Creativity and computation.
Tracing attitudes and motives.

Peter Beyls
Electronic Imaging Department
St Lukas School of the Arts
Paleizenstraat, 70
B - 1210 Brussels
peter@arti.vub.ac.be

Abstract

The present paper suggests a framework for critical evaluation of computational approaches to creativity and aesthetics from a trans- disciplinary point of view. We suggest various models for building systems that exhibit aspects of creative behaviour. The alternatives include traditional knowledge-based systems, constructs inspired by linguistics and the concept of self-organization as perceived in products of natural creativity. The evolutionary paradigm is outlined and its application in evolving cellular automata is briefly documented.

1. Introduction.

This paper aims for the creation of a framework to address the following questions: can we build machines that merely simulate human creative activity or is there potential to emulate true creative thinking in a computer program? How do we build self-reflective programs given certain aesthetic criteria and what are these criteria supposed to look like? Social expression, introspection and the synthesis of meaning while handling multiple views of the same thing at the same time are at the heart of human creative behaviour. All this seems hard to expect from a machine though quite powerful statements have been produced in recent years, in particular by using methods of artificial intelligence and algorithms inspired by biological evolution.

We are forced to study the psychology of creative decision making if we ever want to implement aspects of it in a computer program. These aspects have many faces including the creation of contexts to augment the chances for something interesting to happen, the invention of problems and questions (not answers), the persistence on exploration and flexibility rather than final products and precision, the expression of interest in the meaning of things rather than (or in spite of) their possibly extraordinary visual appeal, and many more.

Because much computer art focuses on the generation of intricate structures does not imply that critique should be limited to formalist criticism i.e. the study of formal relationships or excellence in designed visual organization. Criticism should not be blinded by the complexity of the medium but receptive to the expression of intense feelings, the communication of ideas of truly great intensity and question the relevance of artistic statements i.e. do they say anything on the human condition as it really affects us? Or should critique be focused on the consequences of the ideas and feelings expressed, for instance, by serving some social end beyond the form of the work itself? Obviously, impressive pictures do not provide a useful ground to guarantee artistic integrity. In any case, a work must offer the potential to raise questions and, possibly fundamental questions which trigger a creative response in the observer or listener. Perhaps, because of its interactive nature, the computer is the ultimate channel for introducing augmented responsiveness in the appreciation of artistic statements.

People are creative because they interact as social species, they express an opinion toward society as well as to the self. This global awareness and explicit physicallity of the human body is important to note both as a structural constraint for machine creativity as well as a source of inspiration. Indeed, we have witnessed the emergence of various socially inspired, distributed, computational methods like the actor paradigm, robot ecology, or even massive parallel hardware. On the other hand, physical interaction with an unpredictable environment; the study of artificial life through actually physically building synthetic entities, is a topic of intensive research. As a matter of fact, ever richer cross-fertilisation grows between art and science. Scientific discoveries suggest new imagery as well as new ways of thinking about the laws and processes forming reality as we experience it.

2. Overview.

We shall outline possible computational strategies for the implementation of an artefact generating machine but first we shall trace the attitudes toward the computer as a medium for artistic experimentation and expression. These attitudes range from seeing the computer as a tool for basic, deterministic visualization, to computer-assisted, interactive evaluation of ideas, to the delegation of responsibility and autonomy to a synthetic art system based on a hypothesis of how human creativity works. In addition, it is observed that attitudes range from passive imitation to interactive introspection, from a collaborative, conversational approach to a wish for independent creative autonomy and, finally, from the creation of physical end products to the promotion of ideas and focusing on the conceptual, the immaterial aspects of the medium.

A creative algorithm should exhibit maximum opening towards anything that could possibly happen; we must start by having all degrees of freedom potentially available, it must produce something new given the constraints of an existing aesthetic paradigm and its output must be useful in the sense that it forces interpretation and the synthesis of meaning. It must explore the unpredictable, thus it must be non-deterministic. Evolution theory according to Darwin provides the ultimate example of non-deterministic activity: the genetic materials of life are combined and modified arbitrarily and from a mass of results non-viable products are eliminated by the constraints imposed by the environment. The emergence of life itself has taken nature millions of years though very fast parallel processing computers may detour the very inefficiency of the evolutionary approach. Incidentally, genetic algorithms, mimicking the mutation and cross-over operations in DNA have been used explicitly (Sims, 91a) as well as implicitly, as metaphor (Verostko, 90) to grow synthetic creatures with considerable success. Note that the building blocks under evolutionary consideration need not be limited to formal elements like perceptible structures but may extent to the the procedural modules themselves (programs as genotypes) while the programmer's opinion about the current state of affairs constitutes the aesthetic constraints for deciding what gets a chance to survive.

