositions in their styles. A shape grammar algorithm for Diebenkorn is presented which accounts for the linear facture of his *Ocean Park* series. The problem of shape in Miró's work is tackled, and progress is reported on synthesizing composition in the style of his *Constellation* series. Further uses of algorithmic description of painting styles include mechanical storage, search and retrieval in art archives, attribution studies and diachronic studies of stylistic change.

I. INTRODUCTION

People usually are able to recognize the works of an artist by invoking a familiarity with the formal nature of the artist's style. For experts, this recognition ability may be based on exquisitely detailed knowledge of the plastic elements that the artist uses and on the artist's methods of juxtaposing these elements. Usually, this recognition ability exists at an unarticulated level, easy to invoke but difficult to explicate. This difficulty may be attributable to the lack of appropriate, powerful tools for expressing such insight. Exposition in a natural language, although it may be an adequate method for communicating such insight to other scholars, lacks the precision and completeness common in much scientific communication and particularly in communication with computers. In this paper, we show how to meet the computer's demands of precision and completeness in describing paintings. To test the validity of our descriptions we can then use computer methods that analyze and synthesize compositions. We discuss the painting styles of the contemporary California painter Richard Diebenkorn and the Spanish artist Joan Miró. For Diebenkorn, we exhibit a grammar for the linear facture of his *Ocean Park* series, which enables us to analyze existing compositions and synthesize new ones purportedly in the same style. For Miró, we attack the more difficult problem of shape and suggest that similar methods may be used for describing his *Constellation* series.

It is interesting to describe human creative behavior to computers both for the possibility of simulating that behavior and for the insight that results from the demands for precision and completeness required by the computer. In the description of language behavior, computer science has contributed significantly to our understanding of how natural language is produced and comprehended. This subfield of artificial intelligence has been maturing for over three decades. About two decades ago, we showed that the kind of formal models that had demonstrated their power for language description could also be used, with suitable modification, for image description [1]. The resulting field of syntactic image processing thereafter developed many formal tools for looking at pictures and diagrams in ways analogous to natural language description. Despite the proliferation of these tools, it was not until the late 1970s that they were harnessed to examine important creative activity. A significant step occurred in architectural design when it was shown that a shape grammar could describe, and thereby partially explain, the house plans of Frank Lloyd Wright [2]. This grammar provided a method for analyzing Wright's architectural style. Then, to test the validity of the grammar, new non-extant examples were generated from the rules. The resulting designs were remarkably similar to Wright's own work. Shape grammars have also been applied to other aspects of architecture, such as buildings [3], fenestration [4], furnishings [5] and landscape [6].

Architecture lends itself to this kind of formal approach because of the inherent physical constraints involved in building and because of the discipline's highly developed system of symbolic representation. In turn, it seemed natural to ask whether such methods of formal description could be extended to the fine arts, which do not have an obvious system analogous to architectural drawing. We answered this question affirmatively by showing that a grammar could be written for Diebenkorn's linear compositions [7, 8].

II. A GRAMMAR FOR THE LINEAR COMPOSITIONS OF RICHARD DIEBENKORN

Our grammar uses modified shape grammar rules to describe the composition of a corpus of some 140 major paintings by Diebenkorn, his *Ocean Park* series. The organization of these roughly geometric works is defined by tightly related horizontal, vertical and diagonal lines, with the pentimenti (partially overpainted sketches) figuring significantly in the final network. Though color and its implication of space characterize Diebenkorn's work, we have confined ourselves in this grammar to the linear composition only.

The Diebenkorn grammar is shown in Fig. 1 [9]. The 42 rules shown are an

Joan L. Kirsch (art historian, printmaker), Sturvil Corporation, P.O. Box 157, Clarksburg, MD 20871-0157, U.S.A.

Russell A. Kirsch (computer scientist), Sturvil Corporation, P.O. Box 157, Clarksburg, MD 20871-0157, U.S.A.

Manuscript solicited by Harry Rand.
Received 14 October 1987.

