

Ten Questions Concerning Generative Computer Art

Jon McCormack, Oliver Bown,
Alan Dorin, Jonathan McCabe,
Gordon Monro and Mitchell Whitelaw

27 July 2012

Jon McCormack (artist, academic, researcher), Centre for Electronic Media Art, Faculty of Information Technology, Monash University, Caulfield East, VIC 3145, Australia, Email: <Jon.McCormack@monash.edu>

Oliver Bown (musician, academic), Design Lab, Faculty of Architecture, Design and Planning, University of Sydney, NSW 2006, Australia, Email: <ollie@icarus.nu>

Alan Dorin (academic), Clayton School of Information Technology, Monash University, Clayton 3800, Australia, Email: <Alan.Dorin@monash.edu>

Jonathan McCabe (generative artist), Faculty of Arts and Design, University of Canberra, ACT 2601, Australia, Email: <Jonathan.McCabe@canberra.edu.au>

Gordon Monro (generative artist), Centre for Electronic Media Art, Faculty of Art, Design and Architecture, Monash University, Caulfield East, VIC 3145, Australia, Email: <gordon@gommog.com>

Mitchell Whitelaw (artist, academic, writer), Faculty of Arts and Design, University of Canberra, ACT 2601, Australia, Email: <Mitchell.Whitelaw@canberra.edu.au>

Abstract

In this paper we pose ten questions we consider the most important for understanding generative computer art. For each question, we briefly discuss the implications and suggest how it might form the basis for further discussion.

Introduction

Generative techniques are increasingly being adopted in many creative practices, from the visual arts and design, through music, cinema and text. What is it about the generative approach that makes it so interesting and well espoused across a diversity of practices? With many artists now embracing

generative techniques, it seems timely to identify the core issues and provide a greater understanding of generative art’s distinctive features. We also need to consider how generative art might develop if it is to remain attractive and relevant to creative practices of the future, alongside the broader impacts of new technology on art and creativity.

In this paper we pose ten questions we think fundamental for understanding generative art. They clarify what makes it interesting and explore the long-term implications of its role as a creative methodology. Rather than providing comprehensive answers, we explore the implications of each question in turn, suggesting how it might form the basis for further discussion and reflection.

Definitions of Generative Art

A number of definitions of generative art have been proposed (see e.g.^{1–3}), which situate it within a wider range of artistic activity, and classify it according to media, methodologies or approaches (systems art, interactive art, algorithmic art, software art, artificial life art, evolutionary art, etc.).

While the questions we pose below are predominantly concerned with generative computer art,⁴ generative procedures have a long history in art that predates the computer by thousands of years. Additionally, much contemporary generative art does not involve digital computers at all.⁵ But the computer and associated technological progress bring new ideas and possibilities that have previously been impossible or impractical to realise. This makes generative computer art different from its non-computational counterparts (an issue explored further in Question 4).

In essence, all generative art focuses on the *process* by which an artwork is made and this is required to have a degree of *autonomy* and *independence* from the artist who defines it. The degree of autonomy and independence assigned to the computer varies significantly – from works that seek to minimise or exclude the creative “signature” of the human designer, to those where the computer’s role is more passive and the human artist has primary creative responsibility and autonomy. This variation is mirrored by different views of art within the generative art community, ranging from a perception that art primarily refers to stand-alone art-objects that are evaluated for their formal aesthetic value, to understanding art as an embedded social and cultural activity within which machines are currently unable to participate independently. In this latter view, relations and artistic meaning emerge through a network of interactions between people and their activities.

In contrast to the critical and social analysis that has traditionally surrounded art movements, generative art is understood primarily as a methodology, with little, if anything, to say about the art itself or the motivations of its practitioners. Despite an increasing number of artists calling their practice “generative”, arguably the only thing *all* generative art shares is

this broad, generic methodology. We explore this issue further in a number of the questions that follow.

The Ten Questions

Question 1: Can a machine originate anything?

That is, can a machine generate something new, meaningful, surprising and of value: a poem, an artwork, a useful idea, a solution to a long-standing problem?⁶ Certainly, computers have played a role in creating all these things and more, but how much of the creativity derives from the program and how much from the programmer?