Another class of algorithms imposes constraints at the generative level, and is much more efficient than the generate-and-filter method. Generative grammars, as inspired on the work of Chomsky, are good examples of computationally powerful methods. Once the building blocks and rules designed, the system will run itself and breed a rich variety of instances from which the program (or user) may select at random.

We shall interpret the above methods in the light of the two prevalent lines of research in artificial intelligence today: first, within the framework of knowledge-based systems and, second, in the sense of behaviour oriented experimentation. Expert systems using (aesthetic) knowledge representation in rules and facts and relying on inference as their basic mechanism are exemplary for the first category. In sharp contrast, the principle of self-organization and spontaneous pattern- formation from local operations in permanent interaction with the environment are characteristic of the second.

3. Motivation.

It has been argued repeatedly that the digital medium liberates the artists from the limitations of conventional media (for instance, lack of precision and physical inertia) and that the computer offers "seemingly infinite capabilities". It is stated that the computer provides potential for fundamental novel ways to create new images and transform existing ones. These claims definitely sketch a wrong perspective. When viewing computers as mechanistic extensions of the hand, it makes them an ideal environment to both preserve existing modes of expression and guarantees a fluid continuation of conventional aesthetics. In sharp contrast, the machine should be seen as a liberation of the mind, not the medium. To this end, it is essential to regard computers as perfectly general purpose symbol manipulators, to view them as open-ended systems, as tools in which programs live. To write programs is to invent conceptual machinery that reflects a personal attitude towards the act of creation itself.

Computers are vehicles for exploring ones own belief systems of what it means to be creative. One is forced to program them in order to discover their language and gain freedom of expression. The question of representation of what constitutes creative behaviour, then, becomes an issue of central concern. What is represented will set the scope of the system, how things are represented will determine its flexibility and usefulness. From this it follows that software systems are by no means neutral, e.g. a lot of work with paint systems is very similar and the hand of the machine (programmer), so to speak, is clearly visible. On the other hand, since AI systems are built on descriptions of what is thought to be known in a specific domain, they are, by definition, coloured by the subjectivity of their designer.

Most algorithmic approaches to creativity feature formalist representations borrowed from mathematics, the study of formal languages, physics or computer science. From cognitive psychology we know that emotions and affect play important roles in human motivation, intention and attitude -- all fundamental ingredients of the creative process. Yet, very few attempts have been made to build representations of these higher order factors. Given the difficulty to represent the semantic significance of artistic statements, in practice, syntax is favoured over meaning. However, some research has been carried out -- in the field of computer music -- in an effort to tighten the gap between the organization of structure and its appreciation and interpretation (Katayose & Inokuchi, 1989). This work is inspired by recent advances in AI which tries to model this higher order cognitive activity of the human mind; the theory put forward in (Minsky, 1986) has been very influential.

The insistence on flexibility, rather than on precision, is valued of greater aesthetic relevance by creative individuals. However, CAD systems were designed to automate the design procedure and do not constitute a platform for intimate man-machine cooperation in the sense where solutions emerge from collective evaluation of intermediate results. The same mechanistic view is expressed in the sequencer paradigm of computer music where timing may be excellent but is flexibility very poor.

Too bad, with digital accuracy, analog sensuality is thrown out on the same account. All the elements of appreciation that talk to the body rather than the mind, are very underdeveloped in computer art. Barnett Newman would never have painted large surfaces if it was not his intention to emancipate colour and relate the experience to the topology of the human body. Metaphorically speaking, digital quantization should remain characteristic to the computer and not penetrate our critical abilities.

4. Attitudes.

Four classes of motivations towards machines in the creative process received detailed outline in (Beyls, 1987) and are briefly summarized in the context of the current paper. Class one refers to the mechanistic point of view criticized in a previous paragraph. It implements established modes of thinking and existing ways of doing, it aims the production of unique works of art, it is object-centered. The second class acknowledges the unique features of computers to combine or transform pictorial or musical building blocks from explicit specification by the user. Examples include spectral interpolation in digital sound synthesis and frame interpolation in computer animation. A third class accommodates everything procedural. The design of algorithms and the evaluation of their behavioural potential is the creative act. Here, a more experimental attitude is expressed explicitly. Many procedural programs do not aim for single-ended results but for the creation of families or classes of works: all output is different but exhibits the same qualities. The development of ideas by a process of gradual definition of objectives is unique to the digital medium. Most composer's assistant programs are situated at the intersection of class two and three. Class four is characterized by cognitive modelling of creative activity and is the theme at the heart of the current paper.

