AI & Soc (1991) 5: 183–201


**AI & SOCIETY**

# Can You Make a Computer Understand and Produce Art?

Roberto Maiocchi

*Dipartimento di Elettronica, Politecnico di Milano, Piazza L. da Vinci 32, 20133 Milano, Italy*

**Abstract.** Although artificial intelligence techniques have been successfully applied to reproduce many rational features of human behaviour, a great barrier has been encountered in simulating human activities where intuition and emotion are involved. Art making and viewing are processes where typically rational and mechanical aspects interact with aesthetic and cognitive criteria. Can you make a computer understand and autonomously produce art?

The main purpose of this paper is to present the most relevant approaches in the study of art perception and creation via computer, focusing on the results achieved in artistic computer graphics.

**Keywords:** Artificial Intelligence; Computer Art; Art perception and creation; Aesthetics; Syntactic and semantic properties of artworks

## 1. Introduction

A machine, stated Turing in a famous paper on the nature of intelligence (Turing, 1950), can be said to think if it can reproduce the behavior of a thinking person limited to what can be transmitted by means of a teletypewriter. And if any machine can be constructed to "think", the digital computer is the machine. Research in Artificial Intelligence (AI) has been founded on the dogma that "mental activities are information processing" (Barr, 1983), more specifically that the manipulation of symbols (representational data structures) by suitable computer programs is no more and no less than what minds do.

Many features of human behaviour have now been successfully emulated by means of programs, mainly modeling rational aspects of our mind. Achievements in game playing, natural language understanding, and problem solving are the striking evidence that "computers can think". But much of our conscious intellectual life is based on intuitions and emotions that are less appropriately considered under a logical regimen, and actually AI and disciplines such as

Cognitive Psychology have begun to overlap increasingly in the attempt to bridge the gap of the rational/emotional dichotomy (Anderson, 1984).

On the other hand, computer art represents a historical breakthrough in computer applications. For the first time computers have become involved in an activity that had been previously exclusive to the domain of humans: the act of creation. Since the mid 1960s, the computer has been largely used as a tool for producing art, and programs have been written to assist artists to perform various tasks during the process of making art. Surveys on the history of computer art can be found in (Franke, 1985; Dietrich 1966). A strong feeling of optimism has also been growing in the belief that computers can analyze and autonomously produce artworks. If we think of an artwork as a model of the real world as perceived and created by an intelligent artist (Apter, 1977), why shouldn't we be able to program computers to understand and produce art?

Since art making and viewing are processes where typically rational and mechanical aspects interact with aesthetic and cognitive criteria, adopting the AI paradigm to the study of such processes seems particularly appropriate. There are two areas in which AI techniques have been applied in dealing with art: perception and creation. *Perception* is concerned with the experience resulting from the confrontation with an object or a scene. Perceptive machines maintain a description of the object in its structural terms, e.g., its geometry, and a description of its aesthetic relevance. *Creation* is concerned with the representation of the intentions, intuitions, and skills of an artist. Knowledgeable machines produce artworks showing some aesthetic and semantic ideas.

The purpose of this paper is to give an overivew of some of the most relevant AI approaches in modeling the "kind of computation" performed by our brain both in art perception and creation, focusing on the results obtained in artistic computer graphics. Major efforts have also been produced in computer music (for a complete survey, see Roads, 1985), while just a few interesting applications can be found in poetry (Davison, 1982; Wilson, 1984), sculpture (Mallory, 1969), computer animation (Kahn, 1979; Zeltzer, 1983; Badler, 1989), and choreography (Stadler, 1982).

The organization of the paper is as follows; a brief discussion on the refractions of science into art is the topic of Section 2; achievements in artistic computer graphics are discussed in Section 3; conclusions are drawn in the final section, and future advances are envisioned.

## 2. Introducing Computers into Art

The use of the computer for artistic purposes is the last and decisive step towards technologizing the arts. Whereas instruments have been used, for instance, in the area of music for a long time, their acceptance is much more recent in the visual sector, where the use of unconventional media is still strongly criticized. Indeed, computer art brings various problems to the foreground, both when the computer is employed just as a tool to produce artistic effects and when it is instead used to explore new art forms and aesthetic experiences.

The first question which needs to be answered when analyzing computer-generated artworks is where precisely the art is located, i.e., if the art resides in the program that generates the piece or in the output itself. Analogies with more established art forms have been discussed in detail in (Jankel, 1984). Computer art is sometimes like *photography*, in the sense that there is one original program, which can produce pictures on any number of runs, as there is a photographic negative and many prints. But when a program incorporates randomness, the output of one run is so different from another that the minor variations of photographic prints are insignificant by comparison. Perhaps, computer art is better compared with *painting*, for in the case of painting there is one image whose copies are reproductions of little intrinsic value. Yet it would be hard to identify in computer art an equivalent of the unique artifact produced by the painter's brush. The image on the screen of the computer is not unique, it appears any time the program is run. Again, the program is repeatable, i.e., it can be copied and run on any compatible computer.

*Music* might make a more logical comparison to computer art. A piece of music can be played many times, as a program can be run many times, and each performance may be more or less different depending on the musician (the user of an interactive computer program) and the instrument (the computer system). We can then conclude that art resides both in the program and in the individual realizations of it.

The speed of execution is a second factor differentiating computer art from conventional forms of art. A program may generate dozens of images in a few minutes while, for example, a painter, as prolific as he or she may be, needs several years to work on a theme. Furthermore, manual skills are no longer a precondition for engaging in art. This eliminates that close connection between the creating hand and the material which is considered to be so important by some art experts; the act of artistic creativity shifts from the manual to those areas which can be described as cerebral.

As a matter of fact, the most interesting applications of computer art involve programs that try to capture general laws of aesthetics. Two relevant aspects should be considered in programming the beautiful: the syntactic and semantic properties of an artwork. *Syntactic properties* reside unambiguously inside the object under consideration. Examples of syntactic properties of artworks are the chord sequence of a piece of music, or some geometric patterns of an abstract painting. Syntactic properties are localizable on the surface level of a piece, and once defined are easily describable through a computer program. On the other hand, *semantic properties* depend on some sort of inner meaning housed in the piece itself and raise mental mechanisms that cannot be consciously described. Different levels of interpretation can be established in every person depending on the circumstances that pull out different meanings, provoke different connections with previously memorized experiences, and generally evaluate all deep aspects differently.

