From Fingers to Digits

An Artificial Aesthetic

Margaret A. Boden and Ernest A. Edmonds

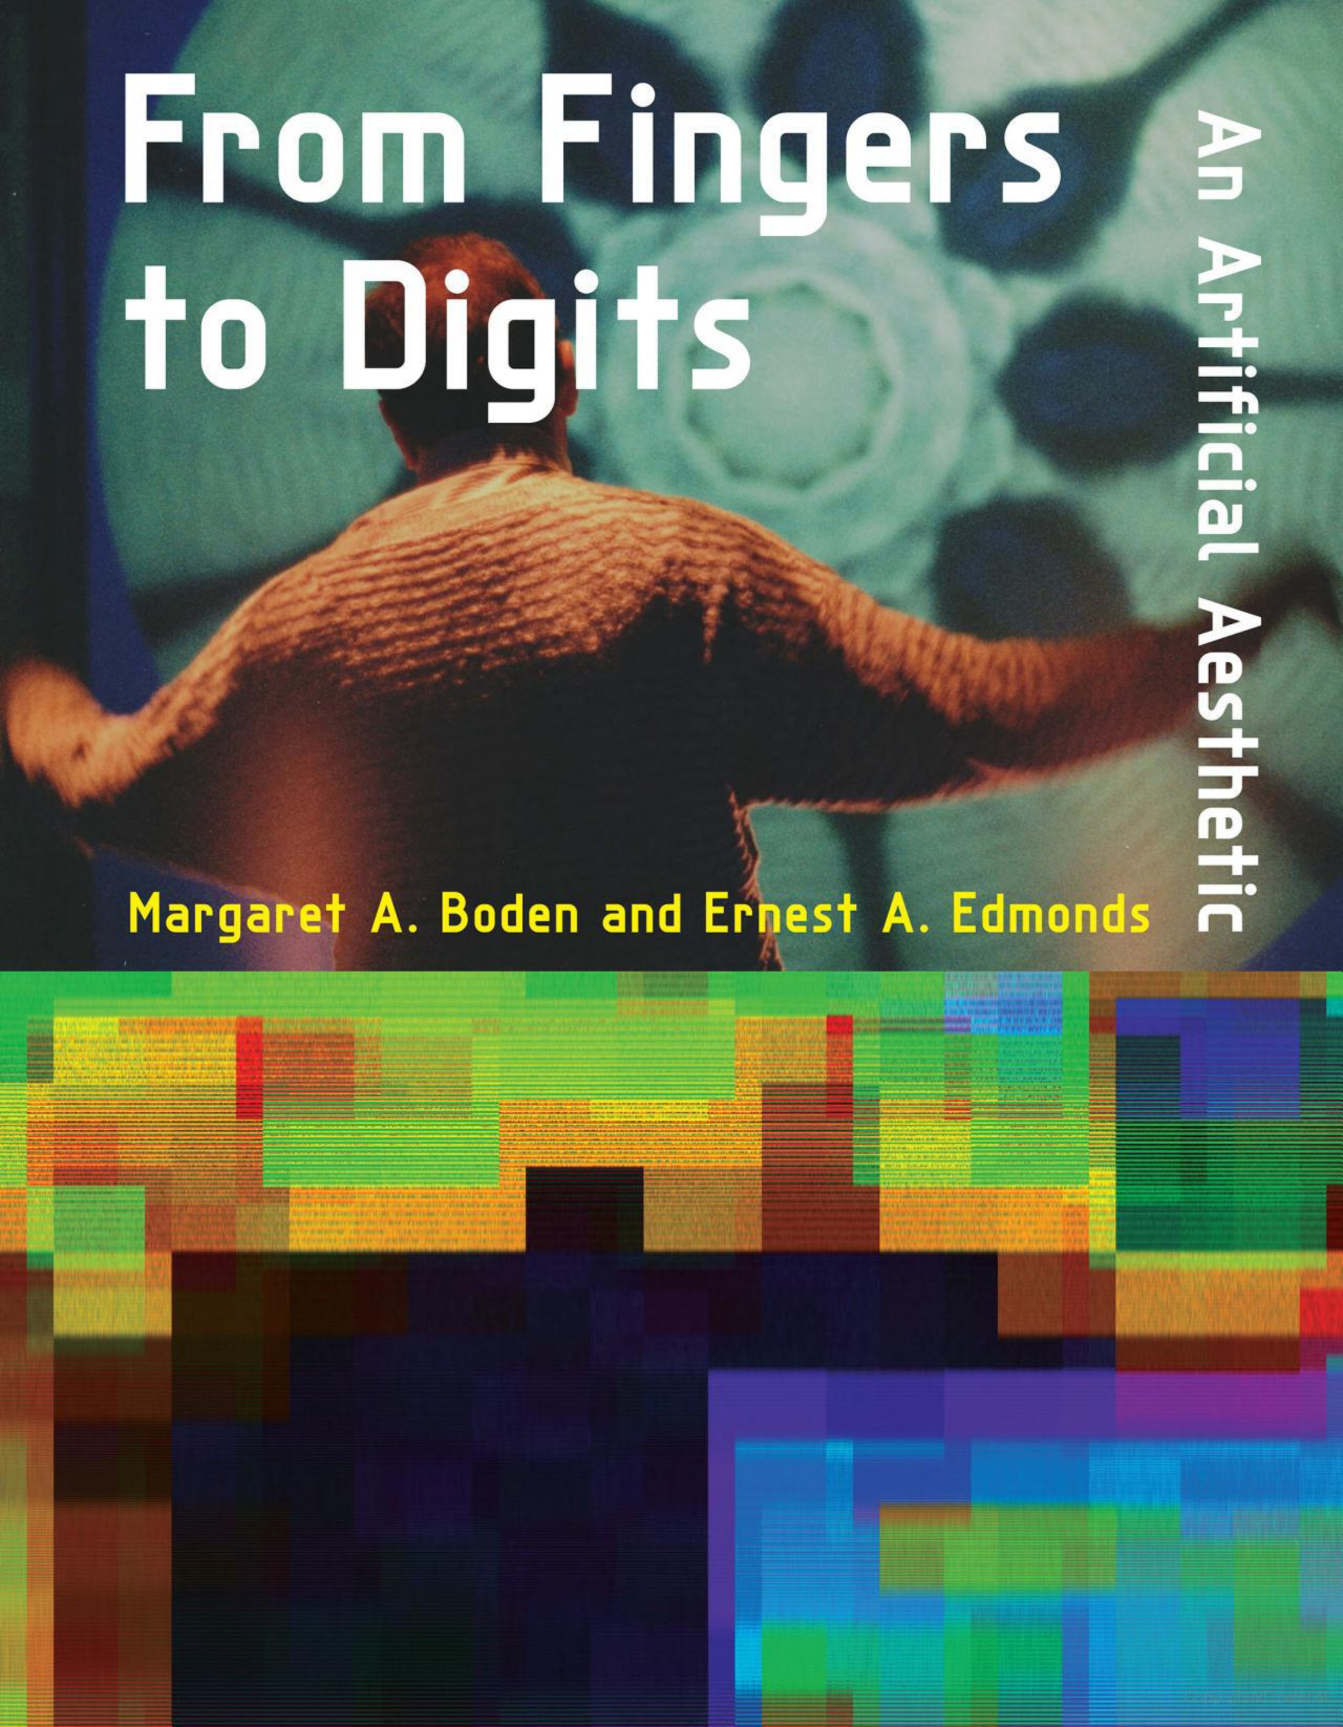

From Fingers to Digits

An Artificial Aesthetic

Margaret A. Boden and Ernest A. Edmonds

The MIT Press
Cambridge, Massachusetts
London, England

This book was set in Stone Sans by Westchester Publishing Services. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Names: Boden, Margaret A., author. | Edmonds, Ernest A., 1942- author.
Title: From fingers to digits : an artificial aesthetic / Margaret A. Boden and Ernest A. Edmonds.
Description: Cambridge, MA : The MIT Press, 2019. | Series: Leonardo book series | Includes bibliographical references and index.
Identifiers: LCCN 2018029639 | ISBN 9780262039628 (hardcover : alk. paper)
Subjects: LCSH: Computer art. | Art--Philosopohy.
Classification: LCC N7433.8 .B62 2019 | DDC 700.1--dc23
LC record available at https://lccn.loc.gov/2018029639

10 9 8 7 6 5 4 3 2 1

Contents

Series Foreword
Leonardo/International Society for the Arts, Sciences, and Technology (ISAST)

Leonardo, the International Society for the Arts, Sciences, and Technology, and the affiliated French organization Association Leonardo, have some very simple goals:

1. To advocate, document, and make known the work of artists, researchers, and scholars developing the new ways in which the contemporary arts interact with science, technology, and society.
2. To create a forum and meeting places where artists, scientists, and engineers can meet, exchange ideas, and, when appropriate, collaborate.
3. To contribute, through the interaction of the arts and sciences, to the creation of the new culture that will be needed to transition to a sustainable planetary society.

