

GENERATIVE SOUND ART AS POETIC POETRY FOR AN INFORMATION SOCIETY

Alice Eldridge

Centre for Creative Research into Sound Arts Practice
University of the Arts, London, UK
alice@ecila.org

ABSTRACT

Recent conversation in both the sonic and generative art communities envision a practice which engages with both formal and cultural or conceptual concerns. At the same time there is a sense amongst practitioners in the generative arts community that there is room to develop the methods of the practice. This paper sketches a simple, practical approach to generative sound art which draws from philosophical observations of the emerging information society and implements a sound-based generative scheme as a simple, literal illustration of a practice which looks outward to society and inward at the materials of its practice. It is suggested that a digital sound art of this nature has potential as a discourse for contemporary society, a poetic playground for coming to terms with the implications and challenges of the information age.

1. INTRODUCTION

Seth Kim-Cohen's recent vision of a conceptual sonic art which engages "both the non-cochlear and the cochlear, and the constituting trace of each in the other." [15], p.xxi. resonates loudly with an ongoing debate in the generative arts. Ten years ago, Florian Cramer surveyed the landscape of software arts practice and drew a rather gloomy distinction between a 'culturalist' software art and 'formalist' generative art [4]. Gloomy, because he saw the concerns and activities of his contemporaries as falling into one or other camp with little cross-talk and felt that neither approach, in isolation, held promise. This analyses was pessimistically reiterated by Troels Degn Johansson, who presented it as the "no future" of software art [14], and again in Inke Arn's distinction: generative art \neq software art. [1]. Mitchell Whitelaw offered a more positive response, suggesting that the "future of software art practice¹ could lie in a fusion of formalism and culturalism" [21], p.1. Analyses in both fields imply a missed opportunity and envision a contemporary praxis which looks both inward (at its materials and form) and outward (at the world and its meaning), toward the creation of works which speak to both the senses and the intellect.

¹This phrase is taken inclusively, to refer to software art, generative art and algorithmic composition, i.e. all algorithmic software-based practices which are concerned with making software which makes something.

These debates in both the sonic and software arts inspire considerations of how practice at the intersect, a generative sound art, might approach this vision. This paper offers a simple, practical illustration. Section 1.1 outlines Whitelaw's *system story*. The exposition of this critical device is an attempt to open up the formal, aesthetic concerns of the generative art community to a broader cultural context, so bridging Cramer's formalist-culturalist divide. Section 1.2 examines current issues in generative art practice. It is suggested that some of the problems currently identified in the field may be a symptom of importing models, methods and their attendant perspectives from the engineering sciences and that the field could benefit from more exploratory, experimental approach to design and implementation of generative systems. Inspiration for an alternative is drawn from Luciano Floridi's observations of the impact of digital technologies in the broader information society. Section 2 introduces a simple class of generative system inspired by Floridi's concept of *ontological friction*. This is presented as possible alternative to the standard practice of algorithmic composition. Rather than sonifying the numerical outputs of a formally implemented algorithm, the generative process is constituted in the sound in which it is also manifest. In section 3, it is proposed that the poetic and creative nature of computer music, and generative sound art in particular, has unique potential as a discourse for coming to terms with the challenges of contemporary society.

1.1. System stories and critical generativity

The potential for software art, in general, to represent and interrogate the cultural complexities of contemporary society has been suggested in the past (e.g. [17]). In adopting formal, complex systems as a basic generative tool, software art has the potential to convey not only an *image* of cultural situations, but more powerfully, to present a systemic abstraction, or model. "Abstract generative art", suggests Whitelaw "performs *cosmogony*: it brings forth a whole artificial world, saying *here is my world* and *here is how it works*." [21], p.5. In collapsing the concept notation and execution, such works are in quite a unique position to illustrate, explore and critique "how it works". In Kim-Cohen's propositional terms, we appear to be tooled up and ready to talk *about*.

Whitelaw observes, however, that this potential remains

largely unfulfilled. Surveying a number of visual artworks he suggests that where complex systems are deployed, the predominant interest is in their formal, generative potential, an inward-looking, utilitarian and non-reflexive concern with the form of the emergent structure. He suggests that what is needed is a critical approach to generative art, in order to open up these formal systems to discussion and critique, drawing out relations between their internal workings and the outside world. Whereas previous discourse focused on the materials and process (e.g. [4]), Whitelaw argues that a culturally-relevant critique must necessarily focus on the formal systems themselves, rather than their sensory outputs.