The mechanistic nature of computing technology leads to the enduring position first attributed to Lady Lovelace (1815–1852): that computers are passive machines that can only do as they are instructed. Many generative artists concur that programming a computer to perform beyond what was obviously encoded in the software’s design is a difficult challenge, but a desirable goal.⁷

There are two common objections to the criticism of a computer program being unable to originate anything it wasn’t expressly programmed to do. The first concerns human ability to know or predict the complete behaviour of any program. Program behaviour, while defined by the program (created by the programmer), typically has a large, sometimes vast, number of executable pathways. This makes it impossible for the programmer to completely understand and predict the outcome of all but the most trivial programs – one reason why software has “bugs”. The second objection arises from the ability of a program to modify itself. Computer programs can be adaptive, they can learn, and so initiate new and potentially creative behaviours.

Computers have already demonstrated the ability to originate something: to exceed their programmer’s anticipations or knowledge. Indeed, this potential for “emergence” is the basis for many an artist’s decision to use the computer. But it is a more difficult problem for a machine to independently originate things of artistic meaning, surprise or value. As computers have developed, we’ve seen our relationship with them change, and the computer’s role shift from that of a “tool” under the direct control of the artist, through to that of a collaborator or creative partner, and potentially, an autonomously creative entity. This suggests a continuum of *creative agency*, assigned in shifting proportions between human and machine, and inversely proportional to the degree of control and intention in the role of the human artist.⁸

Philosophers such as Anthony O’Hear have argued that, no matter how sophisticated or independent, machines cannot originate art because art “in the full sense is based in human experience” and requires a communication

between artist and audience drawn from that shared experience.⁹ Computer works that mimic this communication are only parasitically meaningful as they derive their meaning from an analysis of existing art-objects, not directly from human experience. However, in response, we can see no reason to dismiss outright the possibility of a machine and a human sharing experiences that result in something meaningful and worth communicating.

We should also remember that the creative splendour of human cultures and built environments are collective and cumulative efforts. Individual creativity is arguably weak in the absence of the structures and systems that enable the accumulation of artefacts and information, so competent autonomous computer artists might conceivably require a similar context.

Question 2: What is it like to be a computer that makes art?

If a computer could originate art, what would it be like from the computer's perspective to be an artist? If this perspective was very different to our own, how would we recognise it or comprehend its art? What kind of cognitive or subjective experiences does a computer need before we can consider it an artist? If art is a social exchange, to what kinds of social contexts could computers belong?

The goal of programming a machine to be an autonomous artist seems to impose a double standard: we're asking the machine to be autonomous, yet we're also asking for human creativity, assessed by human standards. If we abandon this second constraint, then we introduce the problem of recognition – what could possibly be the defining characteristics of an autonomous computer artist?

In 1974, philosopher Thomas Nagel asked the question “What is it like to be a bat?”, suggesting that conscious mental states require something that it is like to *be* that organism, something that we cannot directly *know* from our experience.¹⁰ In other words, how do we connect the subjective to the objective, particularly if we want our autonomously creative machines to do the same?

We could broaden the question to ask if conscious experience is necessary for a machine to be an artist. Here two different views of art come into play. If the art object is simply an aesthetically appealing form, then consciousness seems unnecessary. Numerous natural and human-designed systems are capable of creating patterns we find interesting and aesthetically pleasing, without reliance on underlying mental states in their generative mechanism. But is it meaningful to call such systems “artists”?

On the other hand if art requires a social or cultural context in which to operate, it probably also requires conscious intent on the part of the artist. While currently, computers and robots don't actively participate in culture as artists, perhaps one day they may. What are the minimal conditions for this to happen? If we can never know what it's like to be a computer that

makes its own art, then how could any such participation ever be appreciated or understood?

Question 3: Can human aesthetics be formalised?

There are few questions that invoke such polarised responses between artists and scientists. Attempts to formalise aesthetics – using quantitative measures or procedural (algorithmic) techniques – are almost as old and varied as the concept of aesthetics itself. Many artists would argue it is the wrong question to ask. However, similarly to Question 1, unless we think there is something uncomputable going on in the human brain, the answer in principle is “yes”.

Considerable technical research goes into trying to answer this question, at least since Birkhoff.¹¹ If an aesthetic measure or algorithm could be devised, then it could be used to automate the generation of aesthetic artefacts (using evolutionary techniques, for instance). If the formalisation included knowledge of individual tastes and preferences, the artefacts could be tailored differently to each individual.

