GENERAL NOTE

Discovery through Interaction: A Cognitive Approach to Computer Media in the Visual Arts

Peter Beyls

ABSTRACT

The author investigates creativity and intelligence in the process of computer-assisted image making. Exploring the meaning of creativity, he proposes the use of standard methods of symbolic computation to model the creative process. This conceptual approach is compared to other inclinations toward digital media that an artist might exhibit. Proceeding from a discussion of how the manipulation of symbols and processes often reveals new underlying concepts, he provides a few examples of machine-based art that have their roots in knowledge-based decision making.

We are living in a terrible world, and so artists have terrible responsibilities. The world today is one of contradictions, and perhaps as a consequence, I believe that many things can be described and explained using a comparative style of reasoning. Contemporary society is on the edge of producing truly intelligent machines, yet it is unable to support its members in their most basic needs. Computers allow for ever richer personal experiences, yet thousands of children die day after day. We must conclude that we cannot possibly speak of collective intelligence. Since creativity and intelligence are mutually inclusive, this leaves the artist in a helplessly frustrated situation. Nowadays, artists have access to the most sophisticated production media; it seems only natural to expect them to arrive at concepts of equal sophistication. I hope for more powerful statements and fewer pretty pictures.

Tools, concepts and techniques borrowed from the field of artificial intelligence (AI) are relevant in various domains of artistic expression, including image making, from at least two different points of view.

First, from the standpoint of cognition, AI offers concepts that assist our analysis and understanding of machine-based aesthetic decision making and problem solving in such domains as search space definition and exploration, the representation of beliefs with what constitutes exciting image making and the exploration of ideas—without really knowing where one is heading—with the intention of discovering new and interesting concepts or images.

Second, from the standpoint of technology, state-of-the-art AI workstations and powerful languages like LISP [1] offer a flexible environment in which to develop ideas by allowing the gradual specification of objectives in an incremental style of programming.

In general, an exploratory attitude seems to be at the heart of creative behavior. This implies that in the generation of visual complexities we are more concerned with discovery through the interactive evaluation of ideas than with the creation of end products. The goal is to invent hypothetical worlds that exhibit some sort of activity. This activity is interpreted in a process where responsibilities are shared by human and machine, and statements are produced in a common effort. Within these worlds and their representation, knowledge is more important than data, and evolution is of greater concern than structure.

Peter Beyls (artist, educator, researcher), Artificial Intelligence Laboratory, Vrije Universiteit Brussel, Building K, Pleinlaan 2, B-1050 Brussels, Belgium.

Received 7 June 1988. Solicited by Herbert W. Franke.

CREATIVITY

As humans, we are systems interested in ourselves and in other systems that exhibit interesting behavior. Our intention is to set up a framework to describe and analyze the attitudes required to develop behavior-inspired ideas. The complexity of life itself seems to offer a wealth of these ideas. The main goal is the creation of interactions between systems, that is, exchanges of information between the user/programmer and computer programs. These programs are formalizations of what we think constitutes interesting behavior in a particular context. Our current context is computer-assisted creative behavior. We may try to characterize the explicit experimental attitude of creativity in many ways. The following list is designed to offer some preliminary insight into the significance of computers as vehicles for artistic expression:

1. Speculation about the nature of things is more important than their documentation.
2. Creativity does not emerge from isolation. On the contrary, the creative process is explicitly receptive toward the environment. In other words, creativity is responsive rather than introverted.
3. The problem of creativity involves experimentation in new domains, not the conservation or the creation of extensions of known problem domains.
4. Typically, creativity requires the exploration of large amounts of data. There is not really a wish for description.
5. Instead of trying to strengthen the power of ideas that proved useful in the past, self-revision seems to be the basis of creativity.
6. Intensity as a parameter characterizing a perceived message is much more important than intrinsic quality or virtuosity of expression.
7. Discovery of the unpredictable aids creativity more than the confirmation of the predictable. Flexibility is preferable to precision.
8. Creative people prefer taking chances with many things rather than expressing faith in the predetermined.
9. Sometimes, asking the right questions in the process of