Pergamon Press plc.
Printed in Great Britain.
0024-094X/88 $3.00+0.00

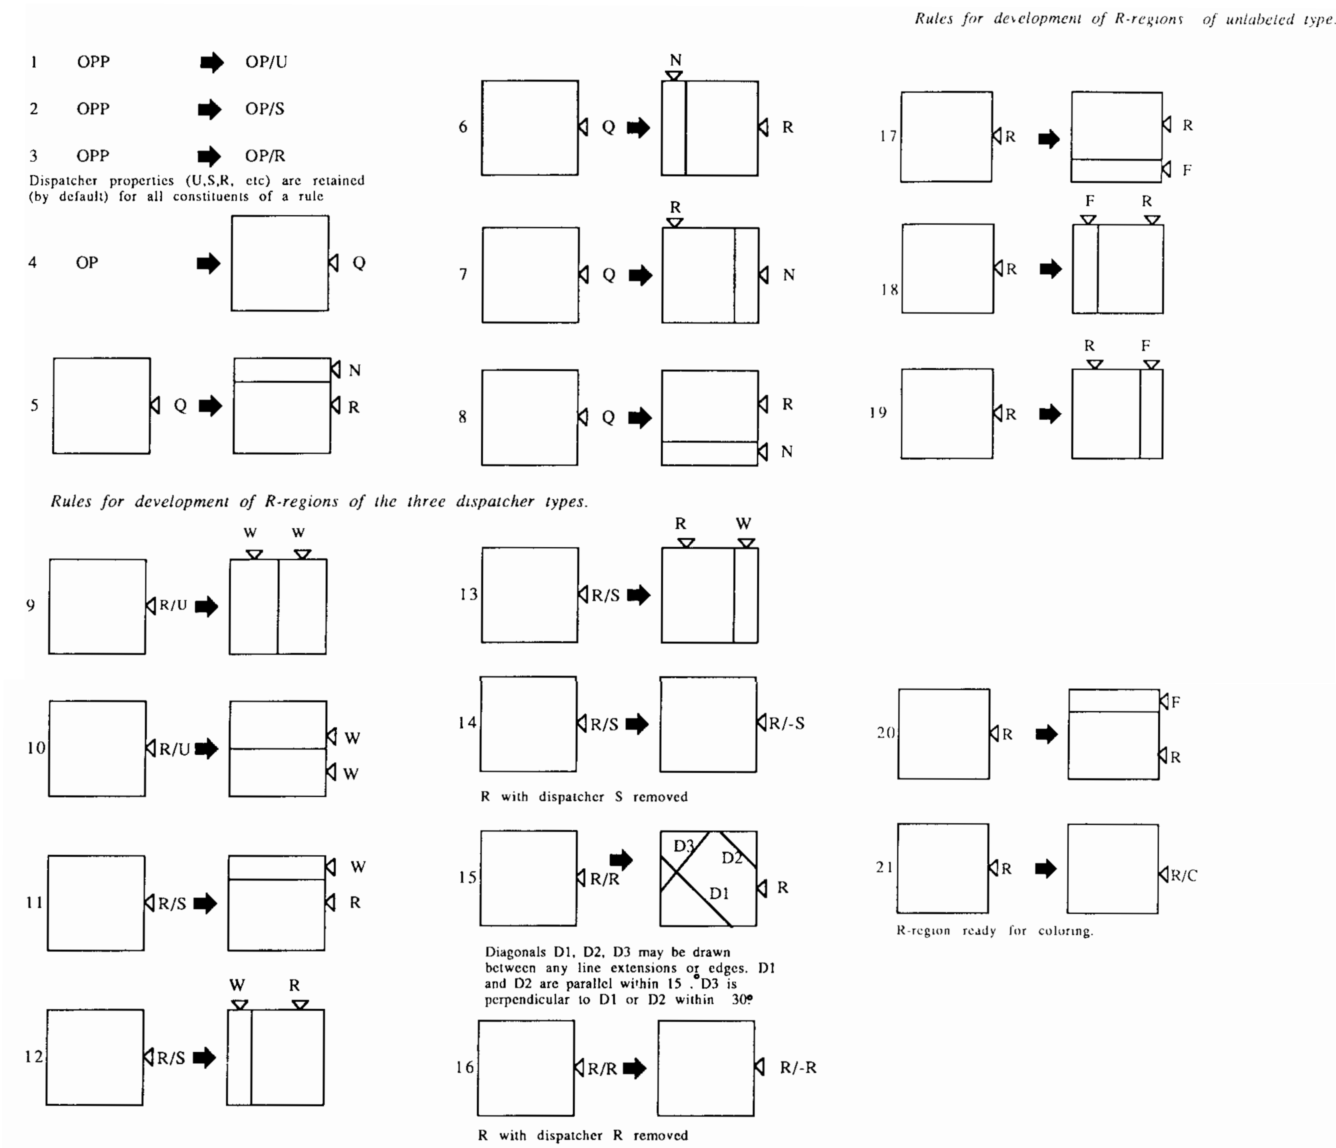

Fig. 1. A grammar for the *Ocean Park* paintings of Richard Diebenkorn.

abbreviation for about three times as many so-called 'productions'. Each production rule describes how a region of the composition may be elaborated, usually by subdividing it to produce other regions. Thus, rule 11 describes how a region labeled with an R/S (corresponding to a 'suburban' part of Diebenkorn's landscape) may be subdivided into two regions, a W region and, recursively, another R region. Rules 11 through 14 show the four alternative ways that an R/S region may be subdivided. The last of these rules, the only non-recursive one, allows the process to terminate and other rules to be invoked to expand the resulting W region.

The rules of the grammar are thus seen to allow optional ways to organize the composition, providing an infinite number of distinct compositions. The grammar is a description of structure in a succinct form. What remains to be shown is that such structure is appropriately attributable to the intended corpus of the *Ocean Park* series. For this, we may consider Diebenkorn's *Ocean Park No. 111*, painted in 1978 (Fig. 2). We can generate the linear structure of this composition by applying a sequence of rules beginning with the symbol OPP for an undifferentiated composition. First, by choosing rule 2, we get the 'suburban' option OP/S. Then, by applying rule 4, we get the otherwise unorganized picture shown in the blank rectangle of Fig. 3. Rule 7 produces the narrow band shown in the second step of the figure. The next step is produced by rule 11 (twice) followed by rules 12, 13 and 14. The fourth step is produced by the recursive application of rule 36 seven times, and the fifth step by rule 20 applied recursively seven times. The final step in Fig. 3 is produced by applying rules 38, 37 (twice), 26, 28, 27 (twice), 30 and 36 (twice). By this sequence of 32 rule applications, we arrive at the final composition, which may be compared with the original painting illustrated in Fig. 2. The similarity in structure is evident.

Such a test, especially when it is applied repeatedly to many examples from the corpus, is a useful corroboration of the correctness of the grammar. But we need a necessary and sufficient test to be assured that the grammar is correct. Sufficiency of the grammar is corroborated by using the grammar not for analysis, as above, but for synthesis. To test for sufficiency, we make random choices wherever the grammar permits. Though the resulting composition is unlikely to arise in the original corpus, it nevertheless may be judged by an informed observer for its consistency with the style of the corpus. It is, of course, the style that we are attempting to describe rather than any one sample of it evident in the finite corpus of actual works. Since, by definition, examples beyond the available ones do not exist, an