However, most of the current research into formalising aesthetics sees aesthetics in (pre-)Kantian concepts of beauty and pleasure. Generative art too has often made pleasing surface aesthetics a principal fetish. Considering aesthetics as a single scalar quantity doesn’t fit with a contemporary understanding of the term, which has advanced significantly since Kant and Birkhoff.¹² Additionally, aesthetics often shifts according to taste, time and culture so quantifying it at any single point is problematic. Rather than asking if aesthetics in total can be formalised, we could ask, “What *kinds* of aesthetics could be formalised?”. Some possibilities include neurological^{13,14} and evolutionary understandings^{15,16}, which have hypothesised basic mechanisms, principles and explanations of beauty, for example.

There is also a difference between aesthetic judgement and aesthetic evaluation. Human artists plan and evaluate their artwork as it proceeds, they don’t necessarily wait until the final work is finished (as the audience must) before considering its aesthetics. How different are these processes from each other and could either be formalised?

Implicitly, any generative artwork “encodes” human aesthetic judgements within its choice of rules and realisation. But even for systems capable of voluminous output (e.g. image evolving systems), the aesthetic variation is far more limited, indicating that aesthetic responsibility in current generative art resides primarily with the artist rather than the system that generates the work.

Question 4: What new kinds of art does the computer enable?

Computation is a relatively new medium for creative expression, and computers are often appropriated for digital art simply as display devices, or for automating prior processes or paradigms.¹⁷ Many widely respected generative artworks, past and present, do not involve digital computers. So what – beyond generating more art – does generative *computer* art bring that is new to art?

Computers allow us to create and manipulate sophisticated processes with an increasing fidelity, flexibility and a level of control that was not possible previously. Computer simulations allow building of “model worlds”, that permit the vivid realisation and expression of ideas and complex scenarios that are impossible in other media, or in reality (one of the reasons cinema has enthusiastically embraced generative computer techniques is that the representational power of the computer exceeds what can be achieved unaided). Dynamic interaction with complex systems simulated in computers has lead to many breakthroughs in human knowledge. Furthermore, networked computers, now thoroughly embedded within human society, have facilitated and determined unexpected cultural, political and social change. Art itself has not been exempt from these changes.

Elsewhere some of us have argued that generative computer art introduces the concept of a computational sublime¹ and that some emergent properties seen in generative systems have only previously appeared in natural systems, if at all. The computer also appears as a destabilising force in contemporary art practice, challenging concepts of authorship and ownership of the art object. Art traditionally requires a mysterious process of creation, unique to the artist, their skill, and their special way of seeing the world. Generative art has explicit mechanisms; if the process is entirely known it can be considered “mechanical” and exactly repeated across boundaries of space, time and culture. If art can be made mechanically, what is so special about artists?¹⁸

Question 5: In what sense is generative art representational, and what is it representing?

Unless software design is conceptualised directly at the level of individual bits, it is impossible to write a computer program without recourse to some form of representation. The nature of programming enforces this constraint.

Generative computer art often draws on ideas and algorithms from the simulation sciences. A simulation involves the representation of important characteristics and dynamical behaviours of some target system. However, few generative artists would view or conceptualise their works as direct simulations of reality. So if not reality, what is generative art representing?

In traditional visual media (painting for instance), works range over a spectrum from photorealism to pure non-objective mark-making, so a variety of engagements with representation in generative art is similarly expected. But a generative artwork has two aspects, the process underlying the artwork and the sensory artefacts it produces. In some “physical” generative artworks the distinction may be ambiguous, but for computer-based works it is very clear. The two aspects may engage with representation differently, and to different extents. The idea of a computer process representing another process in the world is largely new to art. Computer works require a selective mapping to take place between the internal process and the perceptual artefacts or stimuli through which the process is experienced.

This computational process can be driven by external data (e.g. human interaction, weather conditions), or an abstract process (e.g. a point moving in a circle). It may be a “model world”: a inter-connected system with representational relationships to the world (real or imagined), that follow a “system story”.¹⁹ Yet even the point moving in a circle is not as straightforward as it might seem: there is really no moving point in the computer, merely a changing pattern of bits that represents one.

