

Contents

1	Introduction	2
1.1	Intellectual Precursors	4
1.2	Core Concepts	6
2	Evolutionary Search	9
2.1	Genetic Algorithms	10
2.2	Using Evolutionary Algorithms in Music	12
2.3	Approaches to GA Design	14
2.4	Variations on a Theme	17
2.5	Evolution and Usability	17
3	Multi-Agent Compositions	18
3.1	Agent-Based Modelling	19
3.2	Emergence in Human and Artificial Music	20
3.3	Modelling Musical Communities	21
3.4	Agent-Based Models at Different Scales	23
4	Adaptive Collaborators	24
4.1	Adaptive Behaviour and Musical Interaction	25
4.2	Behavioural Objects	26
4.3	Subsumption Architectures for Musical Agents	28
5	Creative Ecosystems	29
5.1	Audible Ecosystems	30
5.2	Computational Ecosystems for Creative Discovery	32
6	Looking Forward	35

Biologically-Inspired and Agent-Based Algorithms for Music

Alice Eldridge^a, Oliver Bown^b

^a*University of Sussex*

^b*University of New South Wales*

Abstract

This chapter examines a range of approaches to algorithmic music making inspired by biological systems, and considers topics at the intersection of contemporary music, computer science and computational creativity. A summary of core precursor movements both within and beyond musical practice (ALife, cybernetics, systems art etc.) sets the scene, before core models and algorithms are introduced and illustrated. These include evolutionary algorithms, agent-based modelling and self-organising systems, adaptive behaviour and interactive performance systems, and ecosystemic approaches to composition and computational creative discovery. We close by reviewing themes for future work in this area: autonomy and agency, and the poetics of biologically-inspired algorithms.

1. Introduction

Whatever vibrates is a musical instrument: whatever is stable is a mechanical brain – the difficulty lies in making a particular one. (Ashby Aphorism 158).

For all of humankind's creative achievements, we in turn were made by a more powerful creative force: biological evolution. Since Darwin and Wallace's great revelation (Darwin 1861), it has become widely accepted that the astonishingly beautiful and complex structure and behaviours of the living world have taken shape through a remarkable process that is mechanical, blind and purposeless. This sublime beauty has inspired art since its beginnings, but whereas we have always incorporated natural *form* in our

paintings, sculpture and music, artists working with code now draw upon *processes* of the natural world.

The arrival of general purpose computers in the middle of the last century transformed not only science, but compositional practice in ways that are documented throughout this book. Of importance to this chapter, it enabled us to harness *behaviours* inspired by natural systems, formalised by biologists and computer scientists into algorithms, in order to develop, perform and compose with software instruments. We now borrow from the designs of specific biological organisms, and the properties and processes of complex natural systems, as well as from the creative mechanism of evolution itself.

In this chapter we examine a range of approaches to algorithmic music making inspired by biological systems. In doing so we cover topics that are located at the intersection of contemporary music, computer science and the study of creativity: optimisation and problem-solving using evolutionary methods; emergence, self-organisation and complexity; adaptive behaviour; and autonomy and self-determination. Section 1.1 provides a brief historical context of the core intellectual, musical and social movements which influence contemporary creative practice. Section 1.2 provides a primer in the concepts and tools developed for the study of systems which are foundational to the specific approaches described in the following sections.

Endeavours in this field are often hybrid and idiosyncratic and cannot be neatly categorised. Nevertheless we organise an overview of the key musical motivations, concepts and computational methods of the field under four themes which map the topics above. Section 2, *Evolutionary Search*, outlines the application of evolutionary algorithms to design issues and opportunities associated with algorithmic music. In Section 3, *Multi-Agent Compositions*, we look at the ways in which agent-based modelling has been used to compose emergent, self-organising music. Section 4, considers how the study of adaptive behaviour has inspired the design and realisation of *Adaptive Collaborators* – interactive software systems which begin to enable active electro-acoustic partnerships. Many of these ideas come together in Section 5, which describes the development of *Creative Ecosystems* inspired by ecological principles. We end by mentioning two themes that will guide future work: autonomy and agency, and the poetics of biologically-inspired algorithms.