Pergamon Press plc. Printed in Great Britain.
0024-094X/91 $3.00+0.00

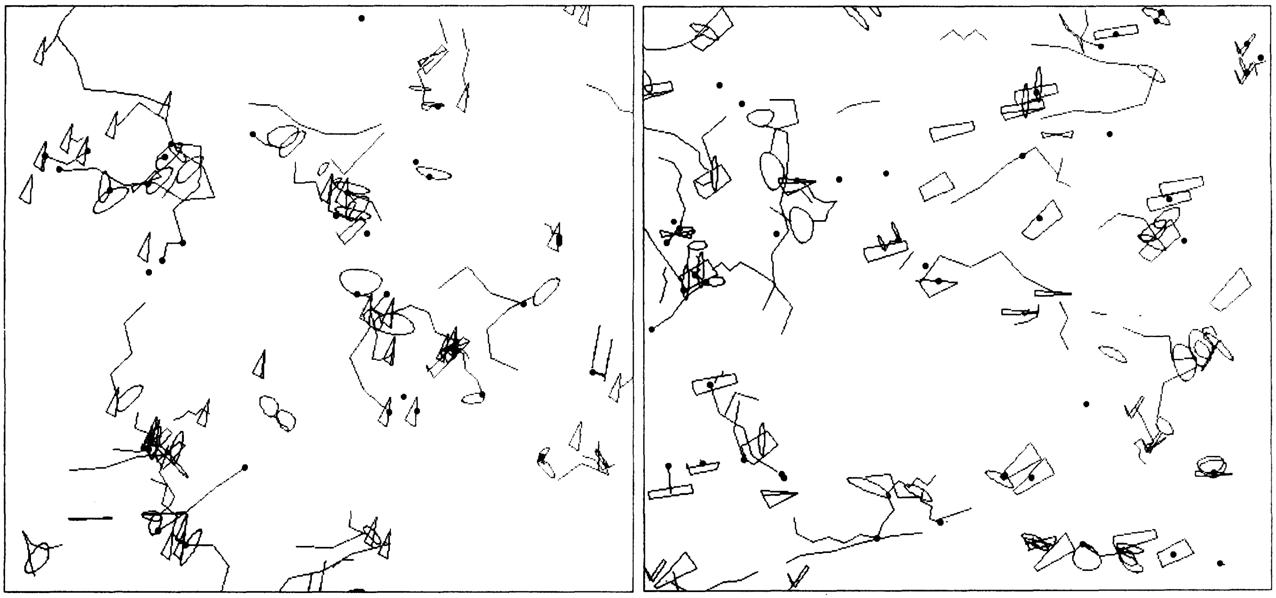

Fig. 1. EVA_1, inkjet prints, 18.5 × 17 cm, 1983. These works were realized on a HP9000.

interpretation seems more relevant than giving answers in the process of explanation. Discovering a problem may be far more exciting than solving a problem.

10. Creativity is more like the synthesis of many things into new concepts than the transformation of existing materials through craft.
11. Focus in a creative search is not fixed but mobile, since the goal in the search tree shifts according to new knowledge acquired as a consequence of searching.
12. Without discipline there cannot be freedom. Likewise, without chaos there cannot be creativity.
13. A creative system adapts to wide swings in context, and stability in the system can only be dynamic.
14. Representations of what we *know* are more interesting than the mere recording of what we see. Our objectives involve the meanings of things rather than appearances.

From Activity to Interactivity

Various writers have tried to classify products in the field of computer art, their criteria reflecting pictorial evaluations. However, Laske [2] distinguishes attitudes rather than appearances. I believe it is useful to sketch briefly these families of attitudes, because they indicate the way from data manipulation to knowledge-based decision making.