Both these aspects have been considered by computer artists who have applied AI techniques in their research, focusing on two types of artworks – perception and creation. Perceptual and generative processes are actually linked in that they both look at the syntactic and semantic qualities of an artwork. So far, in fact, perceptive machines have attempted to give a description of the object in its structural terms, whereas knowledgeable machines have been concerned with the representation of the intentions and intuitions of an artist.

Whether we can actually make a computer understand and produce art, i.e., whether we can mechanize, and to what extent, perceptual and generative processes, has to be considered carefully before analyzing the results achieved so far in this field of research. One of the main theses developed by the philosopher and computer scientist Douglas Hofstadter in his books (Hofstadter, 1980, 1985) is that

"every aspect of thinking can be viewed as a high level description of a system (the brain) which, on a low level, is governed by simple rules . . . The image is that of a formal system underlying an informal system, one that can, for instance, make puns, discover number patterns, forget names, make awful blunders in chess, and so forth. This is what one sees from outside: its informal, overt, software level. By contrast, it has a formal, hidden hardware level which is a formidably complex mechanism that makes transitions from state to state according to definite rules physically embodied in it . . ."

In other words, Hofstadter's statement is that all brain processes, even those showing some degree of irrationality, are derived from a computable substrate: the neural level.

Now many programs which have been developed in AI research deal with, manipulation of images, formulation of analogies, confusion of concepts, blurring of distinctions, and so forth. This does not contradict the fact that they rely on the correct functioning of their underlying hardware as much as the brain relies on the correct functioning of its neurons.

When applying AI techniques to emulate artistic behaviors, the concept of *beauty* needs to be specified in some way. In fact, beauty is a very ill-defined notion involving qualities of intelligence such as learning, creativity, emotional responses, memory, a sense of self. But its appreciation is a brain process, just as proving a mathematical theorem or playing a game of chess. We can thus believe that once some significant features involved in art perception and creation have been formalized, we could have computers developing original thoughts or works of art.

Who should get credit when an AI program comes up with an idea that has not been explicitly implemented in its program? The human will certainly get credit for having invented the program, but not for having had the ideas produced by the program inside his/her own head. In such cases, the human can be referred to as the "meta-author", i.e., the author of the author of the result, and the program as the author. Programmers of such amazing programs must specify something similar to the symbols in our brain and their triggering patterns which are responsible for creating the syntactic and semantic properties of an artwork.

New skills are then required for computer artists. The traditional view of artists as illogical, intuitive, and impulsive in contrast to that of programmers as constrained, logical, and precise, that exactly reflects the separation between art and science in our society, needs to be overcome. If at the beginning of computer art artists usually cooperated with scientists because they did not have any programming expertise, later technological artists with hybrid capabilities also appeared. Artists have to be programmers themselves for not depending on scientists to provide the software tools to realize their aesthetic needs. The essence of computer art should be a balance of *psyche* and *techne*, of right and left brain expressions (Palyka, 1982).

In this view, computer art not only actively encourages the bringing together of the two cultures, the technical and the artistic, but also promotes the investigation of general processes of cognition. And if programs can be written based on some cognition principles, whether they are meant for graphics, poetry or music, computer artists could find it relatively easy to shift from one field to another, or to move into the multi-media sphere, extending their means of expression.

In conclusion, the question arising from computer art is not the replacement of the coventional methods of artistic creation with electronics or a machine. Rather it makes sense to use every possible means to extend the range of artistic expression. The art of every age has used the means of its time to give form to artistic innovation, but the demands of a new medium have to be understood in terms of previous forms of artistic expression before the medium can be freely utilized by the artist and properly appreciated by the viewer. For instance, photography had to rival painting before it could take up a distinct space in the spectrum of the visual arts, and computer pictures have to relate to photography, performance and painting before they can establish their own unique ground. Only when such aesthetics has been developed will there be room to comment on whether, for example, a picture is really a graphic illustration, a piece of fine art, or an avant-garde experiment. Meanwhile, why should not the computer to used as a medium and instrument of art?

## 3. AI And Artistic Computer Graphics

Computer graphics is the creation, storage, and manipulation of models of objects and their pictures via computer. Computer graphics is the most important mechanized means of producing and reproducing pictures since the invention of photography and television, with the added advantage that it allows the representation of abstract, synthetic objects. The applications of computer graphics range from data plotting in business, science and technology to cartography, design, simulation and animation, process control, and office automation. Computers have also been used by visual artists as a tool for producing aesthetically pleasing pictures, raising a debate in the art community on whether computer-generated images can be considered artistic and on the definition of the theoretical foundations of computer art.

Since the beginning, visual artists have also tried to create autonomous art-making programs. The employment of random number generators has been the first technique used with the purpose of generating many different images from one program, introducing change with the selection of certain parameters to define, for instance, location, type, or size of a graphic element. Random numbers served to break the predictability of the computer, but simulated intuition in a very limited way.

Random numbers have also been used in relation to aesthetic rules derived from an analysis of traditional paintings. For example, the mathematician and computer artist Michael Noll took a Mondrian painting, "Composition with

Lines", an abstract, geometric study with seemingly random elements, and from it he extracted some statistics concerning the patterns (Noll, 1966). Given those statistics, he programmed a computer to generate numerous pseudo-Mondrain paintings having the same or different values of these randomness-governing parameters (Fig. 1). He then, showed the results to naive viewers. The reactions were interesting, in that more people preferred one of the pseudo-Mondrains (Fig. 1a) to the genuine Mondrian (Fig. 1b).

This quite amusing fact proves that a computer can certainly be programmed to imitate mathematically capturable stylistic aspects of a given work. Randomness is actually an indispensable ingredient of creative acts, but human creativity does not simply rely on such arbitrary sources. The essence of any artistic act is not just a selection of particular values for certain parameters, rather it is in the balancing of a myriad of intangible and mostly unconscious mental forces, a judgmental act that results in many conceptual choices that eventually add up to a measurable work of art. What is important is the making of an artwork and not the object itself.

Regarding artworks as models created by artists to represent chosen aspects of the real world, in a simplified and selective manner, (Apter, 1977; Cohen, 1987), the implementation of an autonomous program producing art implies the investigation of the cognitive process of representation of the artist, the understanding of his/her model, and the simulation of the artist's representational acts. In other words, an autonomous art-making program has to exercise something very like "human intelligence".