Drawing from the way in which Artificial Life simulations have been critiqued in the humanities (e.g. [12], [13]), Whitelaw develops the *system story* as a critical device. For Whitelaw the ‘system’ about which the story is told very specifically refers to the abstract, formal structures (the objects, relations, actions and processes of the formal system), as distinct from either the language-specific text in which it is implemented, or its material (sonic, visual, sculptural etc.) manifestation. The system story then, is a retelling of the narrative of the entities, relations, ontologies and processual structures of the software system. The system story can be used to engage with the formal object and draw out its implications. It provides a means to “connect - critically, prospectively, speculatively - entities and relations within the system, with entities and relations outside it.” [21], p.3. The cultural critique afforded by system story is the bridge he builds in an attempt to span Cramer’s formalist-culturalist divide.

Whitelaw’s critique ultimately aims to open the door to a more sophisticated, experimental or critical approach to the practice of generative art. Rather than simply modelling the world as we know it, we can begin to envisage ways of realising Manovich’s vision of a generative art which provides not just images of, but *imaginings* of the social, cultural, personal, material etc. systems in which we live. Viewed as a discourse which communicates system stories, generative art has the potential to ‘tinker critically’ and explore the world in which we live.

1.2. Beyond Algorithmic ready-mades

This feeling that there are as yet unexplored opportunities for contemporary sonic and software practices is also evident amongst practitioners in the generative art community. It has recently been suggested that it is time to develop our generative art methods: “Generative art must do more than simply implement formal systems imported from the sciences.”² [7], p.18.

Mathematics and the biological sciences offer a rich compendium of algorithmic toys for software artists. Motivation for the adoption of a particular model is often

couched in terms of the relevance of properties of its particular emergent structure or behavioural dynamics to a specific compositional or interactive task: Xenakis’ experimentations with stochastic granular techniques reflect his interest in the (auto-) creation of “higher order sonories” (e.g. *Analogique A et B*, 1958-59); Blackwell’s adoption of swarm algorithms [3] reflects his interest in the self-organisational forces at play in group improvisation, etc. Under this approach, a formal model is implemented and the numerical outputs are then mapped to particular sonic parameters. The design of musically meaningful mappings from numerical output to relevant parameters is sometimes natural but often far from intuitive. The challenge is framed as that of establishing meaningful correspondences between the underlying principles of the model and the appearance and behaviour of the artwork, and is raising its head again with the new generation of ardent swarm enthusiasts [2]. To call for new approaches to generative design is not to admonish the use of existing algorithms for specific compositional tasks, but to encourage a more experimental approach in order to expand the repertoire and scope of the practice.

Dorin et al recognise the pedagogical value of these ‘readymade’ systems (particle systems, cellular automata, physical simulations etc.) but express concern that the systems “operate as ‘black boxes’ whose internal operations are obscure to the artists and designers who use them” [7], p.18. The recent population explosion in artist-programmers has created lively communities around environments such as Max-MSP and Processing. The open source ethos of these communities coupled with the portability of code, means that algorithms, encapsulated in libraries, are shared amongst practitioners more readily than ever before. We no longer have to implement algorithms from scratch. On the one hand this expedites development, obviating countless wheel-reinventions, reducing laborious debugging of personal implementations of textbook algorithms and so affording more time to focus on the more ‘creative’ side of development - how these algorithms are mapped to specific visual or sonic parameters, for example. The flip side is that in not going through the process of reinventing the wheel, programmers have no need to get to grips with the workings of the wheel. This arguably reduces the scope for creative tinkering and with it, the chances of inventing new methods.

Similar cries for new directions are heard within the Evolutionary Arts (EA) community. Galanter writes that “After a great deal of initial promise and enthusiasm, evolutionary art seems to have hit a premature and disappointing plateau.” [10], p.216. The field, suggests Galanter, is suffering from a conceptual malaise. The malaise is seen to stem in part from two perennial problems: the design of genetic representations which afford boundless innovation and the specification of fitness functions able to drive vast populations over many generations. These technical issues accompany what Galanter sees as a deeper theoretical incoherence. The specification of a fixed fitness function implies a bottom-up approach: a particular class

²This call, as the authors note, can be seen as an echo of Fenton’s warning that art should not become “the handmaiden of science”. Fenton was responding to Jack Burnham’s enthusiasm for technological, process-based, autonomous art in the 1960s, work that shares many attributes with contemporary generative art.

of form or behaviour is preconceived and the genetic operators designed in order to reach that goal. Whilst the EA community looks to the boundless creativity of nature for inspiration, natural evolution is not teleological. Evolutionary pressure is neither fixed nor determined from Above, but arises from interactions between the agent and the components of the ecosystem in which it is situated. Galanter prescribes a “dedication to an aesthetic of truth to process and dynamism” [10], p.213, as the remedy, suggesting that what is essential to generative art is not any particular form of a particular material, but the harnessing of process.