Is there a continuum or a hard distinction between generative art and data visualisation? If generative art uses real-world data, what are the ethical and political implications of the artist’s chosen representations?

Yet another kind of representation may be called representation *in potentia*. Some generative systems have enormous (much greater than astronomical) numbers of potential variations or exemplars. Does the system represent this enormous range in some sense? (C.f. the computational sublime in Question 4).

How can an audience best understand these selective and often obscure processes of representation? As generative art matures, will we encounter a shift from the mimetic to non-mimetic features of process, similar to the advent of modernism in painting? There are many issues surrounding representation in generative art that deserve greater consideration.

Question 6: What is the role of randomness in generative art?

Not all generative art makes use of randomness, but from the musical dice game of Philip Kirnberger (1757) onwards,²⁰ randomness and chance events have played an important role. The American composer John Cage, well known for his use of chance methods turned to a computer program to generate *I Ching* hexagrams.²¹

We can distinguish different sources of randomness in generative art. The first is “pure” randomness, obtained by a physical process such as rolling dice, tossing coins, or by dividing piles of yarrow sticks, as used in generating hexagrams for the *I Ching*. With the use of computers, pure randomness

has been largely replaced by pseudo-randomness, where the numbers are obtained by a deterministic function, but pass statistical tests for randomness. To introduce variation, typically the process begins with a small injection of pure randomness (a seed), such as the exact second the program was started.

What does the use of randomness say about the place of intentionality in the making of art? John Cage wanted to take the artist’s ego out of the production of the work, but in Iannis Xenakis’s compositions “randomness is introduced as a necessary part of a willed product”.²²

How does knowledge of the source of randomness impact on the conception and interpretation of a work? For example, the concept of wind as an element in an artwork, such as in Tim Knowles’ *Tree Drawings*, is very different from that of stock market fluctuations (seen in a work such as Lise Autogena and Joshua Portway’s *Black Shoals*), even though they both have random properties. Are the metaphorical associations of randomness more important in generative art than its sources?

Randomness is often used to “humanise” or introduce variation and imperfections to an underlying rigid deterministic process, as when a sequencer program plays back a musical score with slight random timing variations. What issues surround the use of randomness as a proxy for poorly understood complexity? (See also Question 7.) In science, statistical modelling is a powerful method, but randomness is sometimes used as a way of working round ignorance or incomplete knowledge: if we truly knew all the forces at play when we tossed a coin, we would know the result. If a generative artist has complete knowledge of the art-making process, why resort to randomness?

Question 7: What can computational generative art tell us about creativity?

Creativity is highly sought after. Brains do it, societies do it, and evolution does it, but how do these things give rise to artefacts and ideas that are new, surprising and valuable?

Creativity is sometimes categorised into two fundamental types: *combinatorial creativity*, in which fixed primitive elements are combined to create new structures, and *emergent creativity* where new structures or symbol primitives emerge *ex nihilo*.²³ Generative systems tend to work with the first type, as getting symbols with new semantics to emerge within a running program is currently an open research problem.²³ Defining the fundamental primitives of the system and their inter-relationships is a creative act in itself. In a truly emergent system new primitives emerge that were not explicitly defined when the system was specified – invoking a creativity attributable to the system itself. In a *dynamic hierarchy* multiple new levels form repeatedly as emergent primitives from one level combine to create primitives and a new level that supervenes upon the old one.

Experienced computational artists know of many “cheap tricks”: simple mechanisms that produce effective results but are seemingly too unsophisticated to be comparable with human creative mechanisms. Many of these tricks involve astute use of random variation. Are we prejudiced in dismissing these tricks as poorly creative, particularly if they achieve desirable results? This question is indicative of how we consider and evaluate generative art (Question 8). Which has greater significance, the process or what it produces?

Of course, generative art is not all cheap tricks. The processes underlying works can also be very insightful, leading to highly original, creative work. We have already raised the issue of how natural, physical and chemical processes form the basis for many generative artworks. Does the creativity reside primarily in the original phenomenon, the algorithm simulating it, or in the artist’s interpretation of it? Being explicit about where this creative agency is attributed (and in what proportions) can make the answers to these questions clear, and we would anticipate, bring a greater understanding of creativity through generative art.

Question 8: What characterises good generative art?