The true instrumental power of the digital medium lies in the fact that it offers a play-ground for breeding ideas, a functional environment for experimentation, a field for growing and testing interesting procedural activity. In short, a computer is a conceptual microscope which reveals detailed explanations of how we navigate for surprise (or disappointment) in hypothetical worlds of our own design. The role of inventor and explorer coincides. From this it follows that the programmer and the user must be one and the same person. This form of intimate man-machine interaction (Beyls, 1986) can be potentially dangerous; the artist can get trapped in a perpetual design cycle since all output is, by definition, intermediate output. In such an interactive exploratory attitude of creativity, the objectives of both the programmer and the program may shift according to what suddenly appears as an interesting discovery. One discovers things by testing and doing, not by abstract contemplation. When the (aesthetic) goal or focus is redefined, the exploration will propagate itself and search for that goal. The intimacy of the process signals how computers are highly integrated in creative thinking and is evidence for the deep integration of concept and medium.

But the question remains; what do we start from, where do we start and where are we going. One thing is certain, any computer comes as a preformulated package, in terms of both hardware (e.g. processor speed, memory capacity) and software (e.g. operating system, man-machine interface, available languages). This uniformity will push the artist in a certain direction since one cannot start from scratch, or as (Nadin, 89) puts it "the machine as predetermined makes it a poor substitute for the empty canvas of the painter". It is indeed very difficult to build an empty computer and it should be realized that this has enormous aesthetic consequences.

5. Definition.

Creativity in humans seems to function as a delicate interplay of freedom and discipline. Freedom points to having uncertainty as a fundamental generative principle. From Prigogine we have learned that science is about the exploration of the uncertain and the relationship of man and the universe. Joseph Beuys has formulated the same idea as follows "The artist points towards the totality of the relationship between the physical incarnation of humanity and its total spirituality." (Nairne, 1987). John Cage has commented in a panel discussion "The function of art is to let us experience the general uncertainty without suffering but with enjoyment." (Cage, 1990). Uncertainty is not just identified as a universal underlying principle from the social, scientific and artistic points of view, but recognized as a necessary condition for creative evolution to unfold. So freedom refers to creating a context for the unpredictable to happen. On the other hand, total unpredictability divorced from a critical context does not carry much meaning. Therefore, the limiting forces of discipline have to be put to action. So the combination of uncertainty and constraints constitutes a productive ground for breeding creative ideas. Moreover, for an idea to work creatively it should be new and useful in the sense that its functions as an object of contemplation and provocation.

In computational terms, discipline may be formulated as constraints which act as critics that decide which random ideas get a chance to survive by examining their coherence with a body of existing principles. These principles have been referred to as conceptual spaces (Boden, 1990) or search spaces (Feigenbaum & Feldman, 63). Conceptual spaces stand for a frame of reference given a specific domain, in the arts it signifies an existing aesthetic style or idiom, in the sciences it is referred to as a paradigm. That space both specifies what is acceptable and provides information on the generative mechanisms of the given idiom. However, truly creative ideas seldom fit into a generally accepted, established conceptual space but signal a new style or paradigm. Incidentally, the insistence on experiment and the refusal to rely on historical pressure and to compose from self- designed first principles is at the heart of much 20th century music. So very creative people don't just explore a given problem but create additional problems, they redefine or extend a given conceptual space. This makes the difference between combining existing ideas into new ones using some rule-based method and the redefinition of a paradigm by changing its limits, for instance, by lowering the influence of constraints that define it or forcing additional constraints into consideration. Note that we speak in computational terms: rules are used to infer new and interesting conceptual configurations, while constraints delimit the field of all potentially possible configurations.

Let's consider the following examples where dropping constraints led to something fundamentally new. Mondrian's oeuvre is characterized by a gradual process of giving up representational considerations by using a restricted formal vocabulary in a wish to focus on intellectual order exclusively. Arnold Schoenberg decided to set free harmonic space by dropping the need of a tonal center in music. The quest for the spiritual, dropping representations of reality, the discovery of abstraction pioneered by Kandisky, and later, its evolution to post-painterly abstraction and an occupation with the coneptual most visible in Reinhardt's black paintings. Fontana transformed the canvas into something three dimensional. Marcel Duchamp was perhaps the most radical of all disregarding aesthetics as such by suggesting the ready-made.

The above formulation of creativity is only losely related to associative theories explaining the creative act. The latter define creative thinking as the formation of associative elements into new combinations hoping for a useful result. Of course, the combination of more mutually remote elements increases chances for novel results though they also receive increased critique given the constraints of the given stylistic paradigm. The theory of bissociation of matrixes of thought put forward in (Koestler, 1964) suggests that the creative act involves the linking together of two previously unconnected 'frames of reference' or associative contexts.

Randomness may function as a combinatorial device as in the cut up method devised by William Burroughs. It involves the snipping of newspaper texts and rearringing them at random until a suitable effect is obtained. To cook up fresh ideas, the approach was also explored by many composers ranging from John Cage to David Bowie.