Process is undoubtedly central to generative art, but the endemic malaise identified by Galanter can also be seen as a symptom of a more fundamental incoherence: the mismatch between the problem-solving engineering domain in which artificial evolution methods were developed and the art-making context in which they are now applied. The adoption of evolutionary techniques as search methods can carry with it the problem solving approach of the engineering sciences in which they originate. Such an approach is apposite to design applications (such as automated synthesis e.g. [22]), but arguably conflicts with art-making in which there may not be a specific ‘problem’ to which a formally specifiable solution exists. The malaise which Galanter identifies can be seen as another sign that there is room for a new approach to generative art, one that moves beyond borrowing algorithms and methods from the sciences and begins to tinker with bespoke approaches.

Kim-Cohen talks of an as-yet unestablished category of practice in which the work of art could be conceived of as “the simultaneous creation of a message and the language of the message’s transmission” [15], p.xxiii-xxiv. The central proposition of this paper is that in focusing solely on the formal process in generative art and disregarding the sensory material (in this case sound) in which it is manifest is to miss a trick: that one way to engage both the ear and the intellect may be to develop methods from the ground up in which the process and its sensory manifestation are mutually constituted.

The remainder of this paper aims to provide a simple illustration of one way this approach could be developed. Rather than lifting algorithms from the sciences, inspiration is taken from a recent philosophical analyses of contemporary society. A specific systemic characteristic which is identified motivates the design of a simple generative scheme in which sound is inextricably bound up in the generative process.

1.3. Inspiration the infosphere

Information philosopher Luciano Floridi has written extensively on the impact of Information and communications technologies (ICTs) on contemporary society, highlighting both the metaphysical and practical, ethical implications. The ubiquity of ICTs in many areas of society, he claims, is radically changing not only our interactions with the world, but our essential views of who we are and

the world in which we live. He presents this as the Fourth Scientific Revolution³.

In expounding his position, he coins the term *infosphere* [9] to denote the environment constituted by all informational entities (analogue and digital), as distinct from the biosphere - the region of the planet which supports life. This is us, our laptops, libraries, digital instruments, museums, phones, intelligent cars, their properties, interactions, processes and mutual relations. His claim is that the ubiquity of ICTs is fundamentally transforming the intrinsic nature of social structures and artefacts and our relationship with them. Another neologism, *re-ontologising*, is coined to refer to this process. With these concepts in hand, Floridi claims that ICTs are “*re-ontologising the very nature of (and hence what we mean by) the infosphere*” [9], p.2, and that these changes will bring about profound transformations in and corresponding challenges for information societies in the near future.

The most obvious and significant changes are the transition from analogue to digital data and the exponential growth of our digital realm. Floridi suggests that the radical extent of the impact of ICTs is largely due to the “*fundamental convergence between digital resources and digital tools*” ([9], p.4). The information technologies available (e.g. software, databases, communication channels and protocols etc.) are increasingly similar to, and so fully compatible with, their objects (their data and the conversations and interactions they support). “*This was one of Turing’s most consequential intuitions: in the re-ontologised infosphere, there is no longer any substantial difference between the processor and the processed, so the digital deals effortlessly and seamlessly with the digital.*” (ibid, p.4). For the information society at large, the converging nature of processors and their objects increases the ease with which information can flow in a given environment. Floridi uses the terms *ontological friction* as a measure of the amount of work and effort required to process and transmit information in a given environment.

An analogy can be drawn between the processors and their objects in the infosphere and the algorithmic processes and the ‘materials’ (digital audio and graphic representations etc.) in which they are manifest in the digital generative arts. All are digital in nature, but whilst the ontologies of the former align, in common practice, there is a fundamental rift between the latter, evidenced by the perennial head-scratching over the design of aesthetically meaningful mapping and representation schemes. The analogue of ontological friction in this context is the disjunction between the use of formal algorithms and the sonic material they are intended to organise. The prospect that decreasing ontological friction could increase generative potential is interesting.