When the journal *Leonardo* was started some fifty years ago, these creative disciplines usually existed in segregated institutional and social networks, a situation dramatized at that time by the "Two Cultures" debates initiated by C. P. Snow. Today we live in a different time of cross-disciplinary ferment, collaboration, and intellectual confrontation enabled by new hybrid organizations, new funding sponsors, and the shared tools of computers and the Internet. Sometimes captured in the "STEM to STEAM" movement, new forms of collaboration seem to integrate the arts, humanities, and design with science and engineering practices. Above all, new generations of artist-researchers and researcher-artists are now at work individually and collaboratively bridging the art, science, and technology disciplines. For some of the hard problems in our society, we have no choice but to find new ways to couple the arts and sciences. Perhaps in our lifetime we will see the emergence of "new Leonardos," hybrid creative individuals or teams that will not only develop a meaningful art for our times but also drive new agendas in science and stimulate technological innovation that addresses today's human needs.

For more information on the activities of the Leonardo organizations and networks, please visit our websites at http://www.leonardo.info/ and http://www.olats.org/. Leonardo

Books and journals are also available on our ARTECA art-science-technology aggregator, http://arteca.mit.edu/.

Roger F. Malina
Executive Editor, Leonardo Publications

1 Introduction

Margaret A. Boden and Ernest Edmonds

This book is about computer art and its relations to art of a more traditional kind. Specifically, the book focuses on generative art, in its various forms, but particularly on art for which the artists use programming, computer code, as a significant element of their work. We also deal with the context of such art, so some chapters consider much of what is often called digital art, but the focus remains on the generative. Much has been written about digital art in general, and we do not repeat what is easily read elsewhere except when understanding the context requires it. At the core of computing is software—computer programming and code. Perhaps rather strangely, the role of code in computer art has received relatively little attention; hence it is our primary focus. We aim to highlight important continuities as well as the many exciting differences.

A wide range of questions is discussed, for the book is a collection of chapters—most newly written for this volume—by authors with different backgrounds. Margaret Boden is a philosopher, with a special interest in creativity and in how concepts drawn from artificial intelligence (AI) can help us understand it. Ernest Edmonds is a pioneering computer artist and a professional computer scientist. Although each chapter can stand alone, they are unified by the authors' shared interest in the history and philosophy of computer art.

In preparing the book we interviewed a range of artists and curators and had more informal conversations with many people engaged in the art world, most with a specific interest in computing. Perhaps it is not surprising that we have found significant differences among generations. Whereas older people, including many of those who form the art establishment, see computer technology as almost disturbingly new, and not infrequently question the possibility of computers contributing to art at all, young people take it for granted that computers are an integral part of almost all aspects of life, including art. Senior artists may talk about the new opportunities that the computer gives them, whereas young artists were often writing computer programs years before they ever thought about becoming an artist. For them, using a computer to make art is as natural as using paint.

Notwithstanding the comfort that the young have with computers and computer software, we still find that computer art is largely presented and viewed with suspicion in the conventional and established art gallery world. In this respect, it seems that visual art is behind music in adopting the new technologies. A good deal of modern music involves the use of computers in one respect or another. Apart from the major digital art events that are staged around the world, internationally a small number of commercial galleries specialize in computer art and a small number of public galleries or special wings or units in public galleries also present such work. Venues are certainly growing, but even when the growth is strong, more often than not the computer art is not integrated into traditional art exhibitions. So we find that the problem of computer art (is it *really* art?) is very much a current problem, even though the young may not agree and even though almost every film, much theater, and a great deal of music is today made with the aid of computers. When we come to chapter 14 and hear directly from artists themselves we do not find any doubts at all. However, there is often difficulty in dealing with the established visual art world.

We see, then, that the question of whether computer art should rightly be called art at all is still with us. Are the continuities that link it to older forms of art strong enough to settle that question? Or can it be settled merely by pointing out that several major art galleries have exhibited examples of this work? Is it relevant that other modern technologies (photography and film) have been welcomed by influential art critics, such as Walter Benjamin (1968), as ways of doing—and democratizing—art? Or does the use of computers compromise this field's aesthetic credentials in ways that the use of cameras does not? If so, how? In particular, is the compromise even greater in cases of *generative* computer art, in which the program is left to produce the artwork by itself, with little or no input from a human being? And if not, why are so many members of the general public skeptical about the very idea of computer art? In addressing these questions we show that some of the links with older art forms are quite strong even though they may not often be noted. We also give examples of very strong public interest in interactive art, although the naming of the work as art may not always be particularly evident.

Some closely related questions concern its history. Besides reviewing how the various types of computer art developed, we say something about how they have been met—or rejected—by denizens of the mainstream art world.

Mainstream art comes in many forms: painting, sculpture, ceramics, music, dance, poetry, drama, and more. So too does computer art. But these chapters highlight the *visual* arts. Certainly, the philosophical points made in the more analytical chapters apply to art in general. Moreover, one chapter is largely about music, and several mention examples from other genres, such as literature. Nevertheless, most of the discussion draws on art (whether mainstream or computer based) in the form of painting or graphics.

The six chapters in part II, by Margaret Boden, are broadly philosophical in tone. The first two set the scene for what follows—by identifying distinct types of computer art and by describing how creativity has been studied (or ignored) in AI. The other four focus on links between computer art and traditional views in philosophical aesthetics. Part III—comprising six chapters by Edmonds—considers some historical questions from the point of view of someone who has been a practitioner of this genre right from its start. In addition, his chapters touch on what the experience of creating computer art is like and how it differs from straight painting or musical composition. Part IV contains a single chapter that comments on many of the questions raised in the previous chapters in a set of interviews (conducted by Edmonds) with well-known computer artists, all of whom write code. We present many of the issues discussed in the book in the artist's voice: what it means to make art today using the media of today.

Chapter 2 provides a taxonomy of computer art. This analysis distinguishes different types within the broad genre of computer-based art and outlines some of its links with more familiar forms.

The term "computer art" is sometimes defined more narrowly. For instance, Dominic Lopes (2009) considers only the interactive examples. By contrast, the taxonomy offered here highlights the *variety* of approaches that can reasonably be included under this name. For example, it distinguishes generative art (G-art), computer-generated art (CG-art), interactive art (I-art and CI-art), robotic art (R-art), evolutionary art (Evo-art), computer-assisted art (CA-art), and live coding art (LC-art). As we will see, these subgenres are not mutually exclusive but sometimes overlap.

In addition, chapter 2 indicates what philosophical questions and problems are related to this or that subgenre. In this context, the most important distinction is between CA-art and CG-art.

CA-art, in which the computer is used by the human artist as an aid in the art-making process, raises no philosophical problems that do not also arise given the use of paintbrushes or chisels. That is not true of CG-art, in which a computer program is left to run by itself, with minimal or zero interference from a human being. Here, the relative autonomy of the computer threatens to usurp (aspects of) the role of the human artist. Our main interest, in all the chapters that follow, is in the generative examples. That is partly because Edmonds is himself a CG-artist and partly because CG-art presents many more philosophical puzzles than CA-art does.

Different types of CG-art, moreover, suggest different philosophical worries. So the taxonomy is followed by a sketch of the relevant philosophical landscape. Some of its peaks and valleys are considered more fully in later chapters.

Next, chapter 3, “Explaining the Ineffable,” outlines the history of AI research on creativity.

Much as people often raise philosophical doubts about whether AI systems are really intelligent, so they often doubt whether they can really be creative. Some even question whether AI can offer the *appearance* of creativity. Taking those questions seriously requires one to consider the research mentioned here.

AI research on creativity has blossomed in the last twenty years. But although creativity was identified very early as a key goal for the field, it was not seriously studied for several decades. Insofar as creativity was modeled at all, it was exemplified as verbal or logical problem-solving—hardly ever as music or graphics. Even scientific creativity had to wait many years before becoming an explicit focus for AI.

As it happened, the two people who were most influential in bringing art into the AI arena were both computer artists: the painter Harold Cohen and the composer David Cope. However, AI work on artistic creativity should not be confused with computer art. Often, there is no specifically artistic intention, still less any attempt to exhibit the results in galleries or concert halls. Rather, the interest lies in trying to model the types of thought processes that go on in the minds of human artists when they create. Indeed, Cohen himself, when already a highly acclaimed painter, turned to AI largely because he thought it would help him understand his own creativity. By the time of his death, in 2016, he had come to believe that human and AI creativity are fundamentally different. Even so, he felt that his computer-based work had helped him appreciate this.

That point leads to the key message of the chapter, which is that—contrary to the still-influential Romantic myths—creativity is a psychological capacity that can be understood by science. Specifically, it involves computational processes comparable to those defined in AI.

Three types of psychological mechanisms can produce new thoughts: combining familiar ideas in unfamiliar ways, exploring accepted styles of thinking or making so as to generate novel structures of the relevant kind, and transforming the style to come up with structures of a novel kind (Boden 2004). Each of these three kinds of creativity can be seen in the work of mainstream human artists and in computer art too. And each, as explained in chapter 3, has been modeled in—and clarified by—AI research.