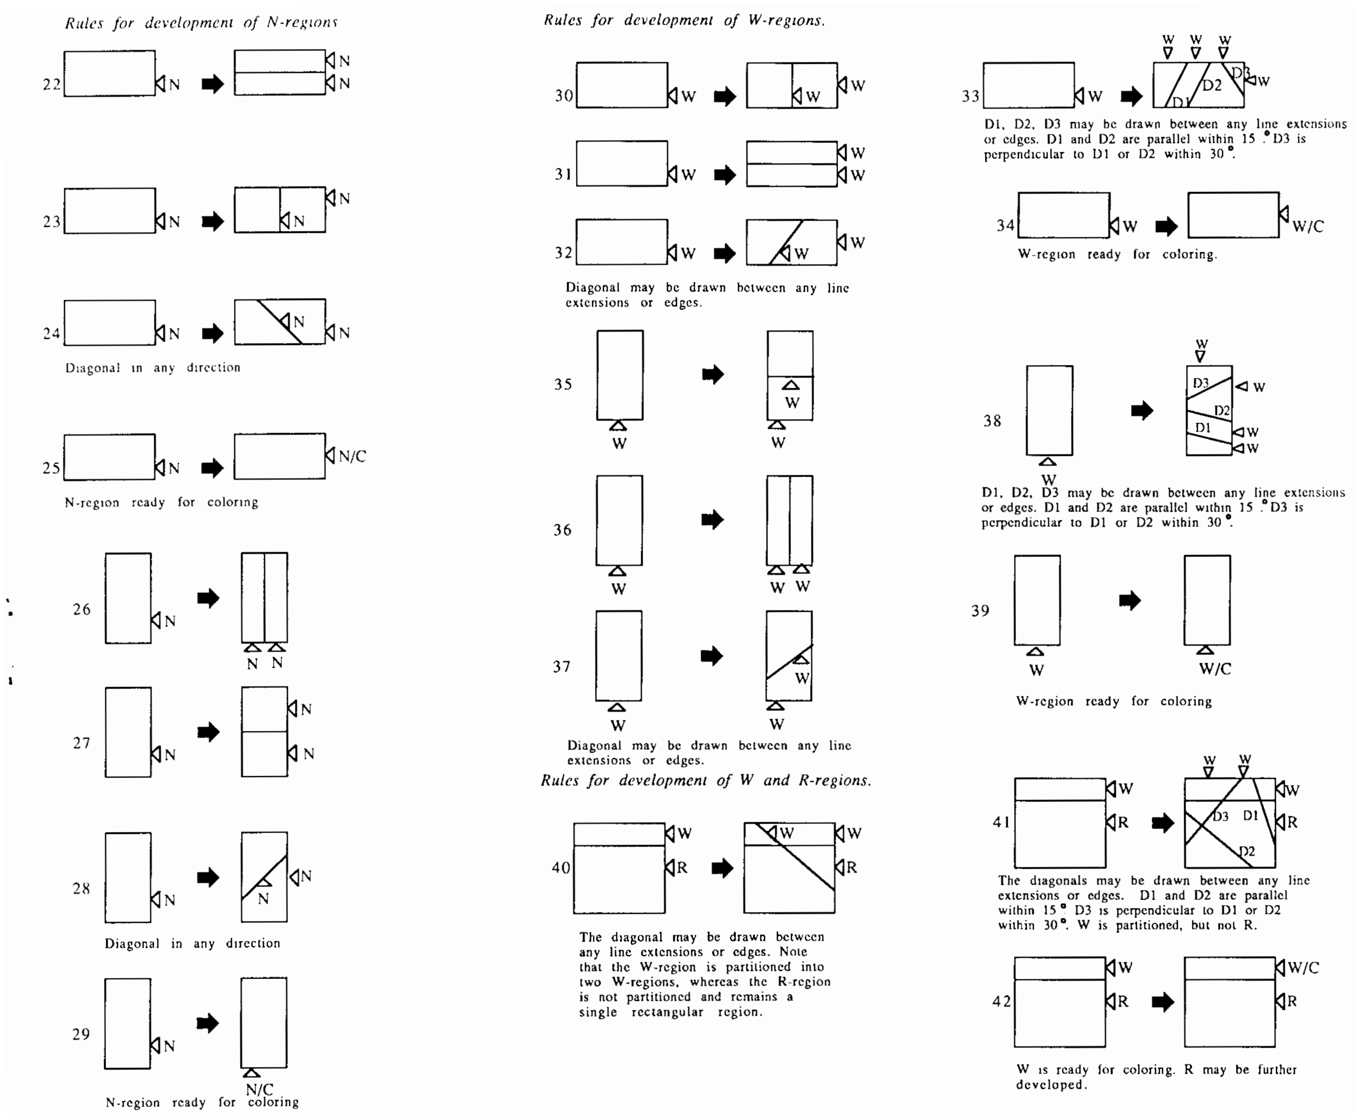

Fig. 1 (cont'd.)

informed observer is needed to judge whether the grammar is sufficient in terms of a plausible extension of the existing corpus of works. Figure 4 shows such examples randomly generated from the grammar. To those familiar with Diebenkorn's work, they appear to be plausible *Ocean Park* compositions. Diebenkorn, himself, when shown Fig. 4a, 'recognized' the composition. At first, we took this to be an indication of the plausibility of the composition. Only later, at a show of his current works, did we see a recent painting that was strikingly similar in composition to Fig. 4a. This painting, *Ocean Park No. 126*, painted in 1984, is shown in Fig. 5.

III. LIMITATIONS OF GRAMMATICAL DESCRIPTION: WHO SHOULD DESCRIBE A PAINTING?

A work of art stands for itself; nevertheless, for many purposes, surrogate descriptions of an artwork are necessary. Such purposes include criticism, cataloging, pedagogy and other scholarly verbal exercises that necessarily invoke language to refer, descriptively, to an artwork. Art history feeds on debate regarding the status of these descriptions and their degree of adequacy in characterizing the artwork. We do not intend to enter this debate here. Rather, we propose to discuss a new issue concerning the description of artworks. The issue arises from the demonstrated possibility of algorithmic description of artworks with computers. No descriptive technique, heretofore, has presented the power of, or allowed the concomitant complexity of, computer descriptions.