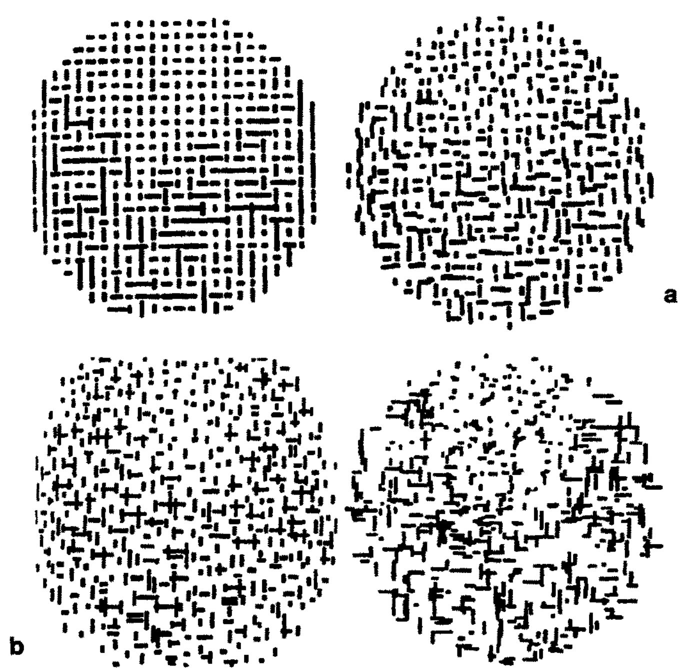


**Fig. 1.** One genuine Mondrian plus three computer imitations (from Hofstadter 1985. Copyright © 1985 by Basic Books, Inc. Reprinted by permission of Basic Books, Inc., Publishers).

AI techniques have already been extensively applied in the field of image processing which treats the converse process of computer graphics: the analysis of scenes or the recontruction of two- or three-dimensional objects from their pictures. Work in this field is focused on the study of certain highly complex tasks of the human visual data processing system. Image processing becomes considerably more difficult when perception of structural features or even semantic classification is required.

In the following, an overview of three relevant studies on art perception and understanding will be given. First, the formal grammar approach proposed by Joan and Russel Kirsh to analyze and reproduce the structure of the paintings of Richard Diebenkorn, a contemporary American painter (Kirsh, 1985, 1986), will be presented, then the problem of representing aesthetic judgments and how this knowledge can be used to create an aesthetic object as addressed in (Gips, 1975) and in (Mazlack, 1981) will be discussed.

A completely different perspective in the role that should be played by computer artists in computer art has been proposed by Harold Cohen with his AARON, a knowledge-based program designed to investigate the cognitive principles underlying visual representation (Cohen, 1979, 1985, 1987, 1988). For Cohen, one of the computer artist's main functions is to teach the computer how to make art by programming: the artist becomes a "meta-artist", and the artist is the computer itself. Under continuous development for fifteen years, AARON is now able to make "free-hand" drawings of people in garden-like settings. Cohen's AARON, the most important example of autonomous program able to "create", is presented at the end of the present section.

### 3.1 The Structure of Paintings

The main purpose of the work by Joan and Russel Kirsh is to investigate the possibility of understanding significant work in art. The formal analysis of an artwork is intended to deal with its visual properties such as color, line, shape, materials and their arrangements, the so-called plastic elements. Excluded are the extrinsic qualities of the work, i.e., feelings, stories, and metaphors.

The approach followed by the authors is the description of the compositional structure of the design of an artwork by means of a *formal grammar*. Borrowing from structural linguistic terminology, they distinguish between a deep structure and a surface structure (Chomsky, 1965). The *surface structure* accounts for many of the observable properties of the finished work, including, for example, texture, variation of media, line quality, and colors and their relationship. The *deep structure* can account for the overall composition and how the work is organized in two or three dimensions. The deep structure can also account for the interesting property of recursion which occurs in certain paintings.

In (Kirsh, 1985, 1986), the authors present a grammar specifying the deep structure for the paintings of Richard Diebenkorn; the choice has been made because the paintings are geometric and appear to be conventionally describable and measurable. The grammar devised consists of a set of production rules in a form similar to both Stiny's shape grammars (Stiny, 1980) and to context-dependent phrase structure grammars (Fig. 2).