There are more and more examples of “generative art” emerging on-line and in galleries and museums. Thousands of new artists and designers are learning and embracing creative computing environments in universities and art colleges, but there is a risk that this may be taught from an overly generic technical perspective. How can we form a more critical understanding of generative art and equip students with both the conceptual and technical understanding necessary to create challenging and innovative generative works?

Why is generative art in need of special quality criteria? Is it better considered alongside other current practices? Consider two important properties that differentiate generative art from other practices. The first is that the primary artistic intent in generative art is expressed in the generative process. This process is what the artist creates, and as such should arguably be the subject of scrutiny in appreciation of what it produces. Secondly, the way this process is interpreted or realised is also the locus of artistic intent, and is intimately intertwined with the first property.

The basis of all generative art resides in its engagement with process. So the locus of artistic intent should include the motivations, design and realisation of the process and these considerations must be an integral part of any critical analysis or pedagogical imperatives. Put simply, the “generative” and “art” parts are inseparable. Process in generative art should be considered the primary medium of creative expression, implying that the exclusive or predominant use of creative software or processes designed by others in one’s generative practice is problematic. This is currently a

TEN QUESTIONS CONCERNING GENERATIVE COMPUTER ART

contentious issue, with debates regarding “algorithmic genericism” of many generative works on the rise.

Somewhat ironically, many of generative art’s most successful processes originate from natural processes or their scientific abstractions. How can these works escape the criticism of naive “algorithmic genericism” or direct re-appropriation? It should be noted that a substantial amount of human effort has gone into understanding and modelling these processes. It would be unrealistic to require individual artists to devise equivalent processes independently. Additionally, a “systems worldview” sees a limited set of fundamental or canonical processes as responsible for an increasingly wide variety of phenomena. The same canonical process can shift from sublimity to parody depending on artistic interpretation, which requires a degree of technical comprehension. Understanding an algorithm’s subtlety or originality opens a fuller appreciation of the eloquence of a generative work. But this is a significant problem for most audiences, reinforced by focussing on the surface aesthetics of the art object as is often seen in computational generative art, where the computational process is rarely directly perceptible.

Our purpose with this question is not to narrowly or subjectively define “quality” in generative art, rather to provoke a better understanding of its defining aspects. These aspects have, to date, received little critical attention. Greater attention will assist artists and audiences to appreciate what makes generative art interesting art.

Question 9: What can we learn about art from generative art?



Figure 1: Driessens & Verstappen: *The Factory* (detail), 1995, installation stills (courtesy of the artists)

Our previous questions have raised issues regarding agency, originality, creativity, authorship and intent in generative art. Clearly these concepts also impact on how we understand art and the art world in general. For

example, can the art world be considered a complex generative system involving many processes outside the direct control of artists, who are agents of production within a stratified global art market? Generative art redistributes traditional notions of authorship and intention, introducing autonomous processes and agents, allowing us to appreciate the systemic aspects of contemporary art production, exhibition and consumption from an illuminating perspective. Such issues are parodied in Driessens & Verstappen’s *The Factory* (Fig. 1) – a generative work where wax sculptural objects are endlessly made, documented, then destroyed on an automated production line. Each sculpture is different, special, and transient, but fabricated autonomously without human intervention.

Much innovative generative practice occurs outside the “precious bubble” of the high art world, in areas of design and contemporary technological culture (e.g. games, cinema, digital music). Many artists work across multiple disciplines and contexts (scientific, artistic, social, technical), making it limiting to consider generative art exclusively from a fine art perspective. The art world has shown a patchy interest in generative art, probably for precisely this reason, which relates to wider tensions between art, technology and mass culture. Additionally, generative art’s emphasis on algorithmic techne and explicable mechanisms alienate it from the mainstream art world, which often remains tied to the “irreducibility of the work of art”.²⁴ What then, is generative art’s place and role within contemporary culture? Is it confined to academic research or just a commercial-art tool currently in vogue, but possibly nearing exhaustion? Is its role simply as a generator of new techniques for application in the design and cultural industries? Is it closer to craft practice than art practice?

Question 10: What future developments would force us to rethink our answers?

The staggering changes brought about by developments in computing technology present us with many opportunities to rethink our relationships to the social and physical world. The complex, emergent nature of these relationships makes prediction of their long term impact difficult. But change will undoubtedly occur, and these changes will force us to rethink our questions.