There is an important caveat here, which applies also to several other chapters (e.g., chapter 5). Whether a computer artwork or AI program can *appear* to be creative is a factual question—to which the answer is often yes. Whether it is really creative is a philosophical question, not a factual one. According to many philosophers, the answer must be no. Irrespective of its actual performance, they say, no computer can provide real creativity.

The reasons advanced for this view are diverse and involve some of the deepest problems in philosophy—including the nature of meaning, consciousness, and moral community. They will be ignored in this book (unless mentioned by any of the artists interviewed in chapter 14), whose focus is on what computer art can achieve in fact. For the record, however, questions about whether computers can be really creative cannot be given a definitive answer, precisely because the philosophical issues involved are so deeply disputed (Boden 2004, chap. 11; Boden 2014).

Chapter 4 considers the broadly accepted assumption that art appreciation normally involves recognition of the making skills concerned. This assumption dates from the predigital age. Indeed, John Ruskin—whose nineteenth-century views on aesthetics are still hugely influential—*defined* art largely in terms of the human skills that it employs (see chapter 7). Even Robin Collingwood, who (as explained in chapter 6) drew a sharp distinction between art and skill, allowed that we often do admire the skill of the person who made the work of art and that we typically value the artifact all the more accordingly. In short, it is nowadays broadly accepted as common sense that art appreciation normally involves recognition of the making skills concerned.

This fact, with its implications for the appreciation of computer art, is the theme of chapter 4. The appreciation of computer art is even more dependent on knowledge of the making process than is the case for traditional (or anyway, non-Constructivist) forms of art.

CG-artists not only *use* the computer but *celebrate* its potential as a medium. Similarly, oil painters may celebrate the colors and textures of their paints and jazz clarinetists the potential of their instrument. Moreover, the different forms of CG-art use the computer in very different ways. It is not merely that individual CG-artworks look (or sound) very different from each other but rather that they can exhibit very different aspects of the indefinitely varied generative potential of the digital computer.

That is all very well, and the multifariousness offered by computers is a genuine cause for celebration. But it raises a problem for the art audience, most of whom have never done any programming and have little or no intuitive sense of what it involves (see chapter 10). Even people who (like most members of industrial societies) use computers very frequently indeed may have no such understanding, because they experience computer programs ready-made. The complex skills that underlie user-friendly applications such as spreadsheets and the commercial search engine are invisible, even unsuspected, by their users. As the computer artist Harold Cohen has put it (personal communication), most of us are *consumers* of programs, not *producers* of them. Thus, we are not well placed to appreciate, criticize, or even compare individual works of CG-art with respect to the programming achievements involved in them.

Up to a point, computer artists and curators can arrange to make these technical achievements more accessible. Several methods of doing this are suggested in chapter 4. But even in a culture whose children are introduced not only to computers but also to simple programming, the extent to which this can be done is strictly limited. The unfamiliar, elitist skills of computer artists are even less imaginable to the general public than the specialist skills of traditional artists are. We can all remember using a paintbrush, and we can all imagine forging wrought iron. Whether we can picture how to program computers to generate interesting graphics or music or to operate robots is another matter entirely.

The conclusion of chapter 4, then, is decidedly downbeat. Computer art can be enjoyed by anyone. Some examples can engage virtually everyone. But it can never be as fully appreciated as the more familiar types of art. It will always carry an aura of magic: awesome, perhaps, but largely unintelligible.

Chapter 5 explores the potential for creativity in one particular form of computer art, evolutionary art (Evo-art in the taxonomy of chapter 2). The key question is posed by its title: "Can Evolutionary Art Provide Radical Novelty?" In other words, is the creativity of Evo-art forever limited in a way that human creativity is not?

As remarked earlier (and in chapter 3), one of the three types of creativity is transformational creativity. In this, some initial style is transformed with the result that structures of a novel kind can arise. In general, the more radical the stylistic transformation, the more it is valued. Although if the change is too radical, it may be rejected. If computer art cannot offer us transformational creativity, then it is less worthy of our appreciation than human art is.

The computer artists who practice Evo-art hope, by using genetic algorithms, to effect stylistic mutations that vary the resulting artifacts (e.g., images) in highly unpredictable ways, perhaps even jumping out of one style into another. If the second style is not just new but radically new, so much the better.

But is that, in principle, possible? Many critics argue that it is not, because the computational potential of any program—including one equipped with rule-changing genetic algorithms—is fixed. It may be huge, even infinite (a point that is important in chapter 4). But there are always some conceivable structures that are excluded. Evo-art, on this view, simply cannot provide *truly radical* transformations.

Such transformations do sometimes happen, but only (so this objection runs) outside cyberspace. For instance, they occur in biological phylogenesis, in which fundamentally new organs (e.g., light sensors) can arise—and once arisen, can then evolve into highly differentiated forms. The explanation given for this is that biological evolution takes advantage of the physical contingencies in the environment. Because cyberspace involves no such contingencies, no opportunity for an evolving computational system to engage with the physical world, this spur to stylistic novelty is forever absent.

Robots, however, do not inhabit cyberspace. They not only take in perceptual data from the real world but are material objects moving around in a material, often cluttered, environment. Chapter 5 describes some work in evolutionary robotics wherein radical (anatomical or perceptual) transformations have happened as a result of previously unconsidered physical conditions, even accidents. It also mentions ongoing research by a leading computer artist who hopes to evolve drawing robots that, as a result of physical interactions, generate fundamental, but aesthetically valuable, stylistic changes.

Biology is not the only noncomputerized forum for radical transformation: history-book artistic (and scientific) creativity in human minds and societies counts too. But here, *physical* contingencies are only rarely relevant. Usually, the relevant environment is psychological or cultural. These factors can be modeled in computer systems. And they do not have to be specifically provided beforehand by the programmer: some (non-Evo) computer artists already allow their programs to take current data—words or images—at random from the Internet. One such example is *The Living Room*, an installation built by Christa Sommerer and Laurent Mignonneau.

In short, the message of chapter 5 is that Evo-art could provide radical transformations. Some of them would doubtless be rejected by us, but that is another matter. It could do this even if its artifacts were situated not in the real world but in cyberspace. The fact (repeatedly cited by skeptics about computer art in general) that *a computer program can do only what it is set up to do* does not exclude this possibility. Certainly, if an Evo-art installation were to benefit from the cultural contingencies available on the Internet, the web's contents would have to be accessed by some instruction in the program. But that is just to say that creativity—even radically transformational creativity—does not happen by magic. Evo-art has an open, and intriguing, future.

The remainder of part II turns to key topics in philosophical aesthetics. Many competing accounts of art have been offered by philosophers, including various modernist and postmodernist positions—not to mention the still-influential views of Plato, Aristotle, and Kant. This rich body of literature is touched on in several chapters but is foregrounded in chapters 6 and 7.

Fashions in theoretical aesthetics change. Constructivism, for instance—which led most of the pioneers of computer art to experiment with the new medium (see chapters 9 and 14)—stressed aspects of the art-making process, and of the finished artwork, that had not been valued before the early twentieth century.

The conceptual art of the 1960s went even further, for it challenged most of the then-accepted assumptions about what counts as art and even what sort of thing an artwork can be (Boden 2007). According to that viewpoint, *many* computer artworks are not art at all—or anyway, not of an interesting kind. Even CG-art, wherein the core of the artwork is often held to be the *process* of art making (see chapter 15), does not necessarily qualify. It does not normally offer a conceptual outrage—a shocking conceptual combination (see Boden 2007)—to the audience. Or rather, it does not offer any *specific* shock to the audience—as when the conceptual artist Walter de Maria, for instance, buried a perfectly crafted and very expensive object (a two-inch-thick steel cylinder, one kilometer long) so as to make it invisible.

For some audiences, however, CG-art *in general* is irredeemably suspect, just because its artworks are made by computers not by human hands. Are those audiences wrong? Perhaps so, but this cannot be taken for granted: the point needs arguing.

In particular, one needs to consider the aesthetic theories—or rather, the largely unexamined philosophical assumptions—that are most widely held by the general public (and within the art world; see chapter 8). These concern the role of *emotion* in art and the role of the art maker as *a human individual*—the topics of chapters 6 and 7.

The theme of chapter 6 is the relationship between computer art and Collingwood's theory of art. Why Collingwood? Well, not because either of us accepts his views on art; we do not. Rather, because it is a common—perhaps the most common—criticism of computer art that *there can be no such thing*, because of the importance of emotions in art. Because computers cannot have emotions, it is said, they cannot be artists either. And Collingwood (1958) is among the most influential of the many philosophers who have claimed that art necessarily involves emotion.

Prima facie, then, his account challenges those who claim that CG-art is, truly, art. Indeed, his ideas underlie much of the hostility to computer-generated art that is found in our culture—not least its uneasy position vis à vis the art world (see chapter 8). If it could be shown that CG-art is not quite as inimical to Collingwood's theory as it appears to be at first sight, the widespread resistance to it would be weakened.