1.1. Intellectual Precursors

Contemporary practice in the area of biologically-inspired computer music can be best understood against the backdrop of a series of interrelated intellectual, cultural and social currents arising in the late 19th and early 20th Centuries. Central to these was a shift from an essentialist paradigm toward the relational thinking championed by the thinkers behind general systems theory (GST) and cybernetics. Systems thinking derives from the work of biologist Von Bertalanffy, who sought to abstract from the intractable messiness of actual biological, social, economic, chemical etc. systems in order to define *general* principles of dynamic interaction. (e.g. Von Bertalanffy 1950). Similarly, cybernetics sought to understand processes of control and communication across electronic, mechanical, biological or economic systems in terms of common principles, such as regulatory feedback. Cybernetics, expounds pioneer Ross Ashby, “treats, not things but ways of behaving. It does not ask what is this thing? but what does it do? ” (Ashby 1956, p.1).

These new paradigms had far reaching influence in society and the arts as well as engineering and science, and underpin contemporary practice both conceptually and methodologically. The conceptual relationship between organisms and machines had been thoroughly explored in western 19th century thinking, typified in popular culture by publications like *Frankenstein*. Darwin’s theory of natural selection (Darwin 1861) sealed the direction of thinking into the 20th century. Cybernetics made the first formal steps towards truly integrating the study of natural systems with the creation of artificial ones. Later the science of complexity and chaos – the new computational magic made famous by Mandelbrot Sets and Lorenz Attractors – in step with developments in theoretical biology, blossomed into the discipline of artificial life (ALife) (Langton 1989). In various ways, each sought to explain complex systems in terms of the interactions between the mutually interrelated parts of which they were comprised. ALife advanced the idea that complex real-world processes could be modelled computationally, imagining that we might not only recreate the phenomena of the biological world, but also uncover principles of life divorced from the biological substrate of life-as-we-know-it – “life-as-it-could-be” (Langton 1989, p.1) – and in doing so aimed to reveal general principles of biology, both natural and artificial.

The inherently interdisciplinary and conceptually profound nature of these movements entailed a close relationship between the sciences and arts. Not only did pioneering practitioners engage across domains, but the cybernetic vision inspired a revolution throughout modern art. Recapitulating the ear-

lier shift from essentialism in the biological sciences, Jack Burnham's 'system aesthetics' (Burnham 1968) and Roy Ascott's 'behavioural tendency' (Ascott 1967) drew attention from self-contained objects characteristic of modern formalism, to a post-modern open-ended and immersive experimentation in which feedback loops stitched previously distinct elements of artist, artwork and observer into an indivisible whole. This new intimacy between technology and arts was celebrated by the landmark exhibition and accompanying book, *Cybernetic Serendipity* (Reichardt 1971) which showcased computer generated work and cybernetic devices from the pioneering players: Stanford Beer, Earle Brown, John Cage, Edward Ihnatowicz, Ben Lapowsky, Frieder Nake, Nam Jun Paik, Gordon Pask, Karlheinz Stockhausen, Jean Tinguely and Iannis Xenakis.

The liberation and expansion of musical sound that was an essential part of Futurism at the start of the 20th century had developed by midcentury into a radical re-thinking of what music could become. Early electronic music collectives such as the Sonic Arts Union (Robert Ashley, David Behrman, Alvin Lucier and Gordon Mumma) and their contemporaries, John Cage, Christian Wolff, David Tudor etc., continued to bring into question essential assumptions about music; breaking free of the western cannon, their music spoke to and drew from world music as well as architecture and science. Their early compositional explorations of principles such as chance and self-organisation continue to inspire the design of digital music systems today (see Section 3.2). These ideas were linked with social and political views that questioned the existing social order: political directions, both to the left and to the right, drew in different ways on Darwinian thinking and its derivatives, through issues such as social Darwinism and sociobiology.

Evolutionary and adaptive software art and music grew into an entity in its own right in the 1990s as part of the second wave of biologically-inspired systems thinking, associated with ALife. Artworks created by evolutionary computing techniques were pioneered by William Latham (Todd and Latham 1991), Karl Sims (Sims 1991) and Jon McCormack (McCormack 1993), and principles of ALife were explored widely by artists in a variety of ways (see (Whitelaw 2004) for a good introduction). Experiments using AI to solve musical problems were well underway (Cope 1996; Todd and Werner 1998) including biologically-inspired cognitive systems such as artificial neural networks.