The *object-oriented* approach aims to realize ideas as objects. The artist produces images that refer only to themselves. More often than not, the artist applies digital technology as a mechanical extension in the process of making images. Commercial paint systems are a case in point here. Objectives are speed of production and 'visual thinking', i.e. fast evaluation of many pictorial alternatives. Strikingly, many artists aim to produce extravagant and abundant visual complexities without worrying about underlying principles. Attention to the external, physical attributes of an artistic statement seems to be the only concern.

In contrast, the *tool-oriented* approach examines the specificity of the new medium of computer art. This happens on at least three levels: (1) New types of imagery are suggested by the technology itself; consider the exaggeration of aliasing in some computer-assisted painting, the use of exotic paintbrushes, etc. (2) Computers offer new means to think about structure; consider the feasibility of simultaneous evaluation of macro- and micro-structures. (3) Computers offer new tools to interact explicitly, a phenomenon that gave rise to a completely new subfield in the performance of computer music—interactive composing.

Tool-oriented artists may use mathematics or computational logic as means to generate abstract images. However, this computational power remains secondary—the artist aims toward amplification of visual complexity. How this is achieved is not seen as a matter for consideration. Image processing in all its forms, as well as commercial computer animation, seem to fit this family of attitudes. Also, the semantic meaning of these images seems trivial; images of transformed realities dominate the landscape. From the point of view of aesthetics and of history, we may notice a new sort of surrealism (and hyperrealism) that obviously lacks the critical social engagement of the original surrealism. Note that popular imagery, such as fractal images and the creation of superrealistic images, seems to be situated between the object-oriented family and the tool-oriented family.

The *process-oriented* approach is much more concerned with problems of behavior on both sides of the computer screen, and with how these two worlds interact. Tool-oriented people would think of an image and then search for a formal method to generate it. Not so with process people, who think of the generative process as a concept that may then be activated. These artists formulate this activity in a program and then sit back to find out what the concept actually implies: learning and introspection through visual feedback. Behavior is central here. Processes run in real time and are interactive and conversational; by their very nature they are experimental and explorative. (Elsewhere I have discussed some of these themes at length [3,4].)

Finally, the *cognitive approach* of AI offers excellent tools to describe any ambiguous process, including the creative process. Some directions of current work in this area are next described here.

From Exploration to Discovery

Research in knowledge-based art aims to shed light on the roots of creative processes and on how these processes work. The central task is to try to understand why people are creative, for example, by building models to verify and simulate aspects of human creativity. It is important to note that behavior-simulating programs are built by a very special type of experimental behavior. First, computer programs are used as idea-amplifiers in that they allow for exploration of ideas by seeing (or hearing) the implications of these ideas. Consider that composers debug programs by listening to them. Frequently, the coded idea (i.e. the program) reveals a hidden potential of the original idea; sometimes it shows that the idea does not lead anywhere at all. Exploration gives birth to side-effects. Very often, these side-effects are more interesting than the initial ideas that triggered them. Exploration gives them a substantial chance to develop. This method has more to it than just trial and error. This is a form of introspection.

As artists, we are interested in spontaneous processes. We know roughly what we are looking for, but the details of these processes remain elusive. When formalizing this approach in a program, our minds are open to more than just this single objective. We know that to look only for what we want means we stay locked within our present knowledge and ideas.

An exploratory attitude is at the heart of any experimental work. It allows us to discover connections between actions, connections that we normally consider nonexistent. Exploration generates search; it opens more precise venues where the exploration becomes more goal oriented. Indeed, artistic activity in general can be viewed as a giant search tree [5]. In the long run, exploration creates strategies for further exploration.

Second, exploration can be embedded in computer programs. The AI community has expressed considerable interest in designing systems that are capable of exploring large amounts of data without being told explicitly what to look for. Systems like AM and Eurisko [6] are classic examples. AM and Eurisko explore data bases, searching for interesting items using heuristics or 'rules of thumb'. Exploration here involves learning since the objective is to discover new (and possibly interesting) concepts while exploring a search space that is complex, immense and irregular. As mentioned before, we may view the artist's output, as well as his or her whole career, as determined by a process of continuous exploration. At all decision nodes, orientation is guided by a new (and possibly just discovered) goal.