Chapter 6 argues that his approach does indeed threaten the status of computer art as art. But that is not because of any reference to emotions in general. There is no reason in principle why computer art could not express or arouse emotion, and some of it—especially computer music—actually does. Rather, it is Collingwood's extreme *particularism* that prevents CG-art from being (on his view) art proper.

He defined art in terms of the expression—and the detailed construction—of some highly particular emotion in the artist's mind, with an equivalent process occurring also in the mind of the audience. He even said that to regard a novel as depicting "the feelings of women, or bus-drivers, or homosexuals" is not to treat it as *art* at all but rather as psychology (Collingwood 1958). As a work of art, it must construct or express the particular experiences of the individual novelist—who may happen to be a woman, a bus driver, or a homosexual but whose work cannot be regarded (qua art) as representative of the group.

Collingwood's view goes against the spirit of most computer art. That is because a large part of the point of CG-art—especially when the originating human artist has a relatively hands-off role—is to explore, and to exhibit, a certain *range* of possibilities. Each CG-artwork is unique, to be sure. But those generated by a given computer program are unified—both perceptually and historically—by it. The audience is expected to notice that unity and to appreciate it. Whether this appreciative response requires knowledge of computing, as well as sensitivity to observable stylistic variation, is discussed in chapter 4.

Given this conflict, the question arises whether we should say "So much the worse for computer art!" or rather "So much the worse for Collingwood!" The alternative chosen in chapter 6 is the latter. For besides undermining the rationale of CG-art, Collingwood's position also conflicts with common intuitions about traditional art.

Shakespeare's plays, for example, are not art proper for Collingwood—although his sonnets may be. They do not, except perhaps occasionally, construct and express some highly specific emotional experience of Shakespeare himself. On the contrary, they often aim at depicting the emotions of some general class of people—monarchs, for instance ("Uneasy lies the head that wears a crown"). Similarly, a painting that aims to make a political point—such as Pablo Picasso's *Guernica*, or Gustave Courbet's *Peasants of Flagey Returning from the Fair*—is seen by Collingwood as propaganda rather than art unless it happens to capture the painter's particular experience and arouses an equivalent emotion in the audience.

The last chapter in part II, like chapter 6, engages with a writer in traditional aesthetics: Ruskin. His approach is prima facie even more inimical to computer art than Collingwood's is—which makes it even more interesting as a philosophical challenge to the field. That is primarily because it stressed the personal individuality, and the creative freedom, of the human artist-craftsman.

Much as we do not share Collingwood's theory of art, we do not share Ruskin's either. But we do admit to having some sympathy with his view, for it has entered deeply into our culture.

In particular, Ruskin's chapter on Gothic architecture in *The Stones of Venice* (1853) was hugely influential. It molded people's ideas about art in general, remaining "the greatest influence on taste" until the early twentieth-century doctrine of pure form

(Clark 1949). For instance, it insisted on the aesthetic importance of the artist's cultural (e.g., religious) motivation and on his chosen methods of hand-working. By inspiring William Morris, it led indirectly to the Arts and Crafts movement. And, again largely via Morris, it also affected the development of socialism in England.

The ideas expressed in that famous essay are still widely current. Many of us today share, or anyway sympathize with, Ruskin's stress on the individuality and fallibility of the working artist and on the importance of the cultural interests that imbue the artwork. Indeed, the suspicion of computer art in the mainstream art world today is largely grounded in aesthetic values inherited from him. Ruskin's essay was an impassioned complaint against the industrial age and especially against the effects of machine technology on the workman's role in art and craft.

High-tech computer art is very different from the medieval masonry so admired by Ruskin. There is no question but that it fails to satisfy his criteria for superior art. However, chapter 7 argues that this new approach is not quite so beyond Ruskin's pale as one might think.

For example, the individuality of the originating artist still shines through, even given the mediation of the computer (see also chapter 10). Ideological values, such as religion and politics, are not typically involved in CG-art—although, in principle, they could be. They are often prominent in other forms of what chapter 2 terms electronic art, or Ele-art: video-art, for instance. But if CG-art normally eschews ideology, it allows plenty of room for infelicities and failures on the part of the originating human artist. And computer-*assisted* art (CA-art), of course, leaves even more room for mistakes by the human individual. It may seem strange to laud mistakes and infelicities as criteria of art: surely they are to be avoided if at all possible? That was not Ruskin's view. And we need to understand his legacy to understand much of the resistance (especially in England) to this new way of art making.

Ruskin identified six principles of the Gothic, which he used to evaluate all forms of art. More accurately, perhaps, he identified twelve such principles: two parallel sets, concerning the form or content of the artwork itself and the moral climate expressed in it by the artist. We will see that computer art can satisfy his six, or twelve, principles only up to a point. To the extent that it fails to do so, computer art would be regarded by him as an inferior type of art.

Or rather, he would have regarded it as inferior *if* he had accepted it as being art in the first place. The final section of chapter 7 shows that he would not have done that. And his reasons for not classifying computer art as art are shared by many people today, some of whom even cite Ruskin in making their argument. In brief, his nineteenth-century ideas remain very relevant to our twenty-first-century concerns.

The six chapters in part III touch on many of the philosophical points already mentioned. But because of Edmonds's long experience as a computer artist, they are more practical and historical in tone.

Chapter 8 explores the sociological relations between computer art and the art world and considers the changes in these relations that have taken place over the last half century. The attitudes of the major art galleries, for instance, are relevant—and, as we will see, deeply ambivalent.

Sociology and history, in these respects, are closely linked to philosophy. One influential account of what it is for something to count as art is that it be accepted as such by the art world (Danto 1964, 1981; Dickie 1974, 1997). On that view, what counts as art depends on the collective response of the socioeconomic network comprising artists, critics, museum curators, gallery owners, dealers, and collectors—and, where relevant, publishers too (Becker 1986, 1995).

Acceptance or rejection by the art world, of course, rests largely in the application of a range of familiar aesthetic criteria (autonomy and authenticity, for instance; see Boden 2010b, 2010c). But it depends also on judgments, and even prejudices, of other types.

For example, one historian of computer art has remarked that, although "artists throughout history have appropriated new technology and manipulated it to serve their own purposes," the technology used in computer art "has tended to be viewed by the art world as deriving from military applications" and that "perhaps this is [one] reason why the art world has been reluctant to embrace it" (Mason 2008, xii). A related source of skepticism is the role played by computer technology in modern methods of mass production (see chapter 14).

Yet another prejudice against this new art form arises (as remarked previously, in relation to chapter 6) from the widespread assumptions that art and art making must involve emotion and that this can have nothing whatsoever to do with computers.

Even social snobbery has played a part. Computer art in Britain was pioneered predominantly in the newly formed polytechnics (although the Slade School of Fine Art is an important exception), which did not endear it to artists and scholars identified with the major art schools and the universities (Mason 2008, chaps. 7 and 8).

Not least, of course, is the always-present arts-science divide. Despite having been included in the first international exhibition of computer art in 1967, and despite having been nominated for an Arts Council of Great Britain award by the chairman of the Institute of Contemporary Arts, an early kinetic sculpture by Edward Ihnatowicz was rejected by the relevant committee because "the Committee took the view that…there was not really enough evidence of its prospects of success as a work of art as distinct from a piece of electronic ingenuity" (Mason 2008, 96).

So in chapter 8, Edmonds recalls some of his experiences vis à vis the art world over the last fifty years and reports on interviews with individuals situated in different parts of this variegated curatorial and commercial world. Some very recent and—to both of us—inappropriate responses to computer art are discussed. Encouragingly, at least a few of the experts had positive and constructive visions to share.

The next two chapters are not only historical but also partly autobiographical. In chapter 9, "Formal Ways of Making Art: Code as an Answer to a Dream," Edmonds provides a historical context for computer art and describes how, as a young man and already a committed painter, he welcomed the new technology of computing as a way of advancing certain artistic concerns that had been informing his work.

The new concepts and capabilities that the computer introduced were very appropriate in a particular art-making tradition that preceded the invention of computers. It can be seen in the discussions of the General Working Group of Objective Analysis in 1921 in Moscow. This was the beginning of the art movement generally known as Constructivism.

The group drew a distinction between composition and construction in making their art. Briefly, composition was seen to be about arranging perceptible forms according to relationship rules, and construction was about making a work according to a plan for its production. The key point was the introduction of the notion of making a visual artwork according to a plan, or underlying structure, that took creative precedence over the final object. Satisfying the goals or constraints guiding the art-making process was considered to be more important than concentrating on the visual (or audible) composition of an art object.

In 1921, of course, the plans were executed by the artists themselves, but in CG-art they are executed by computers. Constructivist work continued through Systems group art and other groupings, leading on to generative art. What art historians term "generative art" includes any approach wherein the realization of an underlying plan, or abstract structure, is highlighted. Full appreciation of the work therefore depends on knowledge of the creative processes involved (see chapter 4). Early examples, although rarely discussed in this context, include the dice music of Haydn and Mozart. With the arrival of computers, the generative artists found the answers to the prayers that their constructivist predecessors had been making.