³The first three being: Copernican (we are not immobile at the centre of the universe), Darwinian, (we are not distinct from the rest of the animal kingdom) and Freudian (we are very far from being purely rational minds entirely transparent to ourselves).

2. SELF-OBSERVING SYSTEMS

Self-observing Systems is an ongoing project which explores generative sound art methods. The concern in the current study is the construction of a simple generative mechanism which exhibits an ontological alignment comparable to that between ‘processor’ and ‘processed’ as seen in *Infosphere* at large, a system which relates a story about the world, through the material of sound. This is approached by implementing a simple system in which the generative process is constituted in the sound.

2.1. A self-determined sample player

The illustrative mechanism is based on a sample player. This familiar unit generator, which is the digital homologue of various physical music players, represents the ‘processors’ in Floridi’s analysis and caricatures the impact of the de-physicalisation of devices more broadly. The sample player is furnished with a listening module. Rather than playing in a linear fashion, its playback rate (and direction) are determined by features of the audio it plays.⁴

2.1.1. Implementation

The system was built in *beads*⁵ and comprises a variable rate sample player, listener module and short audio file as shown schematically in Fig.1. The sample is played at rate, R . When an onset is detected, absolute playback rate is updated as a function of the most dominant frequency in the current audio buffer and its direction reversed. Noise is added by means of a probabilistic update regime. The probability is inversely proportional to the current rate. This balances the effects of playback rate: at slower speeds, significant audio features will be further apart in time (and less prominent). See Fig.2 for details.

The listener module is an onset detector comprising the following analysis chain: short frame segmenter, fast fourier transform, peak detector and spectral difference measure. The peak detector follows the algorithm described in [6]. In this implementation the audio file is not overwritten.

A resynthesis module is included for ‘decoration’ (it plays no functional role in the process), highlighting the moments of change. When an onset is detected, an oscillator bank is triggered, creating a complex chord. The frequency, duration and amplitude of each of N oscillators is set according to the frequency and strength of the corresponding first N partials in the current buffer. In this implementation, $N = 10$.

⁴The basic design principle adopted here is therefore one of circular causality. This is nothing new. The approach lies at the heart of the cybernetic enterprise [20]. Its generative and interactive potential was explored by early electronic pioneers (consider e.g. Gordon Mumma’s ‘Cybersonic’ Horn piece, *Hornpipe* (1967) and it has experienced somewhat of a revival in recent years, following Di Scipio’s Audible Ecosystemic Interface series [5].

⁵*beads* is a Java library for real time audio: www.beadsproject.net




Figure 1. Self-determined sample player schematic

```

for current output buffer
  look for onsets
  if onset occurs, with probability, P
    extract partials F1...n
    update R according to F1
    play buffer at rate, -1 * R_new
  else play buffer at rate, R
    
```

$$R_{new} = \frac{1}{F_{p1} \times rK} \quad (1)$$

Where F_{p1} is the frequency of the 1st partial, rK is a constant and R can assume negative values. $P = f(\frac{1}{R})$

Figure 2. Pseudo code and rate update equation

2.2. Sonic examples

The system is intentionally playful and lively and this is reflected in the output. A range of samples containing varying degrees of harmonic and gestural complexity were explored. The examples given here include a dripping tap, *Sprids*, a field recording of blackbird song, *Balckbrids* and a live example with pizzicato ‘cello, *Spwang*.

What we hear is an engaging and active exploration of the given sonic material with an apparent intentionality that belies the simplicity of the generative process. In *Sprids*, the harmonic simplicity of the original drip develops a coherent harmonic progression as each drip triggers a change in rate according to its harmonic content, altering the pitch of the subsequent drips. Over time, the long term structure (albeit musically trivial) unfolds according to the spectral content of individual sound events.

In *Balckbrids* there is a basic (and comic) exposition of the gestures contained in the original recording. This gestural exposition would be effective in a live interactive setting. *Spwang* provides a short excerpt from initial experiments with improvised ‘cello. In a performance setting, the self-directed exploration of material has the potential for a lively trading of gestures between instrumental performer and digital components of a performance system, an aesthetic design principle which guides many electro-acoustic improvisation systems (e.g. [19]).