EXAMPLES: EVA_1, EVA_2 AND EVA_3

After experimenting with algorithmic systems for over a decade, I conceived the basis for a knowledge-based production system during the early 1980s. Since then, the EVA programs have been under continuous development [7].

The EVA programs grew out of my interest in thinking of images not as computed structures but rather as possible products of a process both defined and activated by a body of knowledge. This knowledge is represented in the system as a list of rules and a collection of things expressed simply as facts. Rules are procedures informing what to do given certain circumstances. The program is goal oriented; it is designed to solve a particular problem using representations of strategies. Much of this work has parallels with expert system design, i.e. the acquisition of domain-specific knowledge from human experts, its representation in forms readable by both humans and machines and, finally, its subsequent exploitation to solve similar problems in the same problem domain.

EVA_1 had rules scattered all over the place; it was poorly structured.

Fig. 2. EVA_2, plotter drawings, 73.5 × 110 cm, 1984. These works were realized on a Symbolics 3600 LISP machine.

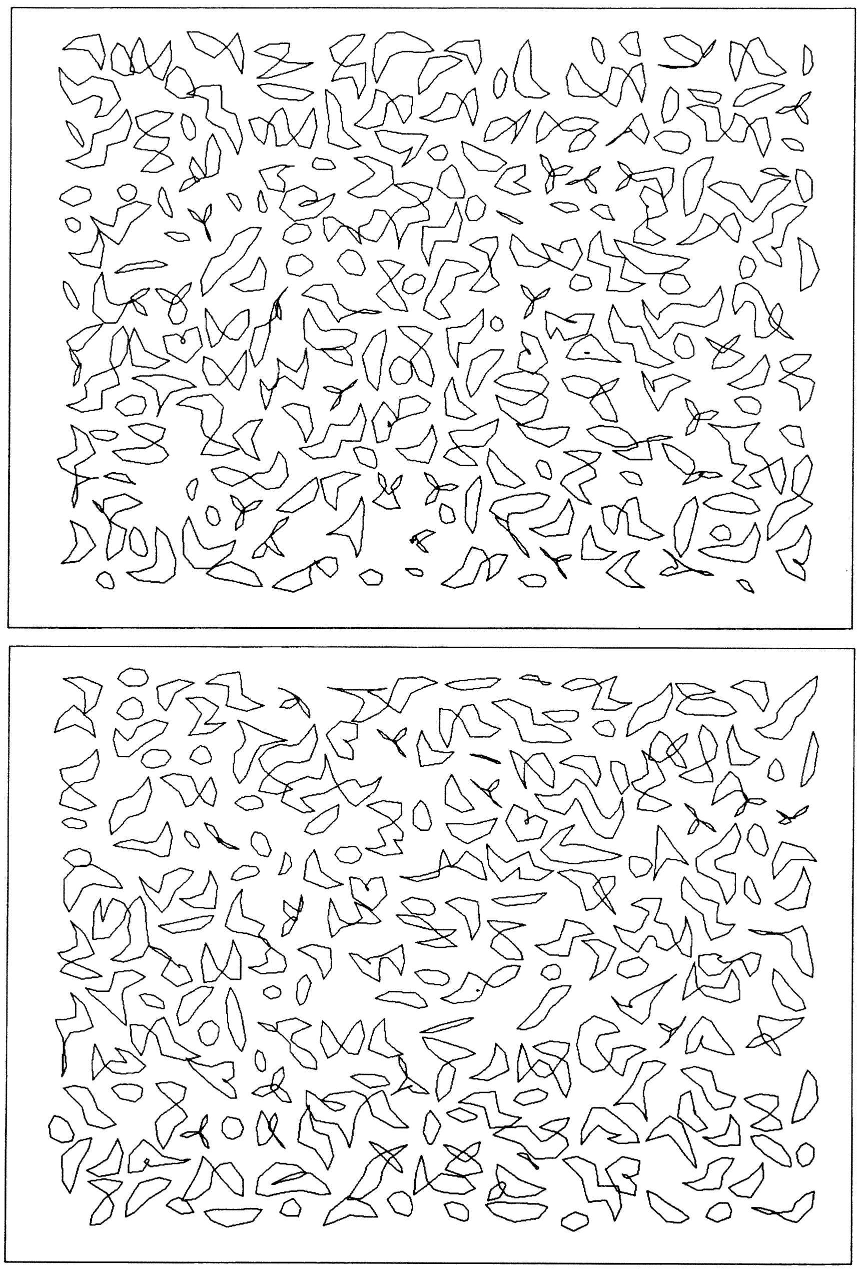

However, there were plenty of new concepts beyond conventional programming style. This program used deterministic formulations as well as stochastic decision making, yet it behaved neither randomly nor as a fixed sequence of actions. The point is that EVA_1 was sensitive to both spatial context and the history of its own activity (Fig. 1).

Spatial sensitivity is guided by rules. The system knows about size, position, density, whether some element is next to another or on top of it, etc. Temporal history, on the other hand, is mainly a product of frequency-based procedures. The system uses the TRESH algorithm [8]. It is easy to understand that the program starts with excessive freedom: Because there is nothing in particular to refer to, anything will do. The drawing emerges in gradual fashion; more objects to refer to emerge, and they increasingly constrain the program's 'legal activities', so that more and more fine-grained control becomes necessary.