Many artists today are using computers to produce serious artworks. Were these efforts merely the attempt by computer scientists to demonstrate a new technology, the results would not have to concern art historians, critics and other scholars in the fine arts. But it is now fairly common for well-trained artists to adopt this new technology as a new medium. And the way in which the medium is used raises the issue of multiple descriptions. An artist who writes a program to produce an artwork is actually producing an unusually complete syntactical description of that artwork. In fact, many such programs (or algorithms) are capable of generating an unlimited number of instances of the artist's intent, each distinguished by some simple set of parameter values that direct the algorithm to choose one path rather than another. All such instances are extraordinarily complete descriptions of the artwork at a formal level, sufficiently complete not only to produce the artwork itself but to distinguish it from all others.

Not all artists using computers reduce their work to complete algorithmic descriptions. One who does is Harold Cohen [10]; his medium is a computer-programming language, and the resultant drawings produced by his Aaron program are instances of his algorithm. Although the drawings are, of course, of visual interest, there is a significant disparity between what viewers/readers can learn

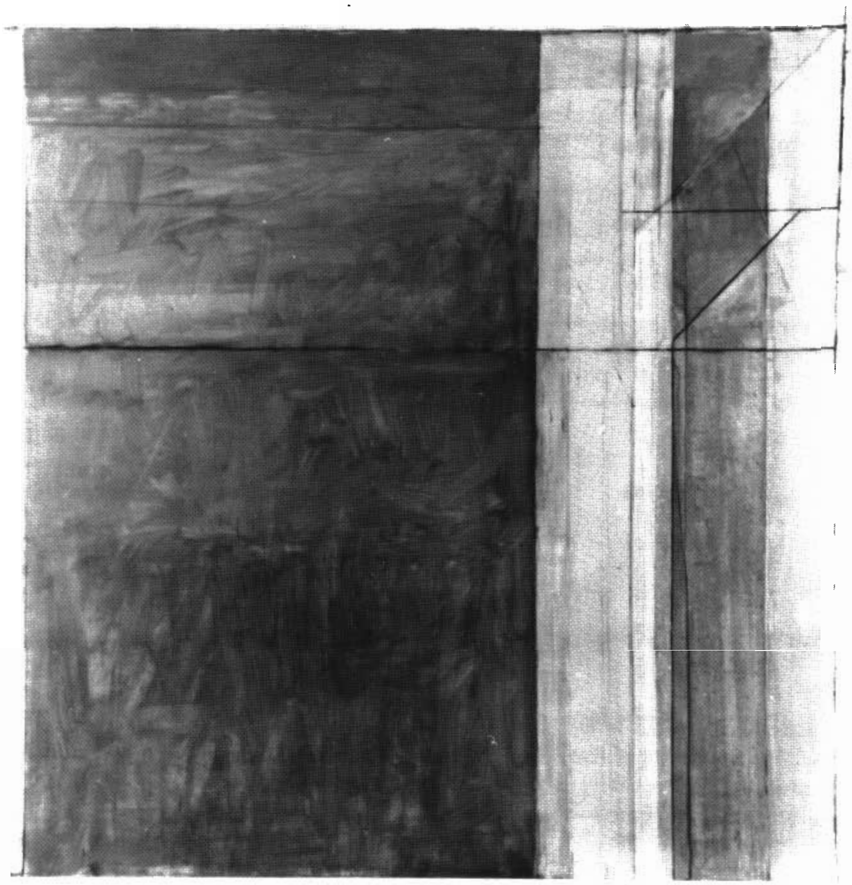

Fig. 2. Richard Diebenkorn, *Ocean Park No. 111*, oil and charcoal on canvas, 336.2 × 336.7 cm, 1978. (Courtesy of the Hirshhorn Museum and Sculpture Garden, Smithsonian Institution, Museum Purchase, 1979.)

from one of the drawings and what they can learn from studying the algorithm.

First, the algorithm provides a scheme for unlimited possible drawings, whereas any single drawing is only one such instance. But sometimes a description of a large class is less perspicuous than a single instance of that class. Different viewers/readers will have different preferences.

A second disparity arises from the 'emergent properties'. Any drawing will suggest many unique or fortuitous relations pertaining to that work's particular configuration. While these properties need not appear anywhere in the algorithm, they are consequences of other decisions made in the generating process: they are emergent properties. The artist who wrote the algorithm probably did not anticipate these emergent properties and, if he or she has not in fact seen the drawing containing certain emergent properties, may not even be aware of them. So the viewer observing the drawing generated by the computer is seeing something not explicitly intended by the artist. Such a possibility is much more remote when reading only the algorithm.

We can see that both an algorithm and its generated artwork are good descriptions of the artist's style, one as a class of works and the other as a specific example. Of