Chapter 9 goes on to give the particular example of the developments in Systems art in the United Kingdom, influenced by Max Bill, for example. It explores how making art with software grew out of that movement. The UK's Systems artists both directly and indirectly informed recent software-based art in the UK.

If programming can aid art, as declared in chapter 9, it can also count as art—or so Edmonds argues in chapter 10. Here, he draws on his long-standing personal experience of computer programming, for both artistic and nonartistic purposes, to emphasize the significance of software in art. The concerns of software-based art include a significant conceptual element. The computer enhances the artist's ability to shape the underlying structures of artworks and art systems. In some artists' hands the software goes well beyond the role of a manufacturing tool. It is also a thinking aid or inspiration and increasingly must be seen as a new medium in itself.

The art in software is increasingly recognized. Software in art has been underrated, however. Even many of the books on digital or computer art pay little attention to the computer code that drives and largely defines such art. That is largely because most people, including most art critics, have little or no experience of programming; to use a computer is not necessarily to program one. The artist's challenge is software, not because it is difficult but because it is the conceptual representation of the new art.

Chapter 11, "Diversities of Interaction," discusses the increasing concern with the active role of the audience that has been encouraged by computer art. The audience can interact with an artwork in many ways. We have long understood that perception is an active process so that even contemplating a Mark Rothko painting involves interacting with it in a certain sense. In that example, of course, there is no material change in the painting when it is looked at. But many artworks do, in fact, materially interact with their environment—for example, with shadows changing as the sun moves round a sculpture or with a breeze changing the configuration of a mobile. With the advent of computers in art, however, there has been a significant growth of interest in observable interactions with artworks.

In these cases, we can see members of the audience move in some way and can also see the artwork change as a result. Such observable interaction can be considered in the light of action-response models. In other words, the causal effects of the gallery visitor's behavior on the artwork are relatively direct. They may even be open to intentional direction by the audience members, in which case the audience actively participates in the creation of the final artwork (see also Boden 2010a; Edmonds 2011).

A further development is based on a systems view, in which the audience and the artwork are seen as collections of systems that exchange information. Stroud Cornock and Edmonds (1973) presented the concept of seeing the artwork, the object itself, and members of the participating audience as interacting systems and suggested that the art system might be seen as the collection of all these things. Going one stage further, Sean Clark and Edmonds showed *ColourNet* at the CHI (Computer-Human Interaction) conference in Paris in 2013 (Clark and Edmonds 2013). This art system allowed separate, possibly mobile, artworks by others to be included within the distributed system. The version demonstrated at CHI 2013 consists of two digital artworks—the core *Shaping Form* component by Edmonds and the *Transformations* component by Sean Clark. *Shaping Form* is projected on a large screen. Participants access *Transformations* on their smartphones. Both components are able to interact with multiple users simultaneously and can work independently. However, they are also able to interact with each other by exchanging color information via the Internet and using that information to influence the colors used.

In works such as *Shaping Form*, an action does not necessarily produce an immediate response. It may change the internal state of the other system and so influence its behavior later in time. The causal influence from audience to artwork is no less real but is much less direct—and often unfathomable.

As for chapter 12, "Correspondences: Uniting Image and Sound," this is the only one that focuses on or beyond nonvisual art. An important form of interaction that has developed considerably as the result of computer programming is interaction between different media. For instance, artworks that integrate visual displays with sound are common today. Such work is often linked with an interest in synesthesia. However, the link is more one of artistic inspiration than science, because the condition (color-and-sound mapping, for example) does not seem to be consistent across individuals.

The key point explored in the chapter is that the computer enables real-time dynamic linking between the generative production of sound and of image. For instance, software technology can analyze live music as it is played, detecting pitch and amplitude (for example) and feeding that information into a computer-based art system that responds to or transforms those sounds—producing music, images, or forms in many other media. A tight integration of media is thus made possible in ways that were imagined but not realized in earlier times.

Chapter 13 "Diversities of Engagement," describes how computer-based interactive art has encouraged a concern for better understanding of audience engagement. At the simplest level, this is a measure of how long someone looks at something: maybe a painting or a sculpture or a video. In the museum world, such simple measures are used to judge how attractive an exhibit is—or perhaps more accurately, how attracted to an exhibit the members of the public are.

The concept of engagement can be elaborated, however. For instance, at least three kinds of engagement can be identified. These are attraction (the work drawing attention to itself), sustainability (the length of time that someone finds it engaging), and repeatability (the extent to which someone wants to experience the work again and again).

Artists who make digital works that take engagement into account become involved in strategies for drawing the audience in and retaining their interest in quite explicit ways. This process can be one that includes an element of audience education. The artwork can be designed so as to attract the public and interact with them in ways that develop interest and engagement. Specific aspects encourage that result. Examples of this can be seen in the curated program of interactive art exhibits and their evaluation in *Beta_Space*, a joint venture of the Powerhouse Museum in Sydney and the Creativity and Cognition Studios at the University of Technology, Sydney (Candy and Edmonds 2011).

The artworks that are made for the new audiences for whom these artists are aiming can be ambient. That is, they may exist in the environment without any specific signals indicating that they should be seen as art. Of course, the context itself can be seen as a strong indicator in this sense. So we might want to distinguish between such art in a museum and art in the street, for instance. Engagement and ambient art are concepts that are important in the digital art world, but they have good continuity with the nondigital world. Dynamic and interactive artwork that is an integral part of an ordinary environment, such as a meeting space, bar, or city square, and that is made by use of computer systems is growing.

Finally, part IV, consisting of only the substantial chapter 14, provides a series of conversations with eleven computer artists. Most focus primarily on the visual arts, although a few of them sometimes integrate or combine vision and music in their work. All the artists are established, even long-established, professionals in the field, although some are known only in very specialized communities. This section of the book presents the artist's voice: the concerns of the book as seen from the point of view of the practicing artist.

In leading these conversations, Edmonds presents the artists with the key questions raised in the previous chapters. As one might expect, there was little or no comfortable consensus on many of these. On some, however, there was a significant measure of agreement. One obvious difference was closely related to age. The older, pioneer artists discovered computers and programming when already into their career and, for various reasons, became engaged in ways that extended or changed their artistic practice. At the other extreme, younger artists report having used computers, including programming them, before they dreamed of becoming artists. For them, the computer did not change their art, because it was always there.

The artists interviewed here work in the visual domain, in performance, and in sound and some work across these media. A strong theme is the importance of programming. The computer code is significant in several ways. For at least one of them, the code *is* the art. For several, writing it is a vital part of the thinking process. They would never make the kind of art that they do without going through the thinking process required by programming. For them, software is a way of both making art and thinking about art. One of the artists puts a significant amount of effort into constructing the technology that is used in making the art. At least two others even find the need to make their own software tools. A performance artist finds that the specification of dance movements can be inspired by and even determined by the code.

Several artists, particularly those involved in live coding art (LC-art, in the taxonomy of chapter 2), are very interested in to what extent the audience should see the code and to what extent it is desirable that they understand it. One of these is designing a programming language with the very purpose of facilitating audience understanding of the software that drives the work.

It is clear from the interviews in the final chapter that computer artists' aims, methods, and beliefs differ significantly and also that they have often changed over the years. We hope that this book will inspire its readers to consider computer art from both historical and philosophical perspectives and, perhaps, also move their thinking on in valuable ways. We must acknowledge that our own thinking has been partly influenced by many conversations with the late Harold Cohen, who sadly died during the later stages of this book's preparation. Stroud Cornock, who also sadly died during these late stages, inspired much of Edmonds's work. Ernest Edmonds also thanks Linda Candy for endless valuable conversations and tireless support.

Perhaps, while considering computer art, some readers might also get involved in doing it. Or at least in hanging a CG-piece on their living room wall. Impractical? Not necessarily: some small wall-mounted interactive displays show dynamically changing images and sounds, triggered by passersby. Let us help bring computer art into the wider art world!

References

Becker, H. S. 1986. *Doing Things Together: Selected Papers*. Evanston, IL: Northwestern University Press.

Becker, H. S. 1995. "A New Art Form: Hypertext Fiction." In *How to Read a Hypertext*, edited by M. Bernstein. Cambridge, MA: Eastgate Systems.

Benjamin, W. 1968. "The Work of Art in the Age of Mechanical Reproduction." Publ. in German in 1936, trans. 1968. Reprinted in R. Kearney and D. Rasmussen (eds.), *Continental Aesthetics: A Reader* (Oxford: Blackwell), pp. 166–181.

