Ten Questions Concerning Generative Computer Art

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

12 April 2012

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. But 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 is timely to identify the core issues, and provide a greater philosophical, conceptual, critical and practical understanding of generative art. We also need to consider how generative art might progress if it is to remain attractive and relevant to creative practices of the future, as it develops alongside the broader impacts of 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

> "The idea becomes a machine that makes the art."
>
> — Sol LeWitt, 1967.

Extended version of: J. McCormack, O. Bown, A. Dorin, J. McCabe, G. Monro and M. Whitelaw, Ten Questions Concerning Generative Computer Art, *Leonardo* (to appear, accepted July 2012), MIT Press, 2012

A number of definitions of generative art have been proposed (see e.g. [3, 12, 19]) which attempt to situate it within a wider range of artistic activity, and to classify it according to media, methodologies or approaches (e.g. systems art, interactive art, algorithmic art, software art, artificial life art, evolutionary art).

While the questions we pose below are predominantly concerned with generative computer art,[1] generative procedures have a long history in art that predates the computer by thousands of years [10]. Additionally, plenty of interesting contemporary generative art does not involve digital computers at all (see examples in [10]). To fully understand computational generative art we must take into account this rich and insightful non-computational oeuvre. But the computer and its 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. Issues of autonomy and independence are raised in a number of our questions. Indeed, arising from our discussions is a distinction between what might be termed "strong" and "weak" generative art. Strong generative art gives creative autonomy and independence primarily to the computer, minimising the creative signature of the human who designed the system. In contrast, weak generative art uses the computer more passively as a tool or assistant, the human artist having primary creative responsibility and autonomy. This distinction is complemented by different views of art within the generative art community. These views range 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.

Beyond a Methodological Description

In contrast to the critical, conceptual and social analysis that has traditionally surrounded art movements, definitions of generative art are primarily methodological classifications which have little, if anything, to say about the art itself or the motivations of those who make it. Despite an increasing number of artists calling their practice "generative" and the collective social association of generative artists in, for example, on-line communities, arguably the only thing *all* generative art shares is a broad, generic methodology.

[1] i.e. generative art that uses a digital computer as the primary (but not necessarily solitary) mechanism for enacting the generative process specified by the artist.

We explore this issue further in a number of the questions that follow.

A descriptive statement is insufficient to define or understand generative art, so the purpose of this paper is to frame generative art in terms of the questions it raises, with a view to assisting practitioners, theorists and critics to gain a better understanding of its significance.

The Ten Questions

Question 1: Can a machine originate anything?

> There is at present no question of arguing from the way that actual machines behave to the nature of art, for there are no machines which can create art. This must not be misunderstood. We have Dr. George's word that it is possible to "programme" a computer to produce a painting or a piece of poetry; but, as he points out, the programme comes from outside the machine, which merely carries out the instructions of the programmer. One would not, then, expect to learn anything about art from the machine's part in the proceedings, any more than from, say, the mechanism of the novelist's typewriter.
>
> — G.H.R. Parkinson, 1961 [23]

Can a machine originate anything? That is, can it generate something new, meaningful, surprising and of value: a poem, an artwork, a useful idea, a solution to a long-standing problem? Certainly, computers have played a role in creating all these things and more, but how much of the creativity comes 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: that computers cannot originate anything. They are machines that can only do what they are programmed to do; the real creativity is in the programmer, not the program.[2]

A stated in the introduction, generative art involves procedures (rules, formal instructions, algorithms) created by the artist, that are enacted with varying degrees of autonomy. Many generative artists concur that programming a computer to perform beyond what was obviously encoded in the software's design is a vexing challenge, but a desirable goal [29].

There are two common objections to Parkinson's "mere instructions" thesis quoted above. The first concerns the human ability to know or predict the complete behaviour of any program. Program behaviour, while defined by the program (and created by the programmer), typically has a large, sometimes vast, number of executable pathways. This makes it impossible

[2] As Parkinson's quote suggests, this is a well discussed problem at least since Turing [26], so we try to keep the discussion focused on the relevance to generative art.

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 and change itself. Computer programs, like people, 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 (and thus, to disprove Descartes' dictum [1]). Indeed, this potential for "emergence" is often the basis for many a generative artist's choice in using the decision to use a computer.