2.3. Other self-observing systems

The principle of self-observing systems is open ended. A further hypothetical example is outlined below.

Self-directed feedback circuits illustrates how a similar generative process can be constituted across all levels of a digital audio system, extending out beyond the confines of a formal algorithm and digital representation of sound through the DAC, speakers, room, mic and back again. It explores ways of auto-maintaining and directing the Larsen effect to create complex, adaptive resonances.

In an analogue system, characteristics of feedback artefacts are determined by the frequency response of the audio system (mic and speakers) and the distance between them. The latter can be modelled digitally using delay lines. In a simple system, the effective gain can be managed by implementing a proportional control algorithm (such as a watt governor) which monitors the amplitude in its input buffers and adjusts the delay time (at audio rates). The gain on each channel can be similarly adjusted. Initial experiments suggest that even this simple mechanism can sustain frequencies other than those promoted by the frequency response of the audio system, creating subtle shifts which are sensitive to tiny changes in the ambient environment. Further variation can be imagined by implementing frequency-dependent delay lines and setting, or evolving, the delay times according to the impulse response of the physical environment, frequency response of audio system and spectro-temporal characteristics of the surrounding acoustic environment. Recent research in automated methods of *avoiding* feedback (outlined in [18]) point to several potentially interesting avenues for exploration.

3. DISCUSSION

The *Self-observing systems* offer a simple illustration of a generative sound art scheme which implements a sonic version of Floridi's ontological alignment. In the design of the generative schema, sound (or its digital representation) is integral to the generative process such that the two are mutually constituted. What we listen to is not the sonification of some numerical process at runtime, but the 'empirical epiphenomenon' of interaction between the two. Under this approach, we circumvent the need to define either mappings to, or representations of, a final sonic manifestation. This is offered in response to calls from the generative art community, that as the field matures, it should move beyond the implementation of algorithms lifted from the sciences; it aims to illustrate how we can explore changes in the broader cultural world, by representing systemic principles of the changing infosphere in the materials, both formal and sensory, of a generative practice.

3.1. Generative sound art as a poetic playground for an information society

The self-observing systems outlined above are also intended as a cartoon-like illustration of how a generative sound practice could reflect upon and explore the implications of significant changes in our broader infosphere. The closing proposition of this paper is that this represents

a potent role for generative sound art in contemporary society.

To return to Floridi's analysis introduced in section 1.3, the metaphysical shift he identifies can be seen in the following key changes: Objects and processes are becoming *de-physicalised*, causing a shift from a materialist to information-based metaphysics in which they become typified, clonable and support independent; As the world fills with intelligent agents, we cease to be discrete individuals and become *networked inforgs*, part artificial, part human, connected informational organisms; *Interactability*, rather than immobility or perceivability, becomes the criterion for existence, even if the interaction is only virtual; The *global infosphere* is merging with the analogue world, creating a world populated by intelligent, active agents which challenge our Newtonian world view.

For the computer musician who has spent the last half century designing intelligent, interactive, modular performance systems toward a networked digital practice, this is nothing new. Artistically, the implications of such developments could even be summarised as an increase in creative opportunity. To the general public, however, this may read as a science fiction scenario. For society at large such changes bring about significant ethical challenges, not least an urgent need to reconcile the technological and natural worlds [8].

The public need for new ways to come to terms with the existential and ethical impact of an increasingly technologised society can be seen in the emergence of new forms of hybridised discourse. *New Scientist*, for example, have just launch a digital publication, *Arc*,⁶ which merges literature (science-fiction) and science (futurology) to explore the impact that technology is having on our lives. "Fiction gives us the chance to explore and be eccentric" says Simon Ings, a novelist, science writer and editor of *Arc*. "If one thing is for sure, the future is not going to be agreed by committee. The future is going to be eccentric. And the best way of predicting the future is to make it up." [16]. New branches of social science are also emerging which adopt the formal modelling and simulation practices similar to those used in generative art practice [11].

The simulation sciences provide ways of understanding the world through modelling. Science-fiction provides a speculative literary discourse for toying with possible futures, but, as Manovich and others have suggested, the generative arts have the potential to speculatively toy with *models* of possible futures. The development of a generative sound art practice has a unique potential as both a poetic (imaginative, symbolic, figurative) and poetic (from the Greek, *ποιητικός*, meaning productive, formative) discourse, one which could engage the ears, intellect and imagination in equal measure.