New paradigms or styles may also emerge from discontent with existing ones. To break out of a feeling of frustration with current global aesthetics artists/composers would apply the negation operator i.e. they would react by doing the opposite. Note that the frequency of this process of switching back and forth is increasing toward the end of the 20th century. Consider the sudden birth of minimal music in the late sixties, composers rejected complexity resulting from mental gymnastics and suggested very simple formulae for motivic development instead. However, later, new-complexity emerged closely followed by new- simplicity.

6. AI approaches.

The goal of artificial intelligence is two-fold. First, to get a better understanding of how human intelligence works. Theoretical models are designed by questioning the functioning of human cognitive processes, including perception, reasoning and problem solving. These models may provide better insight into the nature of human creativity. Second, AI is interested in the implementation of natural cognition in artificial systems to serve various specific engineering goals. AI has also developed very advanced symbolic programming environments of exceptional flexibility e.g. Lisp, that are now very well documented and generally available. Lisp promotes creativity since it promotes interactivity and taking risks.

The idea of exploring conceptual spaces outlined above may be translated to the actual methodology of navigating in a search space. This view of creativity has been described by (Steels, 1986) and (Ames, 1992). In order to find out about the richness of a search space one must describe it in a functional model. The ideosyncratic features of a particular space can be represented in a computer using standard methods of AI. Then, the key idea is exploration. Early work in AI by Newell, Simon and Shaw on the programs Logic Theorist and later General Problem Solver (GPS) are of historical significance. They saw creativity as a special form of problem-solving activity characterized by novelty, unconventionality and the recognition that problem formulation is indeed very difficult. GPS and its predecessor are examples of rule-based systems; the idea is to solve a problem by applying formalized knowledge (represented in production rules), to grow a decision tree and navigate in that tree from given initial conditions toward a solution, known as the goal state. The machine has to keep track of partial solutions and if it gets stuck in an impasse it has to backtrack to a lower level in the tree and restart from there. So the program generates suggestions for getting to the goal-state and evaluates their distance from that goal-state. However, the decision tree may grow out of hand because of a combinatorial explosion, so special short cuts, known as heuristics, are needed.

Heuristics are rules-of-thumb used to detour this problem as it occurs, for instance, in a chess playing program; the number of legal moves is so enormous that it would take too much time to consider them exhaustically. In addition, the pressure on memory, to represent everything possible, would grow out of hand. Therefore, these programs consider only a few moves and counter-moves up to a certain depth of reasoning. The most promising move is selected, the one that brings the current state closer to the goal-state.

The pioneering work of Newell and Simon led to the development of modern expert systems that are capable of handling ill-specified symbolic information about a specific domain. The most impressive example of the knowledge-based approach in the visual arts is the work of Harold Cohen documented extensively in (Mc Corduck, 1991). Cohen's pioneering expert system Aaron creates original simulated free hand drawings from artistic expertise represented as procedural and factual knowledge. The generative behaviour of the program evolved from abstract work to figurative drawings of human figures when Cohen added real-world knowledge about the relationships between parts of the human body. The total knowledge-base represents the artist's global concept space, a style which was developed -- and this is important to note -- before he got involved in computing. The stylistic continuity in his work is not surprising since Cohen was his own knowledge engineer, in sharp contrast to industrial expert systems where the task of knowledge acquisition and knowledge exploration are strictly separated.

7. Exploration.

My own work with expert systems is characterized by viewing them as exploratory devices explicitly. The accent is more on the process of collaboration between man and machine than on the creation of an autonomous art generating system. The artist does not formalize a pre-existing style as or does not manipulate a predefined vocabulary using a collection of previously stipulated rules. On the contrary, one aims at the discovery of a powerful strategy to generate pictures. Focus is no longer on visual structure as such but on methods to think about action and interaction. This requires maximum openness on the part of the programmer and maximum purity (minimum constraints) of the medium, an empty computer would indeed be most welcome.

Many artists borrow ideas from examples in nature. The wealth of imagery in nature makes us question, isolate and formalize the processes that are responsible for this richness. These processes are then used to produce works that relate to their real-world equivalents. Growing trees with a grammar is a good example (Smith, 84). The creation of hypothetical worlds is the other end of the scale: the artist relies on products of his own imagination. This is a definite wish but never exclusive since no