Take, for example, recommender systems – the globalised Internet interpretation of “word of mouth” recommendations for literature, art, music, cinema, etc. Taken to an extreme these could change our approach to Question 3, by actually altering the way our aesthetic preferences are realised and satisfied. Rather than recommending similar or related works, future recommender systems could generate content *precisely* to individual taste, eliminating any need for choice on the part of the consumer.

If the future is home to autonomous machine agents that are capable of

participating in a society as artists, we cannot at present know what form they will take because that depends on the emerging state of technology and of art itself, and how this in turn affects human behaviour. The Western notion of art has been radically transformed from the classicism of antiquity to the relational aesthetics of the present day. Throughout this time, the meaning of terms such as “create”, “generate” and “originate” has been far from stable.²⁵ If the language and computational concepts used in this paper change as dramatically as the concepts and motivations of artists and their audiences over this period, then we should expect our discussions to become a caricature of their time rather than a far reaching analysis of the possibilities of generative art.

Conversely, history shows optimistic speculation about technology is often misfounded. Many automated “creative-decision” systems found in current technology limit human creative choice rather than enhance it. Attempts to achieve open-ended evolution and generative complexity in software has so far proven unsuccessful. Despite 60 years of research, artificial intelligence on par with humans remains only a distant possibility. We should be prepared to concede the diminishing returns and limits of technological progress. Our humanist tendencies to satisfy *our* desires above all others limit what is acceptable and possible in our machines.

Emphasising the processual nature of generative art invokes another problem: has generative art run out of ideas? Question 8 raised the possibility that direct mimicry of processes discovered in other disciplines, or by other artists, lacks innovation. But radically new processes might become increasingly difficult to discover, potentially leading to conceptual stagnation.

Conclusion

Our purpose in devising these questions is not to be prescriptive about what is important in generative art. Rather, it is to motivate a productive critical discussion on what makes generative art an interesting practice, how we differentiate it from other practices, and what the implications of possible technological developments hold for it as a creative methodology.

Generative art’s discipline- and medium-independent methodological focus has given it reach beyond traditional fine art boundaries. This has both advantages and difficulties: advantage in that it is not constrained by individual disciplinary concerns, discourses or trends, but difficult in that it lacks the philosophical, artistic and critical pedigree of a traditional art movement. We hope that our questions are helpful in opening further discourse on generative art and its role in creative cultures past, present and future.

Acknowledgements

This research was supported by an Australian Research Council Discovery Projects grant DP1094064.

References and Notes

- [1] [Jon McCormack and Alan Dorin. Art, emergence and the computational sublime. In Alan Dorin, editor, *Second Iteration: conference on Generative Systems in the Electronic Arts*, pages 67–81. CEMA, Melbourne, Australia, 2001.](#)
- [2] [Philip Galanter. What is generative art? complexity theory as a context for art theory. In *6th International Conference, Exhibition and Performances on Generative Art and Design \(GA 2003\), Milan.*, Available online at \[http://www.philipgalanter.com/downloads/ga2003_paper.pdf\]\(http://www.philipgalanter.com/downloads/ga2003_paper.pdf\), 2003.](#)
- [3] Margaret A. Boden and Ernest A. Edmonds. What is generative art? *Digital Creativity*, 20(1 & 2):21–46, March 2009.
- [4] That is, generative art that uses a digital computer as the primary (but not necessarily solitary) mechanism for enacting the generative process specified by the artist.
- [5] [Alan Dorin, Jon McCormack, Jonathan McCabe, Gordon Monro, and Mitchell Whitelaw. A framework for understanding generative art. *Digital Creativity*, to appear, \(accepted Nov. 2011\), 2012.](#)
- [6] This is a well discussed problem at least since: Alan M. Turing. Computing machinery and intelligence. *Mind*, 59:433–460, 1950.
- [7] [Mitchell Whitelaw. *Metacreation: art and artificial life*. MIT Press, Cambridge, Mass., 2004.](#)
- [8] [Oliver Bown and Jon McCormack. Creative agency: A clearer goal for artificial life in the arts. In George Kampis, István Karsai, and Eörs Szathmáry, editors, *ECAL \(European Conference on Artificial Life\)*, volume 5778 of *Lecture Notes in Computer Science*, pages 254–261. Springer, 2009.](#)
- [9] [Anthony O’Hear. Art and technology: An old tension. *Royal Institute of Philosophy Supplement*, 38:143–158, 1995.](#)
- [10] [T. Nagel. What is it like to be a bat? *The philosophical review*, LXXXIII\(4\):435–450, October 1974.](#)