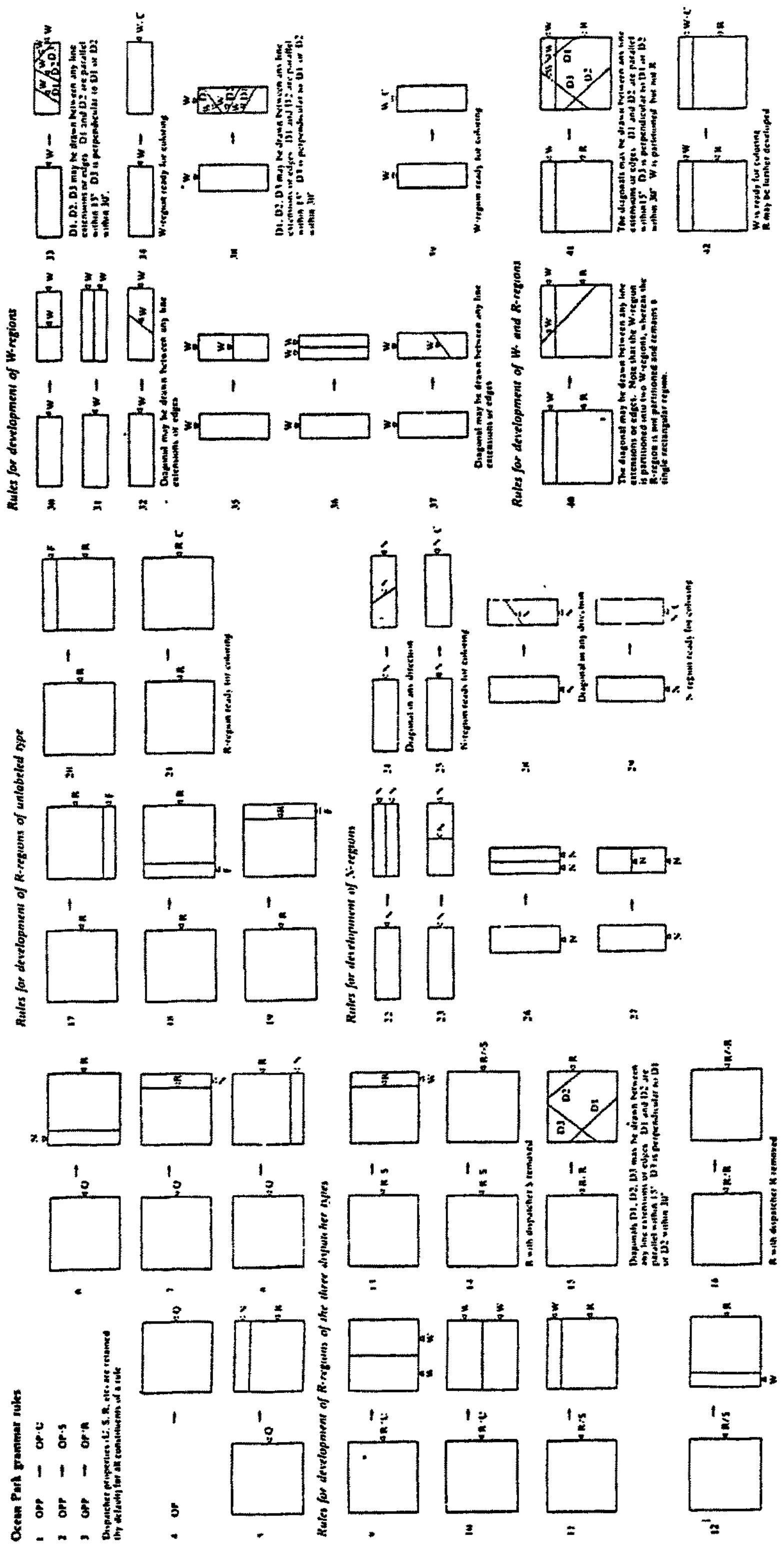

**Fig. 2.** Ocean Park Grammar rules. (from Kirsh 1986. Copyright © by Pion Ltd. Reprinted by permission of Pion Ltd).

As in shape grammars, labels are used as control structures to regulate the applicability of the production rules. For example, in a rule such as OPP→ OP/S, the dispatcher S is a property added when the rule is applied and inherited in all subsequent rule applications unless specifically removed. When the dispatcher appears on the left hand side of a rule, it serves as a condition that must be met for the rule to be applicable.

Two tests have been applied to the grammar to prove its validity: analysis and synthesis. *Analysis* has been applied to an existing painting to determine whether compositional phenomena used by the painter can plausibly be furnished by the grammar. For the analysis, the authors started with Diebenkorn's Ocean Park number 111 (Fig. 3). In Fig. 4, the grammatical derivation of linear composition for the painting is shown.

*Synthesis* has been used to generate compositions to determine whether the grammar specifies particular compositional phenomena that cannot be plausibly attributed to an extention of the painter's style. A synthesis test was applied by generating a linear composition randomly from the grammar. A pseudo-Diebenkorn derived from the following sequence of rule applications is shown in Fig. 5: 2, 6, 17, 17, 11, 31, 31, 30, 38, 37, 30, 31, 30, 30, 30, 32. The generated structure has both a busy and an open region, as in the style of the painter.

The idea behind this approach is very interesting, providing a means to understand the underlying structure of large classes of paintings and suggesting a computational theory of style. In fact, construction of a grammar would be very

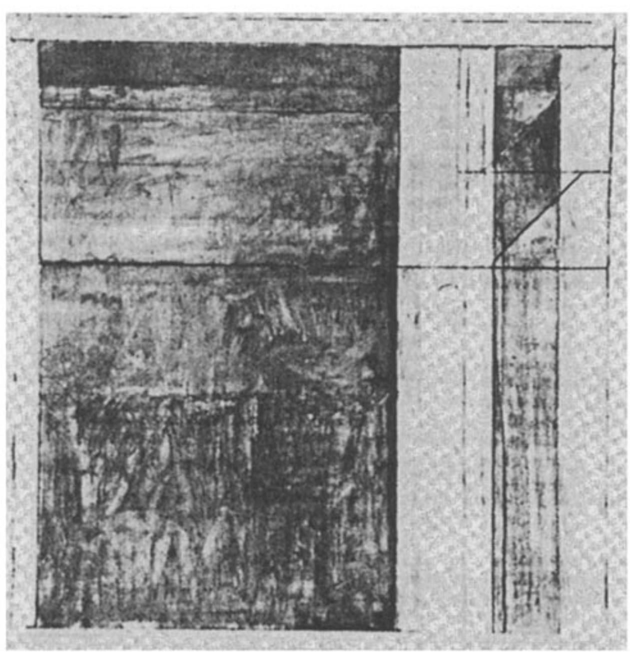

**Fig. 3.** Richard Diebenkorn's Ocean Park number 111. (from Kirsh 1986. Copyright © by Pion Ltd. Reprinted by permission of Pion Ltd).

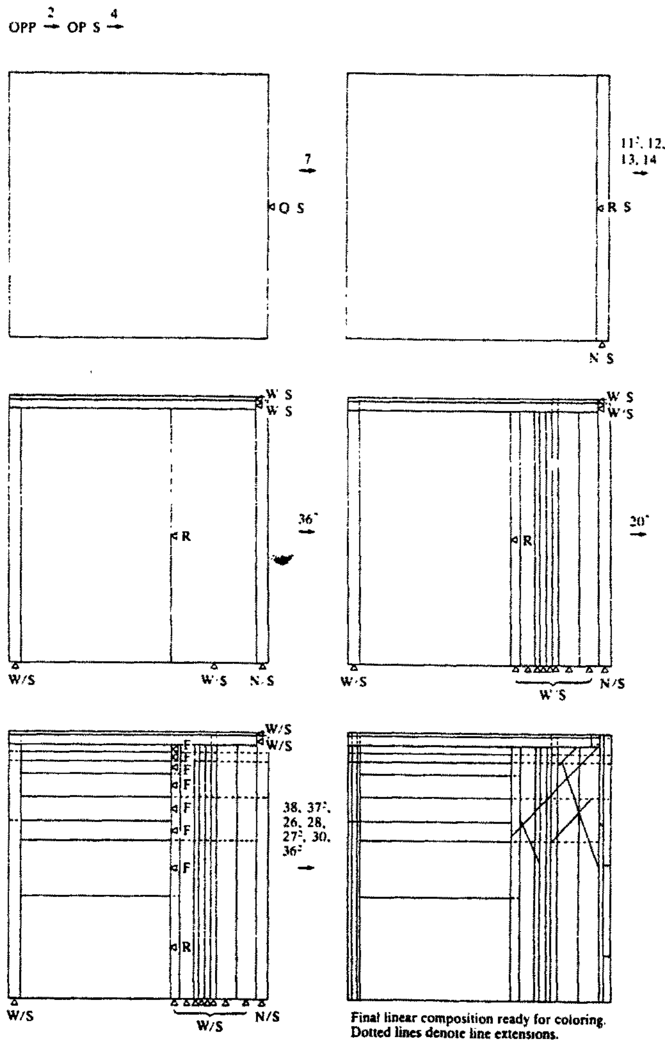


**Fig. 4.** Grammatical derivation for Ocean Park number 111. (from Kirsh 1986. Copyright © by Pion Ltd. Reprinted by permission of Pion Ltd).