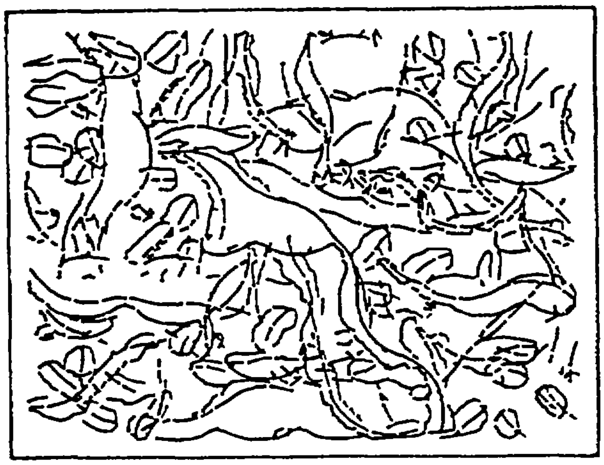

artist escapes the gravity of context. However, starting from simple representations of some hypothetical behaviour in an invented, artificial world and through gradual refinement of the algorithms involved, interesting phenomena (that map to pictures or sound) may emerge. One is involved in a circular process of optimization: evaluation of the continuous feedback from the program shows the road to changing the program in question. Designing such programs is planning for change. Note the relationship of this exploratory attitude with the evolutionary paradigm introduced in later on. A concise, illustrated definition of this experimental method featuring intimate machine interaction is given in (Beyls, 86). Refer to the three accompanying illustrations -- reproduced from a family of large plotter drawings -- for the 1990 state of affairs of EWA, a personal expert system developed by the author.

So the computer offers an instrumental channel for gradual specification of objectives while wandering though a given personal search space. Exploration is guided by what is discovered along the search path and is thus the key to surprise. There has been quite some research in AI to model the exploratory process explicitly in a computer program and in different domains of expertise. The pioneering work of Doug Lenat who designed an automatic mathematician (AM), a program that explored the space of mathematical rules in search of new theorems, was of historical importance. AM considered the beauty (interestingness) of mathematical expressions (using heuristics) it cooked up and consequently focused creatively on the most promising ones through self-modification of the program (Lenat, 77).

(Kahn, 79) has addressed the problem of making aesthetic choices in the context of computer animation. His program ANI is capable of making animated sequences in response to high-level incomplete story descriptions. The user describes personalities, their character and relationships. Suggestions for action are like rules that are combined and modified i.e. explored. The program makes choices for action based on minimization of arbitrariness and maximization of overall coherence.

The previous work is related to Talespin, a system that creates objects judged primarily on their aesthetic appeal (Meehan,76). The program writes stories from the description of animal characters and plans are produced to satisfy their desires i.e. to reach some specific goal. The stories document the problem-solving behaviour of their characters. Many different stories may be generated from the same rules since the objectives of characters are interacting and possibly in conflict. This program offers many possible solutions for the same problem so it may considered creative.

Conflict, confrontation and tension have been recognized as residing in the very conversational process of the artist with a synthetic personality of his own design. (Laske, 90) views composition as conversing with the machine as an alter-ego. The creative process is propelled by the tension that exists between ideas suggested by the program and their interpretation -- the meaning and function they receive in the mind of the composer.

My work in interactive composing has concentrated on the creation of virtual musicians designed explicitly to interact with a single human performer. The musical dialogue is propelled by a conflict the machine- performer aims for the expression of its musical character working simultaneously towards social integration with the human performer. The type of real-time expert systems used here are known as pattern-directed inference systems: they aim at the realization of their embedded goal while, at the same time, keep perceptive channels open to accommodate external stimuli and adjust their plan accordingly (Beyls, 88).

8. Behavioural approaches to creativity.

Observation of creativity in nature shows many processes where coherent structures or complex, coordinated behaviour emerge from the social interaction of many participating organisms as well as their interaction with the environment. This behaviour is referred to as the emergent properties or emergent functionality of a system that only survives if it keeps in continuous communication with its surroundings, just like creativity itself which cannot emerge from isolation. The behavioural approach finds its roots in many different disciplines such as non-linear physics (Babloyants, 86), the study of complex dynamical systems in general and mathematics (Wolfram, 84). In contrast to knowledge-based systems where logical inference over symbolic representations is a basic technique, behavioural systems use local operations on the microscopic level which give rise to complex macroscopic structures given the right environmental conditions. Such systems use distributed representations such as regular arrays as seen in both cellular automata and connectionist networks. The idea is to map the abstract behaviour of a given system to the specific problem at hand. Thus, the creative accent is shifted from designing a global rule-base to inventing interesting local rules -- self-organizing behaviour is suggested as an alternative to knowledge. Since KBS are designed by hand, there is a limit on the maximum complexity a human programmer can handle. In addition, expert systems do not show graceful degradation in performance when situations occur that were not anticipated by the programmer. Also, reasoning or searching over symbolic structures is often very slow which makes expert systems inappropriate for many time- critical applications e.g. interactive composing. All these problems have contributed to the adoption of the dynamics paradigm. For a complete comparative an illustrated discussion of subsymbolic methods and the role of randomness and chance in the context of musical creativity refer to (Beyls, 91).