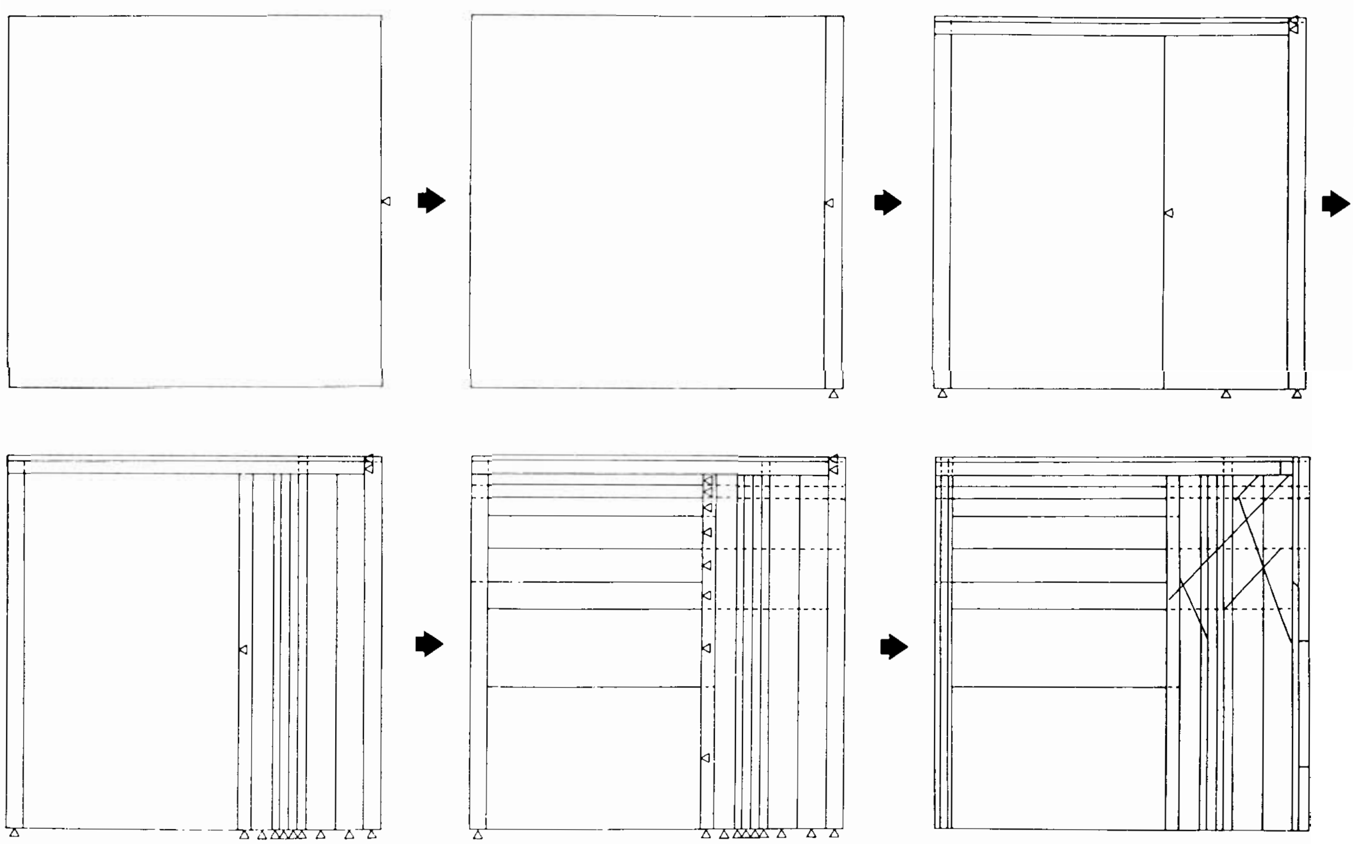

Fig. 3. Steps in the generation of Diebenkorn's *Ocean Park No. 111* from the grammar.

course, the algorithm that is written by the artist supplies a better description of the artworks than one written by an observer.

But is the issue of best description limited only to those few works being produced on computers? The Diebenkorn grammar described above has shown that the same questions arise for art done in traditional media. We have shown that scholars studying paintings can construct computer algorithms similar in purpose to the algorithms written by computer artists.

When it is a scholar, not the artist, who produces an algorithmic description, the issue of emergent properties cannot be ignored. Since scholars generally are not privy to whatever (informal) algorithm was actually used by the artist, they can only infer from examples what such an algorithm might be. Such an inference is unavoidably corrupted by whatever emergent (and unintended) properties are found in the oeuvre from which the algorithm was built. The scholar must be sensitive to the difference between accidental and systematic properties found in the examples studied. Inevitably, the algorithm written by a scholar to describe a conventional painting is a complete description, at least as complete as the algorithm written by the artist who works on the computer—complete, but maybe wrong!

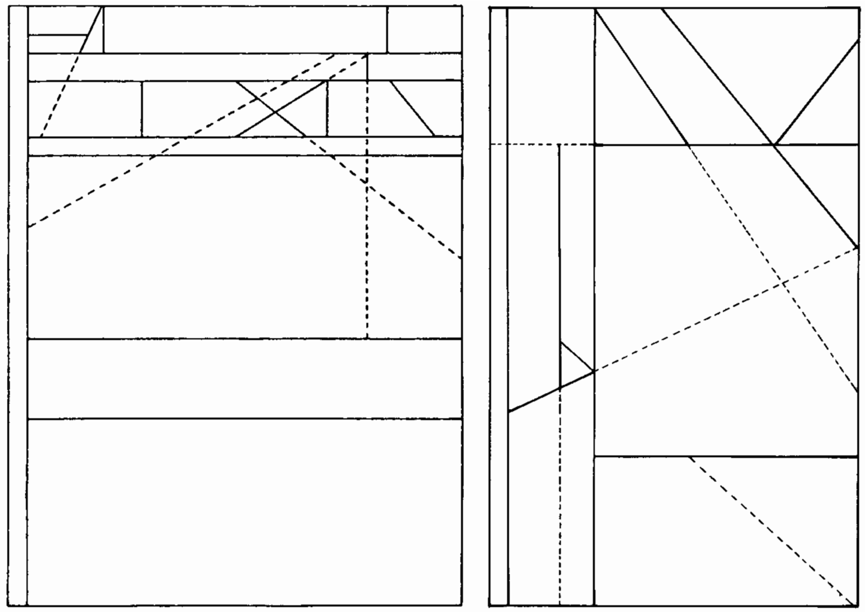

Fig. 4. Two pseudo-Diebenkorn compositions generated at random from the grammar.