Boden, M. A. 2004. *The Creative Mind: Myths and Mechanisms*. 2nd ed. London: Routledge.

image
not
available

II

2 A Taxonomy of Computer Art

Margaret A. Boden and Ernest Edmonds

2.1 Introduction

Since the late-1950s, an ever-diversifying set of novel art practices has arisen that is still little known or discussed in aesthetics and art theory. (For examples of these practices, see Krueger 1991; Wilson 2002; Candy and Edmonds 2002; Whitelaw 2004; Woolf 2004; Popper 2007.) As Jon McCormack, one of the artists concerned, has put it, "Much of the innovation today is not achieved within the precious bubble of fine art, but by those who work in the industries of popular culture—computer graphics, film, music videos, games, robotics and the Internet" (McCormack 2003, 5).

The "precious bubble" of fine art is a (shifting) socially accepted norm. But artists often work outside the norm of their day as famously illustrated by Marcel Duchamp and his ready-mades or John Cage's use of silence. And sometimes the bubble eventually expands so as to engulf the previously maverick efforts. The Impressionists, for instance, no longer have any need for a *Salon des Refuses*. They do not even need a salon: their images assail us every day on calendars and chocolate boxes. Whether the innovations mentioned by McCormack will one day be included in the expanding bubble remains to be seen. Their fate, in this regard, depends partly on how people—both curators and the general public—respond to the controversial aesthetic and philosophical questions raised in section 2.4.

The novel approaches involved here are closely interrelated, both theoretically and methodologically. So much so, indeed, that they are often all lumped together under

---

This chapter is based on M. A. Boden and E. A. Edmonds, "What Is Generative Art?," *Digital Creativity* 20 (1–2) (2009): 21–46. The research for that earlier paper was supported by AHRC Grant no. B/RG/AN8285/APN19307, Computational Intelligence, Creativity, and Cognition: A Multidisciplinary Investigation and also by the Australasian CRC for Interaction Design, established under the Australian Government's Research Centres Programme.

a label, such as computer art, electronic art, process art, or generative art. This chapter aims to clarify how they can be distinguished.

2.2 Origins and Interrelations

From the theoretical point of view, this new art originated in cybernetics and general systems theory. The young painter Roy Ascott, later to be highly influential in the field, identified the novel activity as "a cybernetic vision" in 1966 (Ascott 2003b, chap. 3; see also Mason 2008, chap. 4). And the exceptionally creative cybernetician Gordon Pask was a key influence. Besides producing or imagining some of the first artworks of this general type (in the 1950s), he provided much of the theoretical impetus that inspired the more philosophically minded artists in the field (Boden 2006, 4.v.e; Mason 2008, chap. 2).

Very soon, the "cybernetic vision" was bolstered by ideas about structure and process drawn from computer science. Ernest Edmonds, for instance, turned from paintbrush and easel to the computer in the 1960s; he thought he could produce more interesting art in that way (see Boden and Edmonds 2009, sect. iii). At much the same time, music and visual art was produced that reflected the computational theories of mind of artificial intelligence (AI). Indeed, Harold Cohen, a renowned abstract painter in 1960s London, deserted his previous working practices largely because he felt that doing computer art would help him understand his creative processes better (McCorduck 1991; Boden 2004, 150–166, 314–315).

Over the last twenty years, this artistic field has been inspired also by ideas about emergence, evolution, embodiment, and self-organization. These concepts are borrowed from cognitive science and in particular from artificial life (A-Life). However, the theoretical roots (and the pioneering experiments) of A-Life reach back to midcentury cybernetics and automata theory (Boden 2006, 4.v.e, 15.iv–v). In short, the theoretical wheel has turned full circle.

The methodological wheel, meanwhile, has climbed an ascending spiral. Art practices outside the fine-art bubble are grounded in technologies for communication and information processing whose power and variety have burgeoned over the last half century. Often, this means that the customary lone artist is replaced by a team, some of whose members may be computer scientists or tele-engineers.

Most of them rely heavily on digital computing and in particular on methods drawn from AI and A-Life. Specifically, they have employed symbolic and connectionist computation and, more recently, cellular automata, L-systems (automatically branching structures that botanists use to study plant form and physiology), and evolutionary programming too. This is an ascending spiral, not a linear ascent, because two of those recent methods were foreseen (by John von Neumann) in 1950s cybernetics, and all three had been mathematically defined by the 1960s (Boden 2006, 15.v–vi). But none could be fruitfully explored, by artists or by scientists, until powerful computers became available much later.

The resulting artworks are highly diverse. They include music, sonics, the visual arts, video art, multimedia installations, virtual reality, kinetic sculpture, robotics, performance art, and text. And whereas some of these outside-the-bubble activities place ink or paint onto a surface, others involve desktop visual display units (VDUs) or room-scale video projection. Yet others eschew the virtuality of cyberspace, constructing moving physical machines instead.

The labels attached to these new art forms vary and have not yet settled down into a generally accepted taxonomy. The names preferred by the artists involved include generative art, computer art, digital art, computational art, process-based art, electronic art, software art, technological art, and telematics. All these terms are commonly used to denote the entire field—and although distinctions are sometimes drawn, they are often treated as synonyms.

With respect to the labels computer art and generative art, that was true right from the start. These terms have been used in tandem and more or less interchangeably since the very earliest days. The first exhibition of computer art, held in Stuttgart in February 1965, was called "Generative Computergraphik" (Nake 2005). It showed the work of Georg Nees, who wrote the first doctoral thesis on computer art, giving it the same title as the exhibition (Nees 1969). That thesis was widely consulted by the small but growing community, harnessing the words "generative" and "computer" together in its readers' minds.

Their near equivalence was reinforced in November 1965, when an exhibition (again in Stuttgart) included both Nees's work and the early computer graphics of Frieder Nake. Both men applied the term "generative" to their work and used it to identify art that was produced from a computer program and, hence, at least in part produced automatically. Others who were pioneering the activities outside McCormack's bubble also adopted the term. For example, when Manfred Mohr started producing drawings with a computer program in 1968 he termed it "generative art" (Mohr 1986). And the philosopher Max Bense—who had composed the manifesto for the original Stuttgart exhibition of 1965—was writing about "generative aesthetics" (Nake 1998).

The label generative art is still current within the relevant artistic community. Since 1998 a series of conferences have been held in Milan with that title (see http://www.Generativeart.com/), and Brian Eno has been influential in promoting and using

Such forces also influenced some visual generative art that predated computers. One clear example is Kenneth Martin, whose 1949 abstract painting used basic geometrical figures (squares, circles, diagrams) and rules of proportion (Martin 1951). Later, his *Chance and Order* and *Chance, Order, Change* series combined rule-driven generation with random choice. Although the basic forms were laid down by the rules that Martin had deliberately devised, chance physical events—such as picking a number out of a hat—determined the actual course of the work.

As for generative literature, this too may involve chance events dependent on physical processes. The various versions of Bryan Johnson's (1969) novel *The Unfortunates* are produced in this way. The novel was published as twenty-seven separate sections in a box: all but the first and last were to be read in a random order, decided by shuffling or dice throwing. Many other examples of interactive stories, partly narrated by the reader, have been produced since then (Montfort 2003). Most of these depend not on physical processes but on deliberate voluntary choices by the reader-author; however, some are partly generated by dice throwing and the like.

Arguably, G-art produced by physical forces can be found *inside* McCormack's bubble, too. Given the phrase "at least in part," used earlier to describe a computer's share in Nees's and Nake's work, one might say that Jackson Pollock's paintings exemplified G-art grounded in physics. Although he certainly was not throwing (still less choosing) paint at random, he did not have direct control over the individual splashes as he would have had over marks made with a paintbrush.

Even more control was lost, or rather deliberately sacrificed, when Hans Haacke, in the 1960s, began to exploit and even to highlight the physical behavior of water, vapor, and ice; of waves; and of weather conditions. He wanted to make "something which experiences, reacts to its environment, changes, is nonstable . . ., always looks different, the shape of which cannot be predicted precisely" (Lippard 1997, 38, 64–65). He saw these works not as art *objects* but as "'systems' of interdependent processes" that evolve without the viewer's interaction or "empathy" so that the viewer is a mere "witness." A few years later, Jan Dibbets placed eighty sticks in the sea, a few inches below the surface, and watched them oscillate in the water from fifty feet above: "That," he said, "was the work" (Lippard 1997, 59).

The Surrealists of the 1920s, by contrast, had exploited *psychological* processes—but their work counts as G-art because these were of a relatively impersonal kind. Inspired by Freud, they engaged in automatic writing and painted while in trance states to prioritize the unconscious mind, which Andre Breton declared to be "by far the most important part [of our mental world]." Indeed, Surrealism was defined by Breton as "pure psychic automatism [*sic*] by which one proposes to express . . . the actual functioning of thought, in the absence of any control exerted by reason, exempt from all aesthetic or innocently throwing Metzger's bag of rotting rubbish into the dustbin), the point of the exercise is to remind us of the deterioration that awaits all human constructions—and human beings, too (Metzger 1959, 59; 1965). He was thinking not only of biological decay but also of the terrible destructive power of the Cold War arms race. The artwork is usually assembled by a human artist (or sometimes, Metzger said, by machines in a factory). But it attains its final form, and its significance, through the natural processes of damage and decay.