9. The linguistic approach.

The use of formal methods derived from linguistic theory – in particular the writings of Noam Chomsky -- have been very popular in music (Roads, 85) and, to some extent, in the visual arts (Gips & Stiny, 75). Linguistics describes creativity as the application of a finite set of rules allowing for the construction of an infinite set of possible utterances in any given language (Chomsky, 66). In composition, a generative grammar is used to grow hierarchical language constructs from a few basic starting blocks. The consistent, recursive application of rewrite-rules that define the grammar forces a data explosion. This explains the ultimate computational economy of grammars. In addition, a grammar keeps intrinsic relationships between overall structure and detail. In other words and in terms of music, the surface structure of a piece (events), the intermediate structure (motivic movement) as well as the deep structure (overall composition) are captured in a single, compact representation. Grammars are also used for analysis, for instance to parse natural language. (Cope, 91) uses symbolic pattern-matching to analyse sequences of existing music and extract "signatures" collected in libraries. An augmented transition network (ATN) grammar for parsing phrase-structured information is then used to reorganize these signatures into new compositions. Grammars may be context dependent meaning that other rewrite-rules will be executed according to the evaluation of the neighbourhood in which tokens happen to appear. This provides several alternative choices and is a source of additional complexity.

(Kirsch & Kirsch, 86) tried to describe the compositional structure of works of art in a formal grammar. This work starts from the analysis of geometric designs and produces a database of syntactic units. The deep structures, i.e. the semantic level is not taken into consideration and this points to a shortcoming of grammatical representations.

In conclusion, grammars are powerful devices for designing intricate, hierarchical structures with absolute control over the relationships among their various layers/components. However, grammars function in a single creative idiom: the rules don't adapt and the search space is not extended in the process of recursion.

Finally, an interesting parallel is drawn between the theory of formal languages and so called L-systems, invented by biologist Aristid Lindenmayer. He suggested mathematical models for biological pattern formation and natural growth such as models describing the development of branching

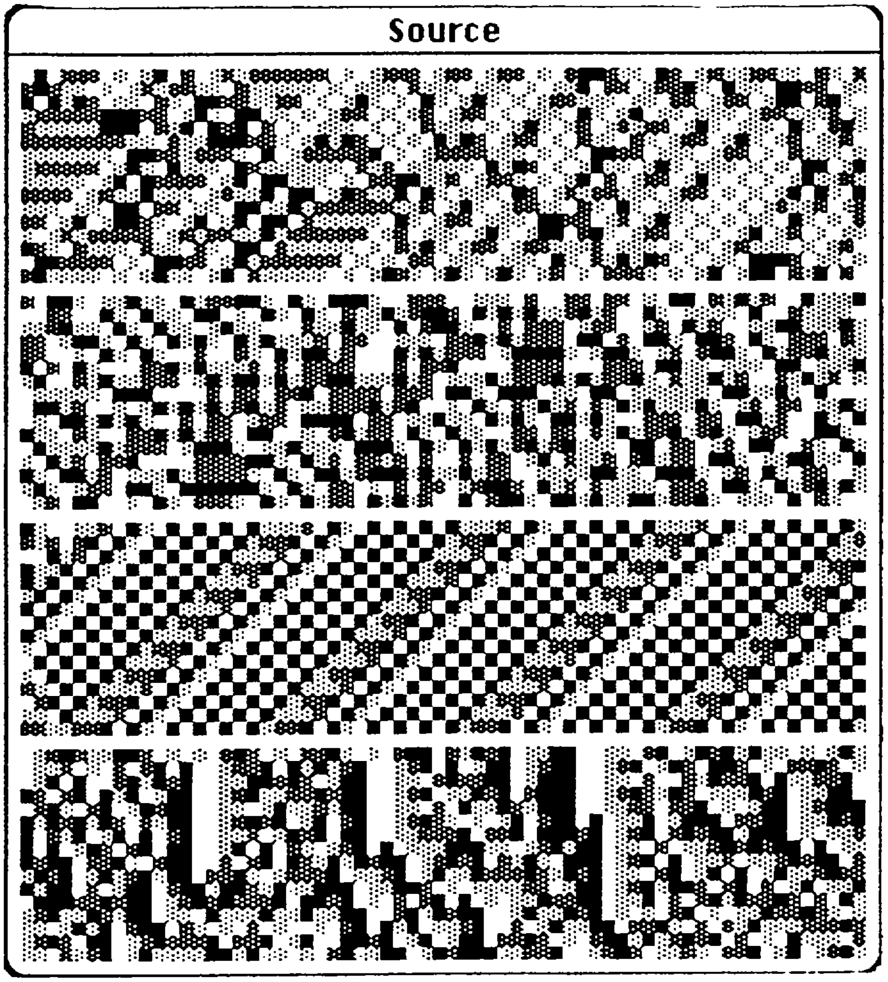

structures in plants. His research currently enjoys a renewed interest in the light of the emerging new discipline of artificial life (Lindenmayer, 68) which brings us to the fourth paradigm for viewing creativity: evolution.