EVA_2 uses a style of object-oriented programming in a LISP environment. The program deals with active objects that communicate by sending messages to each other. The idea was to have a large number of such objects all having similar characteristics but situated at different points in space (stage 1). If they are close enough to each other, objects may associate with a limited number of other objects, this association being represented by a line (stage 2). This gives rise to complex constructs. As a final step, EVA_2 produces contours according to the nature and complexity of the constructs (stage 3) (Fig. 2).

EVA_3 follows similar lines but includes an 'architect'. The architect imposes a macro-structure while other rules fire and generate more detailed structures. This program also simulates bursts of spontaneous energy while drawing lines. Clearly, all of the above programs are a step toward knowledge-based programming. Foremost, the results they suggest are evidence of how one might think about machine-initiated aesthetic decision making in general (Fig. 3).

SPECULATIONS

In 1969 Robert Mallary suggested six levels of cybernetics in an attempt to evaluate degrees of creativity [9]. He spots the interweaving of activity by artist and machine while the two are engaged in a common effort to produce computer-assisted sculpture.

At level 1 the machine acts as a supercalculator, not at all as an electronic embodiment of creative intelligence. Level 2 results as a swing from a mechanistic vision toward a versatile and self-propelled system. At level 3 the machine is programmed to generate its own complete artistic statements. However, the artist still monopolizes the crucial decisions, including those that determine when a work has been completed and is acceptable.

At level 4 the machine's heuristic capabilities are defined. This level reflects a shift in the definition of the role of the machine, from a slave, to a collaborator, to a virtual surrogate of the artist. The implementation of the program becomes analogous to the activities of a living organism: the program itself manufactures contingencies and instabilities and then proceeds to solve them. Also, the system remembers how spontaneous problems were solved successfully in the past; expertise accumulates.

At levels 5 and 6 the artist is no longer needed. Here, we spot what is probably the first artistic statement suggesting the plausibility of 'autonomous agents' in the arts, that is, computational entities that exhibit intelligence in their behavior through self-control, that learn from interacting with their environments.