difficult for drawings of any complexity, and to define a grammar for every artist or class of similar paintings would be an almost infinite task. It is not even clear whether pieces of one grammar could be used in another. As a further critique, no investigation of the author's aesthetic viewpoint is provided. More work could be done in this direction to include semantic properties in the grammar.

### 3.2 An Investigation of Algorithmic Aesthetics

Gips and Stiny have defined a formal structure for theories of aesthetics. In particular, they are interested in using algorithms to model many different

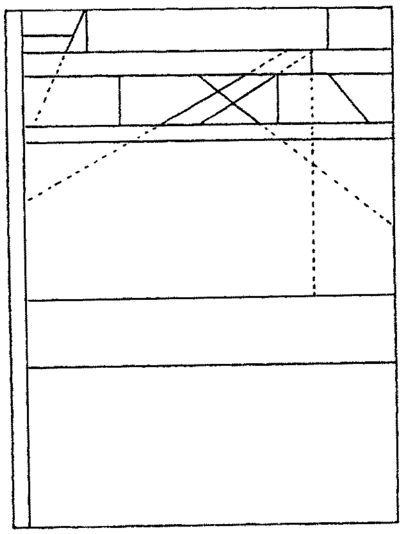

**Fig. 5.** A grammar-generated Diebenkorn-like painting. (from Kirsh 1986. Copyright © by Pion Ltd. Reprinted by permission of Pion Ltd).

aesthetic viewpoints to interpret and evaluate works of art. An *aesthetic viewpoint* determines how an object is understood as a work of art and how the quality of an object is judged when it is understood in this way.

A viewpoint may be thought of as a construct of an observer used to consider objects as works of art and as a construct of an artist to produce new works. There are many possible aesthetic viewpoints. The variety of viewpoints is apparent, for instance, when two different people understand and appreciate the same object as a work of art in two different ways.

To provide a logical framework in which aesthetic viewpoints can be represented, the concept of an aesthetic system is defined. An *aesthetic system* is represented by a 4-tuple $<I_A,R,E,O>$. $I_A$ is a set of interpretations defined by algorithm A; R is a reference decision algorithm which determines if an element of $I_A$ refers to a given object; E is an evaluation function defined on $I_A$; and O is an order in the range of E. In an aesthetic system $<I_A,R,E,O>$, the two initial components are called an *interpretative system*, the final two components are an *evaluative system*.

An interpretation is a pair $<\alpha,\beta>$ where $\alpha$ is an encoding of an input, and $\beta$ is an encoding of an output from the algorithm A when present with input $\alpha$. The interpretive system decides what type of interpretation can be made for an object and which interpretations refer to which object. The evaluative system computes the aesthetic value of an interpretation and ranks the interpretations in order to determine their appropriateness from given aesthetic viewpoints.

The description of an aesthetic system for pictures in Fig. 6, generated using a shape grammar as described above, is given in (Gips, 1975). In the example, the aesthetic system deals with the internal coherence of pictures having generative specifications. For interpretations $<\alpha,\beta>$ in the system, $\alpha$ is a specification of the underlying structure of the picture and $\beta$ is the description of the picture in terms

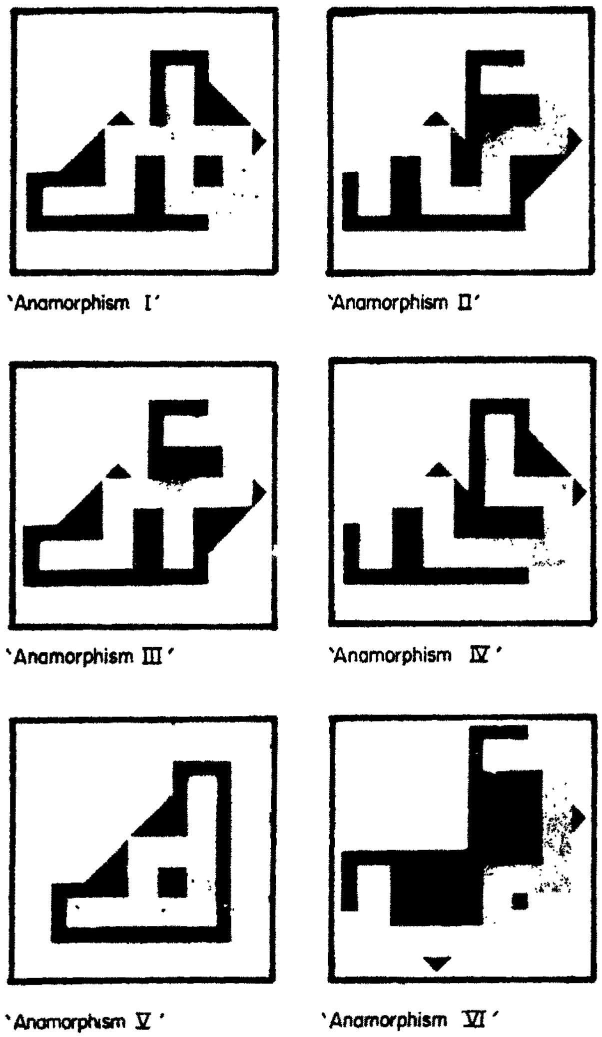


**Fig. 6.** Anamorphism I-VI. (from Gips 1975. Copyright © by Pergamon Press PLC. Reprinted by permission of Pergamon Press PLC).

of its shape and color. The algorithm A embodies the conventions by which $\beta$ can be constructed using $\alpha$. The set $I_A$ is infinite, containing all possible generative specifications and their associated shapes and colors. The evaluation function used in the system assigned high aesthetic value to interpretations having short generative specifications and long shape and color. The pictures were ranked in the following order: I, IV, II, III, VI, V.