10. Creativity as an evolutionary process.

This approach tries to model creative decision making as an analogue of biological evolution. The diversity in living systems is brought about by the process of evolution; the process that alters the genetic information they embody to make them increasingly better adapted to the environment they inhabit. The set of genes of an organism is called its genotype. Organisms search for maximum adaptation by changing their genotypes keeping only offsprings that are well adapted i.e. exhibit maximum fitness. However, the number of possible genotypes is enormous. In addition, the genes inside genotypes also interact in a non-linear fashion, a phenomenon known as epistasis. Notwithstanding this combinatorial explosion of the search space, evolution has been extremely successful in breeding well adapted rich varieties of living systems. This observation led to the introduction of genetic algorithms (Holland, 75): the idea is to generalize critical processes of evolution and apply them to other search problems. In computational terms, a genotype contains a set of instructions, a micro-program. The execution of the genotype produces a phenotype, that is, a certain behaviour, a pattern or some conceptual machinery. The process by which the genotype evolves, generation by generation, under control of the genotype is called morphogenesis.

A genetic algorithm uses genetic operators that interchange and modify the contents of initial (possibly random) genotypes. Any computer program can be considered as a genotype, its output is viewed as a phenotype. The operators include variation (e. g. change part of a rule), cross-over (e. g. replace left-hand side of a rule by the right-hand side of another rule with cross-over points assigned at random), mutation (replace part of a rule by another one available). The resulting phenotypes -- more exactly, their underlying genotypes -- are evaluated for their fitness (i.e. interestingness or aesthetic appeal) and consequently survive or disappear. What remains will breed the next generation and so on.

The idea is to view this circular process as an instrument to create conceptual bridges between elements of the conceptual space determined by a given aesthetic point of view. This leads to discoveries, one finds patterns that were not anticipated by the programmer, complexities are generated without knowledge of how to generate them explicitly in the first place. More than merely exploring search spaces, genetic algorithms may actually transform them. Very interesting work has been done according to this computational paradigm in the field of computer animation (Sims, 91), computer assisted sculpture (Todd & Latham, 92) and a technique of thematic bridging of musical material using genetic algorithms is described in (Horner & Goldberg, 91).

It should be noted that genetic algorithms make the mechanics of evolution explicit. The evolution of art itself may be approached with genetic operators implicitly in mind. Artists have copied rules from each other with eventual random mutations, produced mixes of mutually borrowed rules as well as rejected common rules. This is how styles emerge, evolve and disappear in the course of art history. Most changes are gradual though: existing features are recombined into new works, the structure of the features themselves remain unchanged. However, when the procedural building blocks are mutated or rejected altogether, non- linear jumps occur that signal fundamental new paradigms: in this light consider, again, Duchamp.

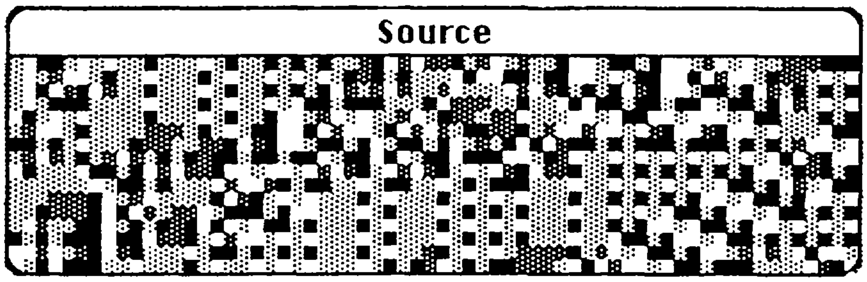

11. Evolution in cellular automata.

As a final example, we shall briefly introduce the interactive evolution of complex dynamical systems formulated as cellular automata in the context of musical composition. A complete description, however, is beyond the scope of the current paper. Cellular automata (Toffoli & Margolus, 87) are very efficient for modelling a wide variety of complex phenomena including musical composition (Beyls, 89). They consist of some regular cellular structure in which the consistent application of simple local rules give rise to surprisingly complex structures. A CA- rule specifies the value of every cell in the next generation from the evaluation of that cells neighbourhood. The starting configuration consists of a random choice from all available values. The idea is to run different rules, evaluate the interestingness of the patterns they produce in time and, apply genetic operators on selected items. For instance, consider the illustration on the previous page showing 4 automata of 64 generations. The neighbourhood is 3, the number of values is 4, the rules (rule 1 to 4, top to bottom) are 80 digits long. The figure above, a single automaton, was produced through a cross- over operation: the first half of rule 1 glued to the second half of rule 3. Results like this are then subjected to morphological analysis, a program extracts and sorts objects according to their complexity. The morphs are then mapped to the musical domain; a polyphonic object results which reflects the behavioural complexity of the automaton.