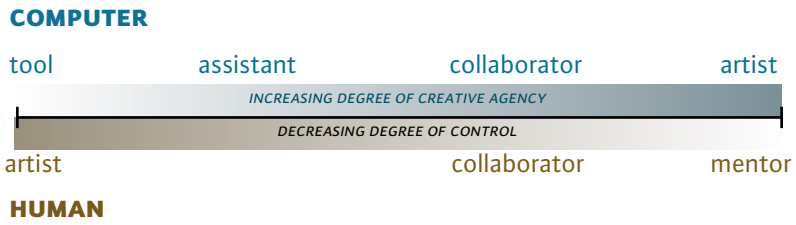

Figure 1: Levels of creative agency between a computer and human

But can a machine independently originate things of artistic meaning, surprise or value? This is a more difficult problem, but equally so for people. As our use of computers has 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 (Fig. 1). This suggests a continuum of creative agency, assigned in shifting proportions between human and machine [6], and inversely proportional to the degree of control and intention in the role of the human artist.

Some philosophers 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 [21]. 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.

Finally, the question has a more tantalising dimension: can a machine originate anything in a radically different sense to that observed in human behaviour? Supporters of the Singularity say so [27], suggesting the possibility of a future in which machine creativity will radically challenge our imagination, and our current concepts of artist, art and creativity.

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

> It would be the realisation that all the "computing power" that resides in the human brain . . . can be bypassed with a handful of state-of-the art chips, and that all that is needed to produce the most powerful artistic outbursts of all time (and many more of equal power, if not greater) is a nanoscopic fraction thereof—and that it can all be accomplished, thank you very much, by an entity that knows nothing of knowing, seeing, hearing, tasting, living, dying, struggling, suffering, aging, yearning, singing, dancing, fighting, kissing, hoping, fearing, winning, losing, crying, laughing, loving, longing, or caring.
>
> – Douglas Hofstadter [15, p. 80]

If a computer can 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 what it does as art? What kind of cognitive or subjective experiences does a computer need to have before we consider it as an artist? If art is a social exchange, what kinds of social context could computers belong to?

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, we then have 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?" [20], suggesting that conscious mental states require something that it is like to *be* that organism, something that we cannot directly *know* from our experience. 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. Again 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 (reaction-diffusion systems, for example) 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 something like a social or cultural context in which to operate, then it probably also requires conscious intent on the part of the artist. This definition sets a high bar for any artificial artist, but while 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?

> Evaluation of a work in progress is directed to how to proceed. Evaluation of a finished work is directed to whether it's any good. The procedures required to satisfy the two are likely to be quite different, even when the same aesthetic is informing the procedure in each case.
>
> — Harold Cohen

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. The question of formalising aesthetic evaluation relates to central questions in the foundations of Artificial Intelligence (formalising human intelligence), to which there is a long, well-known and respected set of objections, and an equally long and respected list of affirmations. Many artists would argue it is the wrong question to ask. However, similarly to question one, 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, and there have certainly been noble attempts, from [2] to [17]. If an aesthetic measure or algorithm can be devised, then it could be used to automate the generation of aesthetic artefacts (using evolutionary techniques or machine learning, for instance). If the formalisation included knowledge of individual tastes and preferences, the artefacts could be tailored differently to every individual, uniting modern mass-production with aesthetic *haute couture* on an unprecedented scale.

However, most of the current research into formalising aesthetics sees aesthetics only 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 (see e.g. [16]). 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, a step forward is to ask, "What kinds of aesthetics could be formalised?". Some possibilities here include neurological [22, 25] and evolutionary understandings [11, 18], which have hypothesised basic mechanisms, principles and explanations of beauty, for example. But these are basic generalisable aspects and principles only, currently only stepping stones towards a more complete understanding.

There is also a difference between aesthetic judgement and aesthetic evaluation. As the quote from Cohen suggests, 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. 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 automated output (e.g. image evolving systems), the aesthetic variation is far more limited, suggesting that (and in contrast with the goals introduced in Questions 1 and 2) 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, but computers appropriated for digital art may simply use them as display devices, or for automating prior processes or paradigms [4]. Generative art predates the digital computer and many widely respected generative artworks do not involve digital computers. So what – beyond generating more art – does generative computer art bring that is new to art? Does the computer change or enable artistic possibilities beyond mimicry, automation and remediation?

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 and understanding. 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 sublime [19] 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 or Promethean process of creation, unique to the artist, their skill, and their special way of seeing the world, qualities that are individual and distinct. 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?[3] Generative mechanisms allow for individual adaptation of the art object, conveniently meshing with the mass-produced, global delivery platforms of the technological age, allowing for difference within indifference: quite the antithesis of conventional art.

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. If generative art is representational, what is it representing?

In traditional visual media (painting for instance), works range over a spectrum from photorealism to pure non-objective mark-making, so a sweep 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 one process (in the computer) representing another (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 a simple abstract process (e.g. a point moving in a circle), or it may be a "model world": a more complex system

[3] Similar questions were raised by photography and cinema.

with representational relationships to the world (real or imagined), that follow a "system story" [28]. 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.