The formal definition of which components are necessary to embody a formal structure for an aesthetic system by Gips and Stiny appears rather complete and extensible to other domains where pictures provide an underlying computational structure. Also in this case, the field of application is limited though, since it does not seem possible to embody in the system less structured cognitive rules.

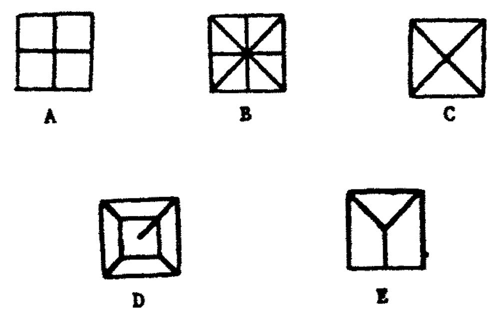

Fig. 7. Basic set of standard patterns. (from Mazlack 1981. Copyright © by IEEE. Reprinted with permission by IEEE).

### 3.3 Representing Aesthetic Judgment

The main interest of Mazlack and Granger in their work (Mazlack, 1981) was in studying the development of a basis for a minimal aesthetic productive capability. In order for a machine to provide this feature, three criteria were identified:

1. Its result must not be precisely predictable.
2. It should operate within broad stylistic conventions.
3. It should reflect a change in the result due to aesthetic judgment from the external environment.

Following these criteria, their investigation focused on the development of a representation providing aesthetic control to an image-producing mechanism from the small set of the simple standard patterns shown in Fig. 7. The construction of the patterns was to be controlled by a representation specifying an aesthetic judgment as to the relative suitability of adjoining local pattern combination horizontally, vertically, and diagonally, and by random selection.

Weights representing aesthetic judgments on all possible basic pattern compatibilities and the effects that would be produced in the image were assigned. The initial weight assignment was modified on the basis of observers' aesthetic reaction, finally resulting in pictures such as the one in Fig. 8.

The method worked on its limited domain and seems to be extensible by enlarging the initial set of patterns. Although the images do not pretend to be works of art, they reflect a beginning of an understanding of how aesthetic judgment could be represented.

### 3.4 AARON, the Knowledge-Based Artist

#### *3.4.1 AARON's History*

A better understanding of Cohen's work can be obtained by following the history of the evolution of AARON. The initial question that pushed the author towards

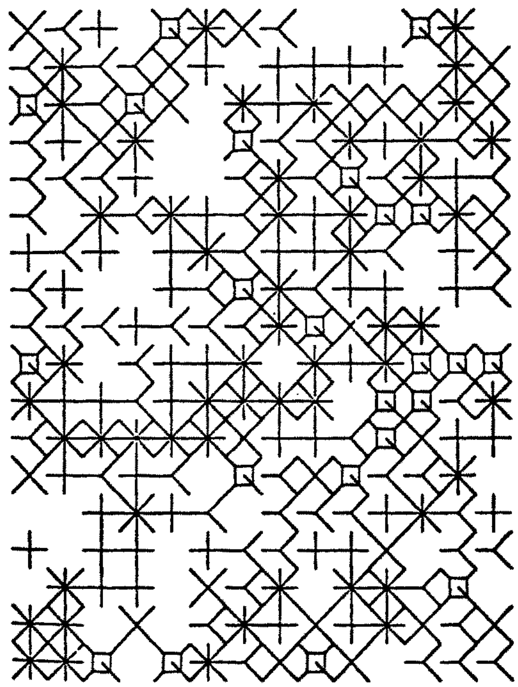

**Fig. 8.** A picture generated though aesthetic control driven pattern selection. (from Mazlack 1981. Copyright © by IEEE. Reprinted with permission by IEEE).

the creation of AARON at the end of the 1960s was how it is that we are able to make sense of systems of marks which were generated within cultures utterly remote from our own, the cultural meanings of which we could not possibly know.

He speculated that a distinction needed to be made between meanings carried by mark systems and the sense of meaningfulness generated by those systems. Meanings, in the sense of transmitted messages, would necessarily be less and less present as their origins become more remote from our cultural location. The sense of meaningfulness, on the other hand, must be generated through non-cultural commonalities between mark-maker and mark-reader. He thus concluded that the non-cultural commonalities reside in the human cognitive system, which was assumed to have been essentially constant throughout human history.

The first AARON represented the attempt to identify and stimulate the actions of a small set of cognitive primitives such as closure, insideness, repetition and division. In its first years of existence, AARON was taught things about the human cognitive system and about drawing, but nothing about the objects of the world that were evidently evoked for viewers by its drawings (Fig. 9). The program succeeded in demonstrating the power of the cognitive system itself, devoid of world knowledge, and the degree to which visual representational systems take form and power from the cognitive system.

Had the body of primitives been enlarged, AARON's drawings would have become richer or more complex. This however was disconfirmed by the results

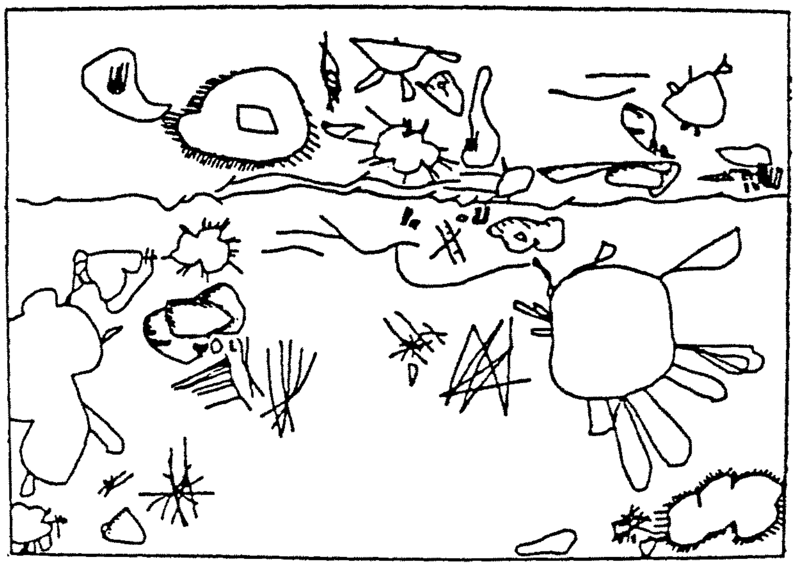

**Fig. 9.** An example of AARON's early drawings. (from Cohen 1987. Photos by Becky Cohen. Reprinted with permission by the author).

obtained in the following years. Cohen started feeling that the program needed some major restructuring in order to further its results, such as the ability to store information across invocations, and the introduction of some knowledge of the world.