Such inside-the-bubble (albeit unorthodox) cases, however, are not what the new artists normally have in mind when they refer to generative art. Their phraseology is borrowed from mathematics and computer science, with which the maverick artists just named were not concerned. These disciplines see generative systems as sets of abstract rules that can produce indefinitely many structures or formulae of a given type and that—given the Church-Turing thesis (Boden 2006, 4.i.c)—can *in principle* be implemented in a computer. The G-art community outside the bubble put this principle into practice. That is, their art making rests on processes generated by formal rules carried out by computers—as opposed to physical, biological, or psychological processes or abstractions personally discovered by conscious thought.

In other words, the instances of G-art that most concern us here are those that are also instances of C-art. They are computer-generated art, or CG-art.

A very strict definition of CG-art would insist that *(df.) the artwork results from a computer program being left to run by itself, with zero interference from the human artist*. The artist (or a more computer-literate collaborator) writes the program but does not interact with it during its execution. In effect, the artist can go out for lunch while the program is left to do its own thing.

Such cases do exist. Cohen's AARON program (described later) is one well-known example. Nevertheless, that definition is so strict that it may be highly misleading. Most people working in, or commenting on, generative art allow a compromise in the core concept so as to include interactive art. That is such a prominent subclass of what is called generative art that, even though the taxonomy given here does not aim to capture common usage, it would be highly anomalous to exclude it.

To be sure, the definition of CG-art does cover most interactive art, because it insists on zero interference from the human artist rather than from *any* human being, whether artist or audience. However, it would be very easy for readers to elide that distinction, which in any case makes a questionable assumption about authorial responsibility (discussed in section 2.4). Moreover, the overly strict definition of CG-art excludes those cases (whether inside the bubble or outside it) wherein artists rely on their intuitive judgment to make selections during an artwork's evolution.

It is preferable, therefore, to define CG-art, less tidily, as art wherein *the artwork results from some computer program being left to run by itself, with minimal or zero interference from a human being*. "Minimal," of course, is open to interpretation. It necessitates careful attention to *just what* interference goes on and *by whom* in any particular case.

Most of what people call computer art is CG-art, in this sense. Indeed, the phrases "computer art" and "generative art" are often regarded as synonyms. Notice, however, that in our terminology not all C-art is CG-art. CA-art is not, because the computer is used as a tool subject to the artist's hands-on control and is of no more philosophical interest than a paintbrush or a chisel.

Admittedly, the distinction between CA-art and CG-art is not always so clear-cut as in the CA-art examples mentioned earlier. We see in chapter 6, for instance, that a program for simulating various painting materials and styles can be run either *more* or *less* autonomously (Colton, Valstar, and Pantic 2008). More, and it counts as CG-art; less, and it is better seen as CA-art—although even so, it relies heavily on CG-processes. In general, CA-art may involve AI agents: programs called on by the human artist to aid in specific ways during the production of an artwork (Boden 1994; 2006, 13.iii.d). The more that the ongoing direction is assumed by the computerized agents, not by the human being, the closer the project is to CG-art.

CG-art is intriguing on two counts. First, the generality and potential complexity of computer programs means that the possible space of CG-artworks is huge, indeed infinite. Moreover, most of the structures in that space will be images or music that the unaided human mind could not have generated or even imagined as the artists themselves admit. CG-literature can be ignored here: unless it is heavily interactive, it is much less successful, because the relevant knowledge of language and of the world is too rich to be implemented in computers.

A skeptic might object that much the same is true of a trumpet or a cello: not even the most skilled stage impressionists could mimic these instruments plausibly. In short, human artists often need help from machines. Trumpets, computers—what is the difference? Well, one important difference has just been mentioned: the generality of digital computers. In principle, these machines can (and do) offer us an entire symphony orchestra and an infinite set of visual images and sculptural forms—indeed, an infinite range of virtual worlds. McCormack (2003, 7) goes so far as to compare this infinite space of possibilities, way beyond our comprehension, with the Kantian sublime.

The second point is even more pertinent. Whereas there is no interesting sense in which a trumpet or a cello can be left to do its own thing, a computer certainly can. And it is part of the definition of CG-art that this happens. As we see in section 2.4, this grammar," based on L-systems generating structures in (simulated) 3-D space, comparable to the decentralized yet organized constructions of social insects such as termites (Jacob and von Mammen 2007).

As for evolutionary programming, this has given rise to an important subclass of CG-art: evolutionary art, or Evo-art. Examples include Karl Sims's *Genetic Images* and *Galapagos* (Sims 1991, 2007), plus many others (Whitelaw 2004, chap. 2). In Evo-art, the artwork is not produced by a computer program that has remained unchanged since being written by the artist. Rather, the artwork is *(df.) evolved by processes of random variation and selective reproduction that affect the art-generating program itself.*

Evo-art relies on programs that include self-modifying processes called genetic algorithms. To begin, a population of near-identical artworks—or to be more precise, the miniprograms that generate them—is produced by the computer. There can be any number, nine or sixteen or even more. In aesthetic terms, these first-generation artworks are boring at best and chaotic at worst. Next, each of these first-generation programs is altered (mutated) in one or more ways by the computer, at random. Usually, the alterations are very slight. Now, a selective procedure—the fitness function (decided by the artist-programmer)—is applied to choose the most promising candidate for breeding the next generation. And this process goes on repeatedly, perhaps for hundreds of generations. Provided that the mutations allowed are not too fundamental (see section 2.4), what ensues is a gradual evolutionary progress toward the type of structure favored by the artist.

Occasionally, the fitness function is fully automatic, being applied by the computer itself. If so, there may be scores, or even hundreds, of siblings in a given generation. This is a prime example of the computer's being left to do its own thing. More usually, the selection is made by the artist or a gallery visitor, for instance, using intuitive, and often unverbalized, criteria. In such cases, the population size rarely rises above sixteen because people cannot take in more than a limited number of patterns at once. In other words, and for reasons touched on in section 2.4, there is usually no *programmed* fitness function. In such cases, the Evo-art also counts as interactive art, or I-art.

One might argue that the suggested definition of Evo-art is faulty on the grounds that evolutionary art need not involve a computer. It is certainly true that the very earliest G-art works of the sculptor William Latham, who later became famous as a computer artist, were drawings generated by repeated dice throwing and hand sketching. At that time, he had no idea that computers might be able to do the job for him—and do it better (Todd and Latham 1992, 2–6). But that is highly unusual: virtually all art that is produced by an iterative process of random variation plus selection is computer based. Indeed, there may be no noncomputerized examples besides early Latham. Someone

Whether it should really be called Goldberg's garden is of course debatable (see section 2.4). To be sure, it was his idea in the first place. But as with all Net-art (by definition), its detailed nature at any time has depended on the individual choices of many other human beings. And many, here, really does mean many. The number of participants-artists involved in the *TeleGarden* had already reached nine thousand by the end of its first year online and mushroomed massively thereafter.

Huge numbers apply also to some of the nonrobotic examples. The literary instances—of which Ascott's global fairy tale was a forerunner—are multiauthor hypertexts. They include narratives composed by many hundreds of participants, offshoots of game-playing MUDs and MOOs (Montfort 2003). Their possibility was glimpsed long ago by Vannevar Bush, whose prescient "As We May Think" (1945), originally written as early as 1937, foresaw not only hypertext but search engines such as Google, too (Bolter 1991; Boden 2006, 10.i.h).

In the cases of R-art previously mentioned, only one robot is involved. Sometimes, however, groups of interacting (distributed) robots are constructed. Such groups usually employ the techniques of situated robotics, wherein the machines respond directly to specific environmental cues, including here *the behavior of other robots* (Boden 2006, 13.iii.b and iii.d). Occasionally, they exploit self-organizing techniques whereby the system gradually reaches an equilibrium state. Futuristic though they may seem, both these methodologies were first used by midcentury cyberneticians: Grey Walter and Ross Ashby, respectively (Boden 2006, 4.viii). One example of the latter type is Jane Prophet's *Net Work* installation (Bird, d'Inverno, and Prophet 2007). One might think of this as a high-tech version of Dibbets's oscillating sticks. But instead of eighty isolated sticks placed below the surface of the sea, *Net Work* consists of 2,500 floating, intercommunicating buoys, each of which is color emitting and wave sensitive. More accurately, it *will* consist of 2,500 such buoys: a miniature version has been tested on the Thames, but it is planned to surround the pier at Herne Bay.

Such mutually interacting robot groups *do not* count as interactive art in the definition of I-art given later unless they are also capable of interacting with the human audience. *Net Work* does have that capability: the audience can affect it by shining torchlight on the buoys or by remote control over the Internet. Other examples of interactive (and interacting) robot groups include Kenneth Rinaldo's works in what he calls ecotechnology. His R-art (and I-art) installation called *The Flock* comprises three wire-and-vine robotic arms suspended from the ceiling that interact with each other and with the moving and speaking human audience. Similarly, his *Autopoiesis* has fifteen robot wire-frame arms distributed around the room that sense the observer's movements and communicate with each other so as to coordinate their behavior in various ways.