The example of external data as a source of generation includes scientific visualisation, where data is (ideally) represented in an informative and unbiased way. 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 Q4).

How can an audience best understand these selective and often obscure processes of representation? As generative art matures, will we see 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?

> "Practice makes perfect, imperfect is better"
>
> — Paul Bley [`http://www.improvart.com/bley/improv.htm`]

Not all generative art makes use of randomness, but from the musical dice game of Philip Kirnberger (1757) onwards [14], 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 [8].

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" [13].

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 is very different from that of stock market fluctuations, even though they both have random properties. Are the metaphorical associations of randomness more important in generative art than its source?

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 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 and bodies 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* [9]. 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 [9]. Combinatorial creativity is often thought of as limiting. If, for example in a musical work, primitives are the seven notes of a diatonic scale, those primitives can never generate compositions that use all twelve notes of a chromatic scale, no matter how sophisticated the generative system. 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 many cheap tricks (or "fast and frugal heuristics"): 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 only cheap tricks. The processes underlying works can also be insightful, leading to highly original, creative work. We have already raised the issue of how natural, physical and chemical processes have formed the basis for many a generative artist's creative inspiration. Does that imply that the creativity resides 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 [7].

Question 8: What characterises good generative art?

> "Technology in alliance with bad taste is the enemy of art most to be feared"
>
> — Goethe

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 new generative works?

Why is generative art in need of special quality criteria? Is it better considered alongside other current practices? Consider two important principles that differentiate generative art from other practices. The first is that the primary artistic intent in generative art is expressed in the design, specification and construction of 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 principle.

The basis of good generative art resides in its engagement with process. That is, this locus of artistic intent should include the motivations, design and realisation of the process and these aspects must be an integral part of any critical analysis or pedagogical imperatives. The "generative" and "art" parts are inseparable. Therefore, process in generative art must 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 contentious issue, with debates regarding "algorithmic genericism" of many generative works on the rise.

Somewhat ironically, many of generative art's most popular 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 generative 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 some degree of technical understanding. Understanding an algorithm's subtlety or originality opens a full appreciation of the eloquence of a generative work. This is a significant problem, 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.

Some might object to our calling for ways to make value judgments about artistic quality. Our purpose here is to provoke a better understanding of generative art and its unique aspects. These aspects have, to date, received little critical attention. We believe that greater attention can help artists and audiences understand what makes generative art interesting art.

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

> At the time we were under the impression that a work of art seemed to be primarily a strategic instrument guaranteeing the continuity of the institutionalised art establishments. New artworks had to be shown every month, and production had to be kept up. The journals gave the most extensive and favourable reports to those galleries and art institutions buying big expensive pages of advertisements in their magazine. The so-called new and interesting seemed to be strongly intertwined with mutual commercial interest. We concluded that the art world was a generative system maintaining itself.
>
> — Erwin Driessens and Maria Verstappen

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?

Artists Driessens & Verstappen suggest it can, involving many processes outside the direct control of the artist, 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 exquisitely parodied in Driessens & Verstappen's *The Factory* – 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 only a patchy interest in generative art, probably for precisely this reason, which relates to wider tensions between art and technology. 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" [5]. 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 relationship to the world. In recent years, digital technology has bound itself to our social organisation and is slowly pervading everyday objects to create an internet of things. These in turn present new niches for technological evolution, social and creative change. 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 (understanding you better than you understand yourself)

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. Consider how radically the Western notion of art has transformed through its evolution 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 [30]. 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.

At present the legal and business worlds are bracing themselves for the impact that techniques of automated invention will have on patents, copyright and competition [24]. As with the historical evolution of the law, we can expect the concepts and terminology that emerge to seep into colloquial language and culture. Our current concepts of art are deeply entwined with the notional semantics of originality, authorship and meaning; concepts that are mediated by socially conferred authority and celebrity. All may shift in response to new media, technological and social configurations, so we may soon have very different questions to ask.

On the other hand, history shows optimistic speculation about technology is often misfounded. The 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 in 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 previously, lacks innovation. But radically new processes might become increasingly difficult to discover, potentially leading to conceptual stagnation. It is also typical for any art movement to have popularity peaks and troughs. Something for the future to determine.

Conclusion

Our choice of these ten questions arose from many hours of discussion and debate between the authors, including difficult choices about what to include. No doubt others will see different questions to be of greater importance, or perhaps dispute the usefulness of drawing up top ten lists at all. Our purpose here is not to be prescriptive about what is important in generative art, rather to ask how we might stimulate productive critical discussion on what makes generative art interesting as an art 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

[1] William Ross Ashby, *Can a mechanical chess-player outplay its designer?*, Br J Philos. **3** (1952), no. 9, 44–57.