⁶<http://www.newscientist.com/arc>

4. CONCLUSION

Now, more than ever, we need new forms of discourse which enable us to come to terms with the complexities of contemporary society. Computer music, and more specifically, generative sound art, represents a possible discourse through which we can model our world reflectively, tinker with these models critically, and present these stories in the material of sound itself: to talk *about* the world *through* sound.

5. REFERENCES

- [1] I. Arns, "Read_me run_me, execute_me: Software and its discontents, or: It's the performativity of code, stupid," in *read_me Software Art and Cultures Edition*, O. Goriunova and A. p. Shulgin, Eds. Center for Digital Æstetik-forskning, 2004, p. 178.
- [2] D. Bisig and M. Neukom, "Swarm based computer music - towards a repertory of strategies," in *Proceedings of the Generative Art Conference*, Milano, Italy, 2008.
- [3] T. Blackwell, "Swarm music: Improvised music with multi-swarms," in *Proceedings of the AISB Symposium on Artificial Intelligence and Creativity in Arts and Science*, 2003, pp. 41–49.
- [4] F. Cramer, "Concepts, notations, software, art," March 2002, last accessed: May 21st 2012. [Online]. Available: http://www.netzliteratur.net/cramer/concepts_notations_software_art.html
- [5] A. Di Scipio, "Sound is the interface: From interactive to ecosystemic signal processing," *Organised Sound*, vol. 8, no. 3, pp. 269–277, 2003.
- [6] S. Dixon, "Onset detection revisited," in *Proc. of the 9th Int. Conference on Digital Audio Effects (DAFx-06)*, Montreal, Canada, September 18-20 2006.
- [7] A. Dorin, J. McCabe, J. McCormack, G. Monro, and M. Whitelaw, "A framework for understanding generative art," *Digital Creativity*, (forthcoming).
- [8] L. Floridi, "The digital revolution as a fourth revolution," last accessed: May 12th, 2012. [Online]. Available: <http://www.philosophyofinformation.net/massmedia/pdf/bbc-1.pdf>
- [9] —, "A look into the future impact of ICT on our lives," *The Information Society: An International Journal Volume*, vol. 23, no. 1, 2007.
- [10] P. Galanter, "Truth to Process - Evolutionary Art and the Aesthetics of Dynamism," in *Generative Art Conference*, Milan, 2009.
- [11] K. Gilbert, N. abd Troitzsch, *Simulation for the Social Scientist*, 2nd ed. Open University Press, 2005.
- [12] K. Hayles, *How We Became Posthuman*. University of Chicago Press, 1999.
- [13] S. Helmreich, *Silicon Second Nature: Culturing Artificial Life in a Digital World*. University of California Press, 1998.
- [14] T. D. Johansson, "Mise en abyme in software art: A comment to florian cramer," in *read_me Software Art and Cultures Edition*, O. Goriunova and A. Shulgin, Eds. Center for Digital Æstetik-forskning, 2004, p. 151.
- [15] S. Kim-Cohen, *In the Blink of an Ear: Toward a Non-Cochlear Sonic Art*. Continuum International Publishing Group, 2009.
- [16] E. Liston, "New scientist's new digital magazine combines science-fiction and futurology," *The Independent*, March 2012.
- [17] L. Manovich, "Abstraction and complexity," 2004, last accessed: May 28th 2012. [Online]. Available: http://www.manovich.net/DOCS/abstraction_complexity.doc
- [18] E. Perez-Gonzalez and J. Reiss, "An automatic maximum gain normalization technique with applications to audio mixing," in *124th AES Convention*, Amsterdam, Netherlands, 2008.
- [19] D. Van Nort, P. Oliveros, and J. Braasch, "Developing systems for improvisation based on listening," in *Proc. of the 2010 International Computer Music Conference*, New York, June 1-5 2010.
- [20] H. Von Foerster, M. Mead, and H. L. Teuber, *Cybernetics: Circular Causal and Feedback Mechanisms in Biological and Social Systems*. New York: Josiah Macy, Jr. Foundation, 1953.
- [21] M. Whitelaw, "System stories and model worlds: A critical approach to generative art," in *Readme 100: Temporary Software Art Factory*. Norderstedt: BoDs, Dec 2005, pp. 135–154.
- [22] M. Yee-King, "An autonomous timbre matching improviser," in *Proceedings of the 2011 International Computer Music Conference*, 2011.