Two main changes improved AARON's output. First, the introduction of a new procedure to the repertoire of the drawing primitives stimulating the drawing skills in young children: as all children begin drawing by scribbling and then by surrounding the scribbles with a closed line, they proclaim the drawing to represent something, so AARON started to make its first attempts at representing something. The more knowledge AARON was given about visual space, the more the entities agglomerated from its primitives took on explicit thing-like characteristics. The need to provide explicit knowledge about those things was becoming clear.

Introducing object-specific knowledge was actually the next major step in AARON's history, late in 1983. Objects such as human figures, trees, and rocks were described to AARON in terms of their structures, not of their appearances. For instance, an agglomeration of closed forms made in the following way – a big one in the middle, a smaller one with markings above it, approximately two appendages more or less hanging from the bottom – was called a figure; a skinny branching structure was called a tree; a big lump more or less rectangular was a rock. Other rules were specified, e.g., rocks may be piled on each other, figures may stand on rocks but not vice versa, and so on. The first result of the introduction of knowledge of objects of the world is shown in Fig. 10.

By any standards, AARON's knowledge of the world was rudimentary. The latest developments of AARON have been directed to increasing its knowledge about the world. In particular, knowledge bases have been specified to describe

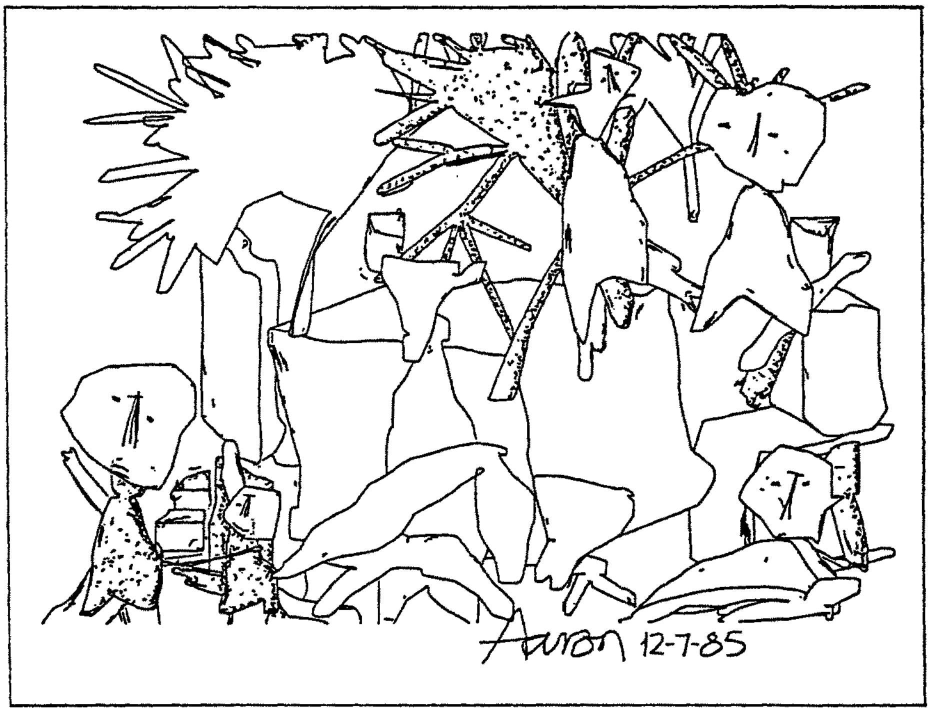


**Fig. 10.** An example of AARON's drawings during the first phase of world object knowledge. (from Cohen 1987. Photos by Becky Cohen. Reprinted with permission by the author).

plants and human figures. Besides the obvious information about how a plant or a human figure is constructed, AARON has been taught rules governing plants' morphology and growth, and rules determining how a human figure can preserve its balance. A late AARON drawing is shown in Fig. 11.

In practice, AARON makes drawings of whatever it knows without requiring any further instructions for the making of a particular drawing, and indeed without possessing any mechanism through which it could take instructions. There is a large consensus that AARON's drawings show a high level of artistic accomplishment. But AARON does not embody any of the aesthetic rules commonly believed to guide the production of artworks, it simply follows simple principles like the "find enough space" rule which governs the drawings' composition. As the author claims "the aesthetics of AARON's performance can be regarded as an emergent property arising from the interaction of so many interdependent processes, the result of so many decisions in the design of the program, that it becomes meaningless to ask how much any one of them is responsible for the outcome . . . If someone else wrote a similar program, I would expect it to exhibit a different identity and a different aesthetic".

The results obtained by Cohen with AARON are exciting, providing a deep insight in understanding the nature of visual representations, i.e., what kind of cognitive activities are involved in the making and reading of representational objects and what artists need to know about the world and about the nature and

**Fig. 11.** An example of AARON's current drawings. (from Cohen 1987. Photos by Becky Cohen. Reprinted with permission by the author).

strategies of representation itself. Further research directions based on the current status of the system could be both augmenting the knowledge about the objects of the world and investigating other cognitive processes such as learning, to provide AARON with the capability of judging its work and modifying itself accordingly.

### *3.4.2 Structure and Knowledge Representation*

In its first version, AARON was mainly a production system, but it also had knowledge encoded in procedural form. The controlling driver of the program consisted of a set of productions having conditions on the left hand side and mainly procedures on the right hand side. The rules corresponded to cognitive primitives involved in human visual perception and representation, and can be put into three classes:

1. distinction between figures and ground.
2. differentiation between closed and open forms.
3. differentiation between inside and outside.

The program was constructed in a hierarchical fashion. The topmost level was called *artwork* and was responsible for upper level decisions such as the overall use of space in the current drawing, i.e., control of the density of information in the picture. The procedure *mapping* is a lower level procedure and is in charge of the allocation of space within the drawing for each individual element of the drawing, i.e., if and where an element is placed in the drawing. Finally, *planning* determines what type of element to draw next.