These levels exclude human interaction from the process of creation; the artist, like a child, can only get in the way. At level 6 the artist will probably not even be able to pull the plug, since the machine now has self-replicating modes of existence and has evolved into a state of pure energy or pure concept. The suggested evolution is, of course, paradoxical. It should not be interpreted as a process of gradual exclusion of human intervention. The contrary is true. The correct understanding of the paradox should point toward the conception of artificial systems where humans and machines cooperate to produce strong artifacts in a common effort.

CONCLUSIONS

Interactive Works of Art That Know about Their Owners

It is easy to foresee interactive works of art that 'know' about their owners. These works might be responsive and alert to their complete environment—that is, all living and nonliving systems with which they interact. These works will have learning capacities and become smarter as they age; we will talk of the age of the systems. It is amazing that so few artists understand the potential of machine-based art, even though the technology continually becomes cheaper, more powerful and more accessible. It is now possible to set up complex experiments using standard microprocessor technology equipped with home-made real-world sensors and activators.

Many of these ideas are being evaluated in the AI subfield of 'autonomous agents', i.e. completely independent entities capable of 'growing up', perhaps just as a child would. At present, it is not clear how—and out of what materials—such works of art would be made. Some think that art will end up as pure concept; others try to explore these ideas pragmatically. I took the latter approach; the computer-controlled environment [10] is an early example of a conversational, interactive, learning work of art. A small computer equipped with visual sensors produces electronic sounds on a multichannel system. The program aims 'to talk in an abstract language' and 'to understand' what the computer sees. If nothing happens from the outside, the program will try to think of something by itself, of some activity related to earlier experiences, or it will change its structure to adapt to a novel context.

The creation of virtual musicians is another exciting challenge of machine-based art. This type of program simulates the intellectual activities of a living musician, including planning, improvisation and musical interpretation. Our laboratory does research in this field, but details are beyond the scope of this paper.

Expanded Dimensions for the Individual

First, much of the potential of interactive technology lies in its introspective character; it offers a practical means to manipulate ideas and concepts in a tangible way. The result of this 'playing with ideas'—some prefer to call it visual thinking—is that the artist gets feedback about his or her own intellectual approach and about the effectiveness of current computational strategies. Very often this feedback results in an amplification of ideas: for every idea implemented, ten are set aside for later consideration.

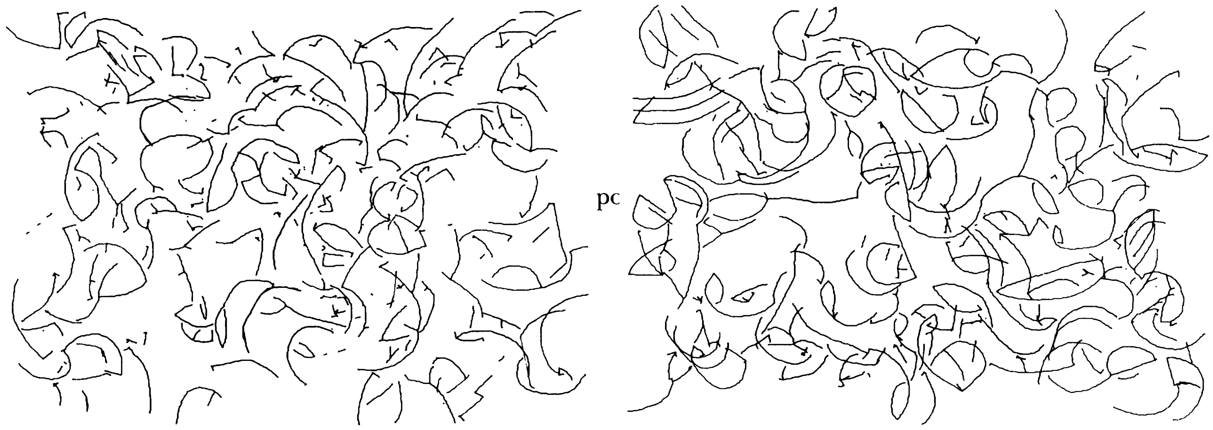

Fig. 3. EVA_3, computer laser prints, 19.3 × 26.7 cm, 1985. These works were realized on a Symbolics 3600 LISP machine.