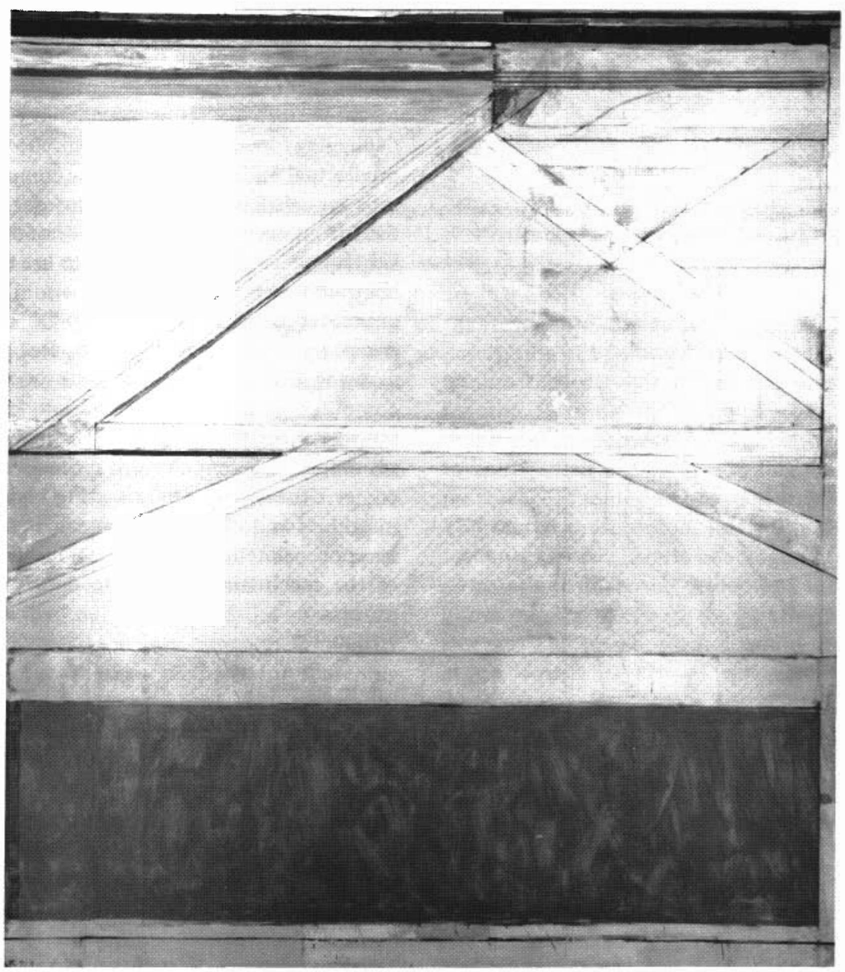

Fig. 5. Richard Diebenkorn, *Ocean Park No. 126*, oil on canvas, 236.2 × 205.7 cm, 1984. (Collection of Donald and Barbara Zucker, courtesy of M. Knoedler and Co., Inc., New York).

IV. MORE COMPLEX DESCRIPTIONS OF OTHER STYLES

We should clarify that a grammar for Diebenkorn's linear compositions corresponds to some inherent limitations in machine representation. A computer can easily accept data about points and their connecting lines. This makes it possible to deal with Diebenkorn's essential or deep structure. But the many other inevitable and inimitable aspects of his paintings, such as color, texture or line quality, present representation difficulties beyond the current scope of our investigation. Furthermore, in studying Diebenkorn's linear compositions, it was possible to ignore shape.

However, for many painters, space is fundamentally demarked by shape. As the next step in our research, we chose to study Joan Miró, who, in 1940–1941, painted a series of 23 gouaches called *Constellations*, which are complex allover compositions based on suggestive and imaginative biomorphic forms. Again, we made many simplifying assumptions, but this time we concentrated primarily on shape in studying Miró's style. We started by drawing on the computer a set of prototype shapes (Fig. 6) obtained from reproductions of Miró's *Constellations*, which we scanned on the computer. Then, aided by a computer-graphics program, we manually traced the outlines of the shapes. We allowed each of these shapes to be modified in three ways: by uniform horizontal and vertical stretching, by size variation and by rotation. This produced an expanded dictionary of prototypes to be used in analyzing Miró's compositions.

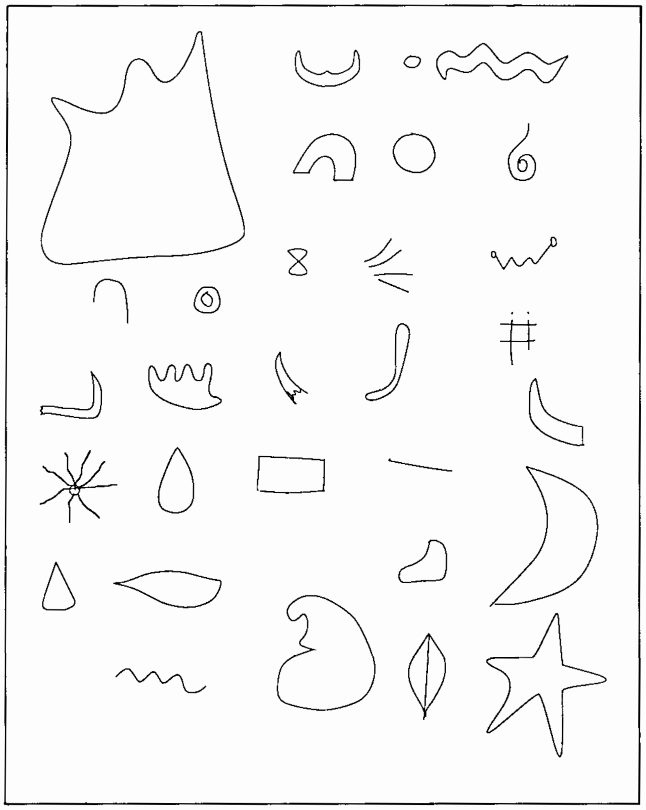

Fig. 6. A dictionary of prototype shapes used in generating pseudo-Miró compositions.

Preliminary to writing a grammar for a painting style, a period of intense and open-minded scrutiny of the works is necessary. For Diebenkorn, we made an analysis of the target compositions in order to test our understanding of his style. For Miró, we chose to begin with synthesis. At the current stage of our research, our understanding of Miró's style has not yet been formalized in a grammar, but we still may use computer methods to aid in the synthesis test. By using a computer to make the transformations on prototype shapes, but retaining manual control of the computer, we synthesized several compositions purportedly in the style of Miró's *Constellations*. One can compare one such synthesis (Fig. 7) with an actual Miró *Constellation* (Fig. 8). Although this test is not directed at validating a synthesis algorithm, it is nevertheless a useful initial step in verifying the informal understanding that the scholar may have. These insights underlie the eventual formal grammar or computer algorithm.