The kind of knowledge of the first version of AARON was thus mainly *procedural knowledge of representational strategies*. The latest versions of AARON embed *object specific knowledge*. Object specific knowledge ranges over four levels of increasingly procedural and context-dependent knowledge:

1. *declarative*: declarative knowledge specifies the hierarchical structure of a figure: e.g., an arm has an upper arm, a forearm and a hand; a hand has four fingers and a thumb, and so on.
2. *functional/structural*: functional knowledge takes its form in ranges given for a figure and how a subfigure is related to its superfigure. For instance, the range of movement of a human arm from the shoulder may describe an arc which must begin somewhere behind the back and below the waist, and must end somewhere behind the back and above the head.
3. *exemplary*: knowledge becomes exemplary when a value for some range specified in the functional knowledge is chosen.
4. *procedural*: procedural knowledge is the executable code for the drawing of an element of the figure to be drawn.

## 4. Conclusions

In this paper, a critical overview of the most relevant approaches to the study of art perception and creation using AI techniques has been given. The only system showing "creative" capabilities in AARON, while the others are mainly concerned with reading and understanding an artwork. Or, in other words, while AARON considers both syntactic and semantic properties of a piece of art, the other systems are just concerned with capturing its syntactic nature.

Can you make a computer understand and produce art? Whatever answer we might agree on, it is certainly clear that the results achieved so far have provided some good understanding and modelling for some art-making behavior. But for computers to act as humans, models of such things as perception, memory, learning, and mental categories are required, and from the state of the art it seems that a long route has to be covered before getting the essence of human artistic behaviors. In addition, both to understand and to produce aesthetic objects one has to take into account the cultural environment in which the artwork is produced and how the artist is influenced by it, the intended meaning of the piece of art, and the definition of new forms of aesthetics.

The issues raised by the research on AI and art are broad and deep, and involved aspects currently studied in several disciplines such as computer science, cognitive science, semiotics, and art history. Apart from the achievements from a strictly aesthetic point of view, the results obtained by computer art may be helpful in the integration of computer technology and humanness. The need for new paradigms involving art and science is evident.

## Acknowledgements

I would like to thank Prof. Harold Cohen for his support and encouragement in my decision to enter the field of computer art.

## References

Anderson, J. (1984). Cognitive Psychology. *Artificial Intelligence*. **23** (1) 1–11.
Apter, M. (1977). Can Computer Be Programmed to Appreciate Art? *Leonardo*. 10, 17–21.
Badler, N. (1989). Artificial Intelligence, Natural Language, and Simulation for Human Animation. In Magnenat-Thalmann, N. and Thalmann, D. (eds) *State of the Art in Computer Animation* Springer-Verlag, Tokyo.
Barr, A. (1983). Artificial Intelligence: Cognition as Computation. In Machlup, F. and Mansfield, U. (eds) *The Study Of Information: Interdisciplinary Messages*. Wiley Interscience.
Chomsky, N. (1965). *Aspects of the Theory of Syntax*. MIT Press, Cambridge.
Cohen, H. (1979). What Is An Image? *Proc. IJCAI*. **6**. 1028–1057.
Cohen, H. (1985). Can Computers Make Art? Tools, Rules, and Other Issues of Art-making Significance. *Nicograph '85 International Symposium*.
Cohen, H. (1987). Implementing an Expert Artmaking System. *Proc. 1st International Symposium on Artificial Intelligence and Expert Systems*. Berlin, West Germany.
Cohen, H. (1988). How to Draw Three People in a Botanical Garden. *AAAI '88*. Saint Paul, Minnesota.
Davison, N. (1982). Designing a Recursive Framework for Evolutive Poetry. *Proc. of the 2nd Symposium of Small Computers in the Arts*. Philadelphia, PA.
Dietrich, F. (1986). Visual Intelligence: The First Decade of Computer Art (1965–1975). *Leonardo*. **19** (2) 159–169.
Franke, H. (1965). *Computer Graphics – Computer Art*. Springer-Verlag.
Gips, J. and Stiny, G. (1975). An investigation of Algorithmic Aesthetic. *Leonardo*. **8**. 213–220.
Hofstadter, D. (1980). *Goedel, Escher, Bach: An External Golden Braid*. Vintage Books.
Hofstadter, D. (1985). *Metamagical Themas: Questing for the Essence of Mind and Pattern*. Bantam Books.
Jankel, A. and Morton, R. (1984). *Creative Computer Graphics*. Cambridge University Press.
Kahn, K. (1979). Making Aesthetic Choices. *Proc. IJCAI 1979*. 448–450.
Kirsh, J. and Kirsh, R. (1985). Computer Viewing Artists at Work. *Proc. of the 5th Symposium of Small Computers in the Arts*. Philadelphia, PA. 65–67.
Kirsh, J. and Kirsh, R. (1986). The Structure of Paintings: Formal Grmmar and Design. *Environment and Planning B: Planning and Design*. **13**. 163–176.
Mallory, R. (1969). Computer Sculpture: Six Levels of Cybernetics. *Artforum*. 29–35.
Mazlack, L. and Granger, M. (1981). Representing Aesthetic Judgments. *Proc. of the International Conference on Cybernetics and Society*. Atlanta, GA.
Noll, M. (1966). Human or Machine: A Subjective Comparison of Piet Mondrian's "Composition with Lines" (1917) and a Computer-Generated Picture. *The Psychological Record*. **16**. 1–10.
Palyka, D. (1985). Computer/Art: Depolarization and Unification. *IEEE Computer Graphics and Applications*. 47–56.
Roads, C. (1985). Research in Music and Artificial Intelligence. *Computing Surveys*. **17** (2) 163–190.
Stadler, C. (1982). Computers and Choreography. *Proc. of the 2nd Symposium on Small Computers in the Arts*. Philadelphia, PA.
Stiny, G. (1980). Introduction to Shape and Shape Grammars. *Environment and Planning B*. **7**. 343–351.
Turing, A. (1950). Computer Machinery and Intelligence. *Mind*. **59** (236) 433–460.
Wilson, N. (1984). Computers and Poetry. *Proc. of the 4th Symposium on Small Computers in the Arts*. Philadelphia, PA.
Zeltzer, D. (1986). Knowledge-Based Animation. In Badler, N. I. and Tsotsos, J. K. (eds) *Motion: Representation and Perception*. Elsevier Science. 318–323.

*Correspondence and offprint requests to:* Roberto Maiocchi, Dipartimento di Elettronica, Politecnico di Milano, Piazza L. da Vinci 32, 20133 Milano, Italy.