- [11] For an overview of attempts at formalised aesthetic measures, see: Philip Galanter. Computational aesthetic evaluation: Past and future. In Jon McCormack and Mark d’Inverno, editors, *Computers and Creativity*, pages 263–302. Springer, 2012.
- [12] Leonard Koren. *Which “Aesthetics” Do You Mean?: Ten Definitions*. Imperfect Publishing, 2010.
- [13] V. S. Ramachandran and William Hirstein. [The science of art: A neurological theory of aesthetic experience. *Journal of Consciousness Studies*, 6:15–51, 1999.](#)
- [14] John Onians. *Neuroarthistory*. Yale University Press, New Haven, 2007.
- [15] Colin Martindale, Paul Locher, and Vladimir M. Petrov, editors. *Evolutionary and Neurocognitive Approaches to Aesthetics, Creativity and the Arts*. Foundations and Frontiers in Aesthetics. Baywood Publishing Co., Inc., 2007.
- [16] Denis Dutton. *The Art Instinct: beauty, pleasure, and human evolution*. Oxford University Press, 2009.
- [17] J. David Bolter and Richard Grusin. [Remediation: understanding new media](#). MIT Press, Cambridge, Mass., 1999.
- [18] Similar questions were raised by photography and cinema.
- [19] [Mitchell Whitelaw. System stories and model worlds: A critical approach to generative art. In *Readme 100: Temporary Software Art Factory*, pages 135–154. Books on Demand GMBH, Norderstedt, 2005.](#)
- [20] Stephen A. Hedges. Dice music in the eighteenth century. *Music and Letters*, 58(2):180–187, 1978.
- [21] William Brooks. John cage and history: Hymns and variations. *Perspectives of New Music*, 31(2):74–103, Summer 1993.
- [22] Paul Griffiths. *Aleatory*, Groves Music Online, 2012. URL [`http://www.oxfordmusiconline.com/`](http://www.oxfordmusiconline.com/).
- [23] Peter Cariani. [Creating new informational primitives in minds and machines. In Jon McCormack and Mark d’Inverno, editors, *Computers and Creativity*, pages 393–428. Springer, 2012.](#)
- [24] P. Bourdieu. [The rules of art: Genesis and structure of the literary field](#). Stanford University Press, 1996.
- [25] R. Williams. [Keywords: a vocabulary of culture and society](#). Oxford University Press, Oxford, 1976.

Author Biographies

Jon McCormack is an artist and ARC Australian Research Fellow at the Centre for Electronic Media Art at Monash University. Since the late 1980s he has worked extensively with computer code as a medium for creative expression and explored creative methodologies using computation.

Oliver Bown is an electronic musician, programmer and researcher in computing, evolutionary and adaptive systems, and music. He is half of the duo *Icarus*, and a member of the improvisation collective *Not Applicable*.

Alan Dorin is an academic at Monash University where he is an Artificial Life researcher and generative artist working in electronic media. His interests include ecosystem simulation and agent-based modelling, artificial chemistry, self-assembling systems, the evolution of complexity, the history and philosophy of science and art, and the links that bind all fields together.

Gordon Monro is a generative artist whose work ranges across sound and music, digital prints, abstract videos and computer-based installations; his work is often inspired by scientific or mathematical concepts. Gordon is currently undertaking a PhD in the Faculty of Art and Design at Monash University in Melbourne.

Jonathan McCabe is currently investigating the application of multiple concurrent Turing pattern formation processes to generative art and design. He is collaborating on new works with the designers 'Nervous System' and the audio/visual collective 'Plink Flojd'. He is affiliated with the Le Méridien group of hotels as a member of the 'LM100'. He is enrolled in the PhD program of the Faculty of Arts and Design at the University of Canberra.

Mitchell Whitelaw is an academic, writer and artist with interests in new media art and culture, especially generative systems and data-aesthetics. He is currently an Associate Professor in the Faculty of Arts and Design at the University of Canberra, where he leads the Master of Digital Design degree.