The concept of the ecosystem, which had been established in the 1930s as a critical unit of study for ecology (Tansley 1935), drew the attention of ALife

arts practitioners with the promise of generative autonomy, engendered by processes of feedback, coupling and coevolution. Ecosystem-based creative works emerged, quite literally, with work such as McCormack’s *Eden* and the ecosystemic approach to composition pioneered by Agostino Di Scipio discussed in Section 5.

The 21st century saw the mass adoption of computing technologies in the developed world, centred around the Internet. Increasingly productive programming languages and methods emerged, including languages and development environments specifically designed for creative coding. These factors have combined to give greater power to individual creative programmers working in music and the arts, and as the theorist De Landa proposed, creative coders now routinely ‘hack’ the rules of thermodynamics, mathematics and biology as if they were malleable artistic materials (DeLanda 2002). Simple biological models may now be considered part of the creative coding canon: cellular automata, swarm models and evolutionary algorithms are routinely included as examples in creative coding environments, such as *Processing* and *OpenFrameworks* (see e.g. Shiffman et al. 2012).

With this whirlwind tour we have begun to map how the work of the pioneering thinkers of the early 20th century has influenced how we approach biologically-inspired and agent-based algorithms in music today. Our aim in this chapter is not to give a comprehensive overview of each field (pointers are given to many excellent references which provide these), but to give a flavour of the challenges and opportunities afforded by biologically-inspired and agent-based computing in digital music making. We proceed with the introduction of some key concepts.

1.2. Core Concepts

“There exist models, principles, and laws that apply to generalized systems or their subclasses, irrespective of their particular kind, the nature of their component elements, and the relationships or ‘forces’ between them. It seems legitimate to ask for a theory, not of systems of a more or less special kind, but of universal principles applying to systems in general.” Von Bertalanffy (1950, p.32).

The approaches to biologically-inspired composition discussed in this chapter draw heavily on the conceptual and methodological tools of general systems theory, cybernetics, chaos theory and the study of complex systems as

well as more specific fields of ALife, connectionism and artificial intelligence (AI). Although aims and techniques vary in each specific field, some key concepts and modelling principles are shared. In this section we outline some of these foundational concepts.




Figure 1: Trajectories showing a) the point attractor of a damped pendulum and b) cyclic attractors of an undamped pendulum.

Central to modelling and understanding systems, from simple particles to complex ecosystems, is the notion of *system state*. The set of all possible states of a system is called its *state space*, and a systems approach advances by studying the *trajectory* of the system states through this space.

Take a single damped pendulum, swinging back and forth in a single plane. Its state can be defined in terms of its current position and velocity. Due to the forces of friction and gravity it swings with ever decreasing energy, eventually coming to rest (Figure 1a) at its stable equilibrium point or *point attractor*. For this closed system (closed referring to the fact that there are no other forces acting on it) the pendulum is attracted to this point no matter where it starts in the state space. If, on the other hand, the pendulum had zero friction, then it could swing forever. In that case we would find multiple *cyclic* attractors, each described by the set of states that the pendulum passes through in its swinging cycle (Fig 1b).

A certain class of systems – such as a double rod pendulum or the famous Lorenz system (Fig 2) (Lorenz 1963) – exhibit chaotic behaviour, which is




Figure 2: A 2D projection of the strange attractor of the Lorenz equation.

exemplified by state trajectories that are close, but never actually repeat. These paths are known as *strange attractors*. Chaotic systems are sensitive to initial conditions, meaning that similar starting states can tend towards wildly different outcomes: “the present determines the future, but the approximate present does not approximately determine the future” (Lorenz cited in Danforth 2013). Note that we are still talking here about closed, deterministic systems; neither external stimuli nor randomness are necessary to produce incredibly rich system behaviours, even when the system might consist of a very simple update rule. This principle has been of great interest to algorithmic musicians wishing to achieve rich, complex outcomes with algorithmic efficiency.

States that are not an attractor state will typically lie in the *basin of attraction* of an attractor state. A helpful metaphor is to think of how rain falling on sloping ground has an attractor state depending on where it lands: rain in the Alsace region in France enters the Rhine river basin, and runs into the North Sea; rain in the Sancerre region enters the Loire river basin and ultimately runs into the Atlantic Ocean. The lie of the land is much like the set of state-space trajectories of a system. This landscape metaphor is similarly useful to conceptualise evolutionary search through a *fitness landscape*, which we return to in Section 2.