At this point it is useful to explain that the process of discovering an algorithmic basis for explaining an artist's style is neither inexorable nor does it suggest mechanization. It requires knowledge, assumptions and hard work. Perhaps after many successful attempts at writing grammars and other algorithms for explaining style have been made, the process will become sufficiently routine as to suggest mechanization. But at this time, it appears to us that serious fine art is too rich and varied to allow a mechanical discovery of its structure. The computer participates in this process only by offering scholars a powerful, expressive tool, which processes their insights and allows them to explore the consequences of those insights. It is an intellectual debugging tool that encourages ambitious intellectual undertakings. But the computer's knowledge is cumulative, insofar as data from each study is incorporated in subsequent studies. If one were to use the computer merely as a tool, one would be guaranteeing that the ability of the computer would remain static. But by assuring that results accumulate in the computer in a prescribed fashion, one allows for the possibility that the insight residing in the computer will successively come closer to approximating one's insight. If and when this happens, it will be appropriate to reexplore the question of the mechanization of the discovery process.

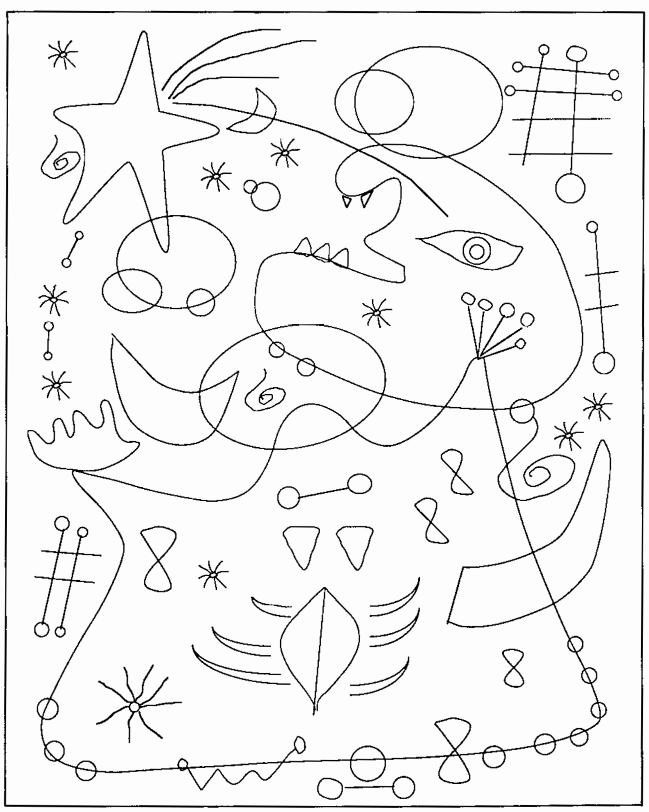

Fig. 7. A pseudo-Miró composition generated from the set of prototype shapes in Fig. 6.

V. USES FOR ALGORITHMIC DESCRIPTIONS OF STYLE

The discipline imposed by using grammars or other algorithmic methods to describe a style to a computer results in a description that can have a clarity transcending that usually obtained with more discursive methods. Furthermore, the consequences of such a description can readily be seen with the aid of the computer. The benefits for scholarship

Fig. 8. Joan Miró, *The Beautiful Bird Revealing the Unknown to a Pair of Lovers*, gouache and oil wash on paper, 46 × 38 cm, Montroig, 23 July 1941. (Collection of The Museum of Modern Art, New York. Acquired through the Lillie P. Bliss Bequest.)

are obvious. But there are several additional uses that such algorithmic description can provide, both of a practical and of an intellectual nature.

Art archives are generally collections of objects that are surrogates of the original work. Though these surrogates may range from models to photographs to textual descriptions, the textual descriptions are always the key to the storage and retrieval of the surrogates. The reason is our common familiarity with systematic methods for searching through texts with tools that use a range of technologies, from the manual to the electronic. Strangely, text is the weakest of the media for describing art objects since it imposes an unnatural linear form on art objects, which are two- or three-dimensional. Storage and retrieval would be more efficient if the structure of the information used in searching corresponded more naturally to the structure of the things it described. Here is where the methods discussed above are of possible significant use. By storing a structural description of a painting in the form of its grammatical derivation, one can search for the painting in terms of a large number of properties that would be hard to anticipate if one were encoding them for storage in some conventional archive. The search prescription, in the kind of system we foresee, would be a fragment of a painting analyzed with the same grammar used to encode the painting for storage. The search method would then match structures just as we currently match text words when searching for documents. Of course, such a system could not be built until the whole corpus of art objects had been described to the computer with the methods we have presented. For an eclectic collection of art objects, this appears impractical; but for a relatively homogeneous collection, the possibility of describing the collection in algorithmic fashion is a practical alternative, which offers powerful searching capability.

Another use for the algorithmic methods we discuss here occurs in aesthetic theory. For a long time, people have been associating quantitative measures with art objects. The motivation for such interest probably comes from the supposition that quantitative methods, so successfully adopted in the sciences, can bring similar benefits to the arts. But even in the sciences, recent technological advances in areas like computer graphics have shown that non-quantitative tools can sometimes be more powerful than the more arithmetic ones. And in the arts this is surely true. It is usually assumed that, by forgoing quantitative methods, one thereby forgoes the possibility of precision and reproducibility of results. However, we have seen that the grammatical methods described above allow for both precision and reproducibility comparable to the standards usually met in scientific investigations, except there are no numbers! The measures are expressed in data structures, like tree diagrams, and the operations performed on them are the analogs of the arithmetic operations commonly performed on numbers. So measurement is not lost but rather transformed in such a way as more faithfully to mimic the art process being measured.