[2] George David Birkhoff, *Aesthetic measure*, Harvard University Press, Cambridge, MA, 1933.

[3] Margaret A. Boden and Ernest A. Edmonds, *What is generative art?*, Digital Creativity **20** (2009), no. 1 & 2, 21–46.

[4] J. David Bolter and Richard Grusin, *Remediation: understanding new media*, no. xi, 295, MIT Press, Cambridge, Mass., 1999.

[5] P. Bourdieu, *The rules of art: Genesis and structure of the literary field*, Stanford Univ Pr, 1996.

[6] Oliver Bown and Jon McCormack, *Creative agency: A clearer goal for artificial life in the arts*, ECAL (2) (George Kampis, István Karsai, and

Eörs Szathmáry, eds.), Lecture Notes in Computer Science, vol. 5778, Springer, 2009, pp. 254–261.

[7] ———, *Taming nature: tapping the creative potential of ecosystem models in the arts*, Digital Creativity **21** (2010), no. 4, 215–231.

[8] William Brooks, *John cage and history: Hymns and variations*, Perspectives of New Music **31** (1993), no. 2, 74–103.

[9] Peter Cariani, *Creating new informational primitives in minds and machines*, Computers and Creativity (Jon McCormack and Mark d'Inverno, eds.), Springer, 2012, p. (in press).

[10] 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).

[11] Denis Dutton, *The art instinct: beauty, pleasure, and human evolution*, Oxford University Press, 2009.

[12] Philip Galanter, *What is generative art? complexity theory as a context for art theory*, 6th International Conference, Exhibition and Performances on Generative Art and Design (GA 2003), Milan., Available online at `http://www.generativedesign.com/` and at `http://www.philipgalanter.com/downloads/ga2003_paper.pdf`, 2003.

[13] Paul Griffiths, *Aleatory*, 2012.

[14] Stephen A. Hedges, *Dice music in the eighteenth century*, Music and Letters **58** (1978), no. 2, 180–187.

[15] Douglas R. Hofstadter, *Staring emmy straight in the eye–and doing my best not to flinch*, Virtual Music: Computer Synthesis of Musical Style (David Cope, ed.), MIT Press, Cambridge, MA, 2001, pp. 33–82.

[16] Leonard Koren, *Which "aesthetics" do you mean? : Ten definitions*, Imperfect Publishing, 2010.

[17] Penousal Machado, Juan Romero, and Bill Manaris, *Experiments in computational aesthetics: An iterative approach to stylistic change in evolutionary art*, The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music (Juan Romero and Penousal Machado, eds.), Natural Computing Series, Springer, 2008, pp. 381–415.

[18] Colin Martindale, Paul Locher, and Vladimir M. Petrov (eds.), *Evolutionary and neurocognitive approaches to aesthetics, creativity and the arts*, Foundations and Frontiers in Aesthetics, Baywood Publishing Co., Inc., 2007.

[19] Jon McCormack and Alan Dorin, *Art, emergence and the computational sublime*, Second Iteration: conference on Generative Systems in the Electronic Arts (Alan Dorin, ed.), CEMA, Melbourne, Australia, 2001, pp. 67–81.

[20] T. Nagel, *What is it like to be a bat?*, The philosophical review **LXXXIII** (1974), no. 4, 435–450.

[21] Anthony O'Hear, *Art and technology: An old tension*, Royal Institute of Philosophy Supplement **38** (1995), 143–158.

[22] John Onians, *Neuroarthistory*, Yale University Press, New Haven, 2007.

[23] G. H. R. Parkinson, *The cybernetic approach to aesthetics*, Philosophy **36** (1961), 49–61.

[24] R. Plotkin, *The genie in the machine: How computer-automated inventing is revolutionizing law and business*, Stanford University Press, Stanford, California, 2009.

[25] V. S. Ramachandran and William Hirstein, *The science of art: A neurological theory of aesthetic experience*, Journal of Consciousness Studies **6** (1999), 15–51.

[26] Alan M. Turing, *Computing machinery and intelligence*, Mind **59** (1950), 433–460.

[27] V. Vinge, *The coming technological singularity: How to survive in the post-human era*, Vision-21: Interdisciplinary Science and Engineering in the Era of Cyberspace (G. A. Landis, ed.), NASA Publication CP-10129, 1993, pp. 115–126.

[28] Mitchell Whitelaw, *System stories and model worlds: A critical approach to generative art.*

[29] _____, *Metacreation: art and artificial life*, MIT Press, Cambridge, Mass., 2004.

[30] R. Williams, *Keywords: a vocabulary of culture and society*, Oxford University Press, Oxford, 1976.