Second, current technology permits the creation of a personal assistant, shaped explicitly to assist the particular needs of a single individual. These programs are designed to function as intelligent partners in aesthetic decision making and have access to knowledge about the person who designed the system. Design of intelligent workstations for composition is already at the heart of much musical research today.

Deeper Integration of Concept and Medium

Consider the pioneering work on the interaction of art and AI of both Harold Cohen [11] and Edward Ihnatowicz [12]. More than a decade ago, Cohen designed the first expert system specifically to create images from artistic objectives. Some 20 years ago, Ihnatowicz designed and built the first intelligent sculptures, computerized mobiles that were receptive to their environments and able to express 'personal character' while interacting with their surroundings. The exceptional degree of innovation and outstanding quality of this work is in dramatic contrast with today's situation.

Ironically, from the point of view of aesthetics much computer art today is extremely traditional, although many artists like to describe themselves as innovators and their work as beyond stylistic constraints dictated by tradition. The community of artists using advanced technology is strongly polarized. Many artists are fascinated more or less exclusively by technical particularities of the medium. The artistic meaning and significance of the potential to create vehicles for exploration, are not fully understood.

While global awareness of potentialities, which requires continuous reorientation and adaptation to fluctuating conceptual considerations, is at the heart of true (artistic) intelligence, only a handful of computer artists seem to grasp this fundamental principle and act accordingly, witness their own creative output. This is a sorry state of affairs. The key problem here is that most artists do not realize that a computer is a general purpose machine. It becomes a medium by virtue of being programmable. Natural language techniques from AI may assist those who are artistically sophisticated but too technologically naive to express themselves. However, I do expect a fundamental breakthrough, a deeper integration of concept and medium, when more artists learn how to program and to express themselves in a symbolic programming environment.

Acknowledgments

Many colleagues and friends were instrumental in shaping the conceptual and practical framework that made this research possible. I thank Ludo Cuypers for guiding me through the intricacies of LISP, Luc Steels for invaluable support, Maria Dimou for expert proofreading and, most of all, Edward Ihnatowicz and Harold Cohen for providing instructive and critical feedback.

References and Notes

1. LISP (LISt Processing) is a general purpose programming language that represents expressions as lists. It is the standard language for research in artificial intelligence.

2. Otto Laske, "Definitions of Computer Art", *Computer Music Journal* (Winter 1984).

3. Peter Beyls, "Interactive Media", *Da Vinci Magazine* **1**, No. 1 (Brussels) (1986).

4. Peter Beyls, "Intimate Machine Interaction", in H. W. Franke, ed. *The Visual Computer* (Berlin: Springer Verlag, 1986).

5. Luc Steels, "Zoekprocessen in kunstcreatie" in *Geest, computer, kunst* (Utrecht: Stichting Grafiet, 1986).

6. Douglas Lenat, "Why AM and Eurisko Seem to Work", *AI Journal* **23** (1984).

7. For more details about the development of EVA programs, see Beyls [3].

8. The TRESH algorithm (from ThRESHold) is a procedure for complex stochastic decision making, organized as a network of interconnected software modules. The modules act if they are activated above their threshold value. Interaction between the modules may result in complex sequential performance while critical tuning is achieved from specification of the threshold values. See Peter Beyls, *TRESH*, exh. cat. (Antwerp: ANHYP, 1985).

9. Robert Mallary, "Six Levels of Cybernetics", *Artforum* (Fall 1969).

10. Peter Beyls, *Computer-Controlled Environments* (Gent: Museum of Modern Art, 1985).

11. Harold Cohen, "What Is an Image", *Proceedings of the International Joint Conference on Artificial Intelligence* (Tokyo, 1979).

12. Edward Ihnatowicz, *Cybernetic Art, A Personal Statement* (London: by the author, 1986).