We can see a simple example of the use of such a measure by reexamining our derivation of Diebenkorn's *Ocean Park No. 111*. The description of the derivation shown in Fig. 3 furnishes a measure, which is the sequence of rule applications used in producing the composition. This data structure (a list) contains 32 items (the rule instances used). If we wish, we can perform useful operations on this data structure. For example, we can obtain a conventional information measure from it. Each rule application is chosen from a set of several options. If there are r(i) rule options available when the ith rule is chosen, then by summing the logarithms to the base 2 of the r(i) for all the 32 rules used in the derivation, we arrive at the common (binary) information measure for the composition. For *Ocean Park No. 111*, this value is 67.6 bits. In other words, the memory needed to specify *Ocean Park No. 111* with respect to the grammar of Fig. 1 is about 68 bits. To store a scanned image of the same painting on, for example, a video disk, would require about ten thousand times as much memory. And although such expensive storage would allow one to retrieve an image of the painting rather than a diagram, the diagram could be 'understood' by the computer whereas the scanned image could not. We speak of understanding here in the sense that the computer could answer questions about the structural description given by the grammar but not about a scanned image.

Connoisseurship relates to another use of algorithmic descriptions of style as, for example, in attribution studies. Many disciplines contribute to attribution analyses, ranging from those that draw upon materials science to those that are based on iconography. Falling somewhere within this range is the practice of stylistic analysis, which draws on formal properties of the work under study as determinants of attribution. Without presuming to judge how strong a contribution such formal analyses make to attribution, we believe that they can be enhanced significantly by the use of algorithmic descriptions. Not only can algorithmic descriptions be of specific use in single cases, but they can also be cumulative. Once prepared, an analysis can be invoked for subsequent studies without recapitulating the whole process. We are suggesting, of course, that the same type of cumulative process that serially allows scientific studies to build on each other can benefit connoisseurship, at least to the limited extent that formal analysis is of use.

Even though we have discussed algorithmic grammatical methods in the context of describing existing art, the studio artist can also use some of these techniques. When we presented Diebenkorn with the Diebenkorn grammar, he immediately questioned whether the sequential process whereby the grammar generates a composition could be varied. He inquired about the possibility of using the analysis of an existing composition and following optional paths not traveled in the actual composition. This can be done readily and offers endless development of many facets of a picture. In general, whether an artist is given a grammar derived post hoc by a scholar or whether the artist generates the grammar in the initial creative process of inventing a painting, he or she can pursue a large number of different compositional derivations with the help of a computer.

An interesting issue relates to the criteria used by the artist for selecting some compositions within a style and rejecting others. Perhaps semantic grounds form the basis for such a choice. Perhaps the criteria are still more complex than would be suggested by intentionality. In any case, the formal definition of a style given by a grammar provides a rich basis for exploring such questions, both from the viewpoint of the studio artist and from that of the scholar.

A study by T.W. Knight [11] suggests yet another use of formal algorithmic analyses. By making a diachronic analysis of the works produced by two De Stijl artists during two different periods of over two decades, she was able to show transformations in the grammars that defined the artists' styles over these periods. The transformations differed significantly. One artist's grammar systematically added more rules over a 21-year period, while the other's successively subtracted rules during a later 20-year period. It is tempting to explain such transformations as owing to the maturation process of artists. In this case, one artist became more constrained, whereas the other became less so. The important idea here is that grammars can be serious objects of study as representations of the styles they purport to describe. And issues that can be resolved on the basis of analysis of styles can be resolved by analyzing the grammars of those styles.

VI. CONCLUSION

We have shown how the art historian's insights into the formal elements of a painter's style can be described to a computer. This description in the form of a grammar or other algorithm can be validated by analysis and synthesis tests. Once validated, the grammar becomes a machine-manipulable form of stylistic description with many uses. These include mechanical storage, search and retrieval in art archives, various measurements of artworks, attribution, creation of alternative compositions in a fixed style and diachronic studies of stylistic change. But most importantly, we have shown how a scholar's insight into the nature of style can be made more readily available to others for probing still further into the mystery of art.

REFERENCES AND NOTES

1. R.A. Kirsch, "Computer Interpretation of English Text and Picture Patterns", *IEEE Trans. Electronic Computers* **13**, 363–376 (1964).
2. H. Koning and J. Eizenberg, "The Language of the Prairie: Frank Lloyd Wright's Prairie Houses", *Environment and Planning B: Planning and Design* **8**, 295–323 (1981).
3. F. Downing and U. Flemming, "The Bungalows of Buffalo", *Environment and Planning B: Planning and Design* **8**, 269–293 (1981).
4. G. Stiny, "Ice-ray: A Note on the Generation of Chinese Lattice Designs", *Environment and Planning B* **4**, 89–98 (1977).
5. T.W. Knight, "The Generation of Hepplewhite-style Chair-back Designs", *Environment and Planning B* **7**, 227–238 (1980).
6. G. Stiny and W.J. Mitchell, "The Grammar of Paradise: On the Generation of Mughul Gardens", *Environment and Planning B* **7**, 209–226 (1980).
7. J.L. Kirsch and R.A. Kirsch, "The Structure of Paintings: Formal Grammar and Design", *Environment and Planning B: Planning and Design* **13**, 163–176 (1986).
8. J.L. Kirsch and R.A. Kirsch, "Computer Grammars for the Syntactical Analysis of Paintings", in I. Lavin, ed., *World of Art: Themes of Unity in Diversity, Acts of the XXVIth International Congress of the History of Art* (University Park: Penn. State Univ. Press, 1987).
9. The Diebenkorn grammar is described in detail in Kirsch and Kirsch [7]. The heavy arrows in the grammar denote replacement operations in the production rules. The light triangular arrows are used by labels to point to the regions they denote. The same convention is used in Fig. 3.
10. H. Cohen, B. Cohen and P. Ni, *The First Artificial Intelligence Coloring Book* (Los Altos, CA: Kaufmann, 1984).
11. T.W. Knight, "Transformations of De Stijl art: The Paintings of Georges Vantongerloo and Fritz Glarner", *Environment and Planning B: Planning and Design* (forthcoming).