Aknowledgments.

I wish to express my gratitude to Luc Steels for discussions and inspiration, to Ludo Cuypers for sharing his expertise in Lisp programming and to Karl Sims for encouragement to explore the concept of artificial evolution.

References.

- Ames, C. 1992 Quantifying musical merit. Interface Journal of new music research, Vol. 21, Swetts & Zeitlinger, Holland

- Babloyantz, A. 1986 Molecules, Dynamics and Life: An introduction to the self-organization of matter. Wiley Interscience, New York

- Beyls, 1986 Intimate machine interaction. The visual computer, Vol. 2, nr. 3, Springer International

- Beyls, P. 1987 Discovery through interaction: A cognitive approach to computer media in the visual arts. Leonardo Vol. 24 nr. 3, Pergamon Press

- Beyls, P. 1988 Introducing Oscar. Proceedings of the ICMC, Feedback Editions, Cologne

- Beyls, P. 1989 The musical universe of cellular automata, Proceedings of the ICMC, Ohio State University, Columbus

- Beyls, P. 1991 Chaos and creativity. The dynamic systems approach to musical composition. Leonardo Music Journal, Vol. 1 nr. 1 Pergamon Press

- Boden, M. 1990 The creative mind: Myths and mechanisms. Weidenfeld & Nicolson, London

- Chomsky, N. 1966 Topics in the theory of generative grammars, Mouton, Den Haag

- Cope, D. 1991 Computers and musical style. Vol. 6 The computer music and digital audio series. A-R Editions, Inc. Madison

- Feigenbaum & Feldman, ed. 1963 Computers and thought, Mc Graw Hill, New York

- Gips & Stiny 1975, An investigation of algorithmic aesthetics. Leonardo Vol. 8 Nr. 2,

- Holland, D. 1975 Adaptation in natural and artificial systems, University of Michigan Press, Ann Arbor

- Horner & Goldberg, 1991 Genetic algorithms and computer-assisted music composition. Proceedings of the GA-Conference, UCSD, 1991

- Kahn, K. 1979 Making Aesthetic Choices Proceedings of the International Joint Conference on Artificial Intelligence, Tokyo

- Kirsch & Kirsh 1986 The structure of paintings: Formal grammar and design. Environment and planning B: Planning and design nr. 13

- Katayose & Inokuchi, 1989 Proceedings of the Second International Workshop on AI and Music. IJCAI-89, Detroit

- Laske, O. 1990 The computer as the artists's alter ego. Leonardo Vol. 23 nr. 1

- Lenat, D. 1983 The role of heuristics in learning by discovery: Three case studies. in Machine Learning: An artificial intelligence approach.

- Lindenmayer, A. 1968 Mathematical models for cellular interaction in development, parts I and II. Journal of Theoretical Biology 18, 280-315

- Mc Corduck, P. 1991 Aaron's code. W.H. Freeman, San Francisco

- Minsky, M. 1986 The society of mind, Simon & Schuster, New York

- Meehan, J. 1976 The Metanovel: Writing stories by computer. Yale University Computer Science Research Report 74.

- Nairne, S. 1987 State of the art. Chatto & Windus, London

- Roads, C. 1985 Grammars as representations for music. in Foundations of computer music, MIT Press, Cambridge

- Sims, K. 1991 Interactive evolution of dynamical systems. Proceedings of the first European conference on artificial life. Paris, 1991

- Sims, K. 1991 Artificial evolution for computer graphics. Computer Graphics, Vol. 25, nr. 4. Siggraph 91, Las Vegas

- Smith, A. R. 1984 Plants, fractals and formal languages. Computer Graphics Vol. 18 nr. 3 Siggraph '84 Proceedings

- Steels, L. 1986 Zoekprocessen in kunstcreatie. Geest, computer, kunst. Hagoort & Maessen, eds. Stichting Grafiet

- Toffoli, T. & Margolus, N. 1987 Cellular Automata Machines. MIT Press, Cambridge

- Todd, S. & Latham, W. 1992 Evolutionary art and computers. Academic Press, London

- Verostko, R. 1990 Epigenetic painting: Software as genotype. A new dimension of art. Leonardo Vol. 23, Nr. 1

- Wolfram, S. 1984 Computer software in science and mathematics. Scientific American, September 1984