This brings us to the tenth category: interactive art, or I-art. In this genre, the human audience, which need not include the artist, is not a passive observer but an active participant. Audiences are never wholly passive, of course, because art appreciation involves active psychological processes. Indeed, Duchamp (1957) went so far as to say, "The creative act is not performed by the artist alone; the spectator brings the work in contact with the external world by deciphering and interpreting its inner qualification and thus adds his contribution to the creative act." Even for Duchamp, however, the spectator's contribution concerns only the work's "inner" qualification (its role, he said, is "to determine [its] weight on the aesthetic scale"). The work's perceptible nature—or, many would say, the artwork itself—does not change as a result. In interactive art, by contrast, it does.

In I-art, then, *(df.) the form or content of the artwork is significantly affected by the behavior of the audience*. And in CI-art (i.e., the computer-based varieties), *(df.) the form or content of some CG-artwork is significantly affected by the behavior of the audience*. Again, I am speaking intuitively here: worries about just what counts as the artwork are left to the next section. The word "significantly" is needed, even though it is a hostage to interpretative fortune, so as to exclude performance art, because performance is usually subtly affected by audience reception. As for the word "behavior," this must be interpreted with generosity. In CI-art it covers voluntary actions (such as waving, walking, touching the computer screen, or choosing a plot line within a story), largely automatic yet controllable actions (such as the direction of eye gaze), and involuntary movements (such as breathing). It even includes arcane physical factors such as the radiation of body heat.

Occasionally, the interaction involves not the audience but the physical environment: aspects of the weather, for example, or wave movements. Some installations in city squares respond to the ambient temperature and rainfall, a gentle drizzle causing changes different from those seen in a downpour; others focus on the changing patterns caused by waves on the sea. Strictly speaking, such cases fall outside CI-art as it is defined here, unless—which is usually the case—they *also* involve interaction with the human audience.

CI-art is generative art *by definition*. But it is not generative in our strictest sense as AARON is. Although the artist can go to lunch and leave the program to do its own thing, the audience cannot. However, it qualifies as CG-art in the broader sense because the artist has handed over control of the final form of the artwork to the computer, interacting with some other human being. The degree of control attributable to the audience varies: they may not realize that they are affecting the artwork or (if they do) just what behavior leads to just which changes. We see later that this variability is an important dimension in the aesthetics of CI-art.

I-art is not an entirely recent phenomenon; remember Haydn's dice music, for instance. But it became prominent in the mid-twentieth century. This was often justified in political terms: I-art was seen as offering valuable *human-human* communication, in societies in which the sense of community had been diluted (Bishop 2006). It was made possible largely by cyberneticians such as Pask applying their theory of communicative feedback to art and by the new electronic technology developed in World War II.

That is not to say that all these I-art efforts were examples of Ele-art. Many artists, indeed, eschewed such technology for (countercultural) ideological reasons: it was too strongly linked with the military-industrial complex. Even Ascott's first I-art had nary an electron in sight; it consisted of canvases with items or images on them that could be continually moved around by hand, so that the viewer of the resulting collages was their *maker* too (Mason 2008, 54–58). *SAM* and the *Senster* were early examples of I-art that did use electronics. But as we have seen, they did not involve computers.

Today's I-art, however, is overwhelmingly computer based. That is because the *generality* of digital computers enables them, in principle, to support an infinite variety of human-computer interactions.

The types of interaction explored in CI-art are already widely diverse—hence the inclusiveness of the term "behavior" in the definition. The by-now countless examples range from interactive CD-ROMs viewed on a desktop and altered (for instance) by touching the screen (Leggett and Michael 1996), to room-sized video or virtual reality (VR) installations—such as Christa Sommerer and Laurent Mignonneau's *Trans Plant*. In this case, a jungle gradually appears on the walls as the audience moves around the enclosure: grass grows when the viewer walks, and trees and bushes when the viewer stands still; the plants' size, color, and shape depend on the size and body attitudes of the human being; and the color density changes as the person's body moves slightly backward or forward. *Trans Plant* is driven by the viewer's movements, but some CI-artworks are modified also, or instead, by the sound of human voices or footsteps. This is reminiscent of the *Senster*—but these computer-generated changes are much more varied and complex than those that could be engineered in the 1960s by Ihnatowicz.

Sometimes, the relevant interactions involve online access to the Internet. This is true of Net-art in general, of course, wherein the artwork itself exists only by means of the Internet. But it is also true when the artwork being shown or generated on the walls of a gallery is enhanced by the automatic incorporation of items that happen to

The thirteenth category, LC-art, covers cases of live coding, wherein *(df.) the art object is generated on the fly by code that is written in real time by the artist*.

Most LC-art is music, but occasionally images or texts are generated instead. The reason is that producing an interesting or attractive artwork is only part of the LC-artist's goal. And appreciating the art object for its own sake is only part of what the LC-audience is intended to do. In addition, they are supposed to read the code (which is displayed in real time, while it is being written) and appreciate how it informs the resulting artwork. Why did *that* instruction lead to *that* observable feature? And *what sort* of music is likely to result from *that* new piece of code? It is even more difficult to do this when both code and result are presented visually than when one is visual and the other auditory. So music is the dominant medium in LC-art. Indeed, very often live coders do not place that much emphasis on the exact manner and detail of the visual presentation of the code, hence indicating indirectly that the music is primary.

Even in the case of music, however, only a highly computer-literate audience can achieve the desired experience. It follows that the problems of art appreciation discussed in chapter 4 are especially great for LC-art. Moreover, it takes a special sort of coding to enable the audience to do so. Some LC-artists use existing programming languages, which may already be familiar to the audience. But others design novel programming languages, so that the successively appearing lines of code can be immediately seen to imply particular musical results.

LC-art is a type of CG-art because it involves programming by the artist and the generation of the LC-artwork by the computer. But it is similar to CA-art in one way: the human artist has direct control over what the program is doing at every moment. It therefore raises fewer philosophical questions than other categories of CG-art as CA-art in general does too (see section 2.4).

The final category is 3DP-art, which involves *(df.) the construction of artworks by means of 3-D printing*. Using this new technology, objects are built by laying down successive levels of some hard-setting fluid. This computer-controlled process is being used to produce toys, jewelry, artificial organs (e.g., ears), and even guns. But it is also beginning to be used by artists (as yet, only very few).

In principle, 3DP-art could be *generative* in the sense required for CG-art. For instance, it is suggested in chapter 6 that some future version of a program like *The Painting Fool* might employ 3-D printing to build brushstrokes in the chosen style, having appropriate 3-D texture as well as the requisite colors.

A team of artists and engineers in the Netherlands is already copying a Rembrandt self-portrait and his painting *The Jewish Bride*, as well as Van Gogh's *Flowers in Blue Vase*, all in three dimensions, not just two. However, that project counts as CA-art rather than CG-art: the desired results were initiated by Rembrandt and Van Gogh, and the team is reproducing *only* those. Similarly, many 3DP-artists who aim to produce their own, novel artworks are using the technology in much the same spirit as Hockney uses Photoshop. In theory, they could make the works without the technology, but in practice the facility that it offers makes the construction of the sculptures realistic. See, for example, the exhibition *Beyond the Buzz*, shown at the Minneapolis College of Art and Design in 2015 (http://mcad.edu/events-fellowships/beyond-the-buzz).

Achieving truly generative 3DP-art involves added challenges. These challenges do not involve concerns for computation, as such, but concentrate on relationships between the virtual and the physical and the ways that different physical objects can relate to one another through their shared virtual digital models. One of the few examples of attempts to meet such challenges can be found in a 2010 exhibition displayed in Auckland, New Zealand (*Hybrids* at the Media and Interdisciplinary Arts Centre; http://shura.shu.ac.uk/4391/). The key point about this work was its emphasis on the dialogue between 3-D printing and traditional media. One of the curators, Ian Gwilt, has specifically concentrated on hybrid collections of artworks in which 3-D printing forms one aspect and is both influenced by and influences the artwork (see, e.g., Gwilt 2013).

In sum, the fourteen definitions in this taxonomy are as follows:

1. Ele-art involves electrical engineering or electronic technology.
2. C-art uses computers as part of the art-making process.
3. D-art uses digital electronic technology of some sort.
4. CA-art uses the computer as an aid (in principle, nonessential) in the art-making process.
5. G-artworks are generated, at least in part, by a process not under the artist's direct control.
6. CG-art is produced by leaving a computer program to run by itself, with minimal or zero interference from a human being. The stricter definition of CG-art (art produced by a program left to run by itself, with *zero* interference from the human artist) was deliberately rejected.
7. Evo-art is evolved by processes of random variation and selective reproduction that affect the art-generating program itself.
8. R-art is the construction of robots for artistic purposes, in which robots are physical machines capable of autonomous movement or communication. C-art that uses robots but has an artistic focus not on the robots themselves is excluded.
9. Net-art is the generation of the artwork on the Internet, by multiple human interactions with the computer and indirectly with each other.