Systems thinking can be applied to biological systems at a range of levels: a cell is a system of chemical and energy transfer; the heart is an open

oscillatory system; an organism is a self-regulating system that operates to keep certain critical parameters within acceptable boundaries (to stave off death – the ultimate point attractor); an evolving population may arrive at an *evolutionarily stable state* or shoot off on a trajectory of runaway coevolution. Although these are very different in their details, systems thinking enables a common language for the study of system behaviours across domains.

Such generalist thinking is alluring for musicians. The language of state, trajectory, transition rules, attractors, basins of attraction, etc. resonate with sonic and musical experiences and concepts and inspire new approaches. Software models of complex, dynamic systems offer rich possibilities for musical composition and interaction, for the imitation of existing musical styles, the creation of esoteric new forms or as frameworks for human-computer musical interaction. In cases such as the modelling of human rhythmic perception using oscillatory models (Large and Palmer 2002), the system dynamics link explicitly with psychological theories of music perception. This state-based approach to studying systems is foundational to many of the topics discussed in this chapter.

2. Evolutionary Search

“With the help of an electronic brain the composer turns into an astronaut pressing buttons of his musical spaceship to introduce coordinates and keep the course of his vessel on its journey through constellations and galaxies of sound, controlling from his easy-chair what the imagination of yesteryear could have envisaged only remotely in its wildest dreams.” (Xenakis 1971, p.124)

Accounts of journeys through unimagined and uncharted territories appear throughout early electronic and digital music discourse, as well as recent generative art and computational creativity literature (McCormack and Dorin 2001). Computational algorithms have clear musical potential, but vast swathes of these spaces of possibilities contain nothing of interest, and efficiently navigating to the interesting areas is a non-trivial task. Searching for solutions to problems is a well developed field in AI, and its application in the arts also has a rich history. The search for fruitful solutions cannot be random, and as one of a number of directed algorithmic search strategies, artificial evolution promises a powerful vehicle for discovery (see McCormack 2008, for a detailed overview of the challenges and conceptual issues involved in algorithmic creative search).

The theory of evolution by natural selection (Wallace 1858; Darwin 1861), radically transformed our understanding of nature by describing a seemingly simple, blind process that explains the origins and development of Earth’s biological complexity. In the reproduction of biological organisms, they observed, there is a predominant continuity of form and behaviour (heredity), but this overall continuity is corrupted by minor random mutations (variation) that occur naturally and that may accumulate over time into radical morphological changes. Whilst the majority of these reproductive mutations are detrimental, certain variations in an organism’s ‘design’ improve its survival rate and reproductive success, relative to its peers. Those better-off variants by definition are prone to grow in number, whilst the weaker variants dwindle. They posited that over time, this alone sets the sufficient conditions for new species to form and develop sophisticated adaptations to their environments. Intense competition for resources enhances this evolutionary effect as weaker variants are rapidly displaced by their stronger counterparts. The combined result is *natural selection*, in which an invisible hand creates life forms intricately adapted to flourish in their present environment, *as if they were designed to do so*. The discovery of the genetic system (Mendel 1866) provided the underlying mechanism for heredity, mutation, and sexual recombination that is critical to modern evolutionary theory.

2.1. Genetic Algorithms

John Holland was the first to recognise that the power of this process could be harnessed in computational models as a search tool in optimisation and introduced the genetic algorithm (GA) (Holland 1975). Many variants followed, ultimately abstracted and generalised into a category known as population-based, metaheuristic optimisation (see Eiben and Smith (2003); Mitchell (1998) for general introductions and Burton and Vladimirova (1999a); Husbands et al. (2007) for musically oriented outlines).

To illustrate how evolutionary concepts can be adapted for optimisation tasks, consider how we might go about designing a paper airplane (Figure 3). We could take a pile of paper and randomly fold pieces to create an *initial population*. A description of the points at which we folded the paper – e.g. x,y coordinates of start and end points of each fold, and so on – represent the *genotype* (roughly the underlying ‘design’) of each individual plane, the resultant physical form being the *phenotype* (the actual form). This population of phenotypes is then evaluated by assigning each candidate a fitness score according to how successful a solution it provides. In this case,

we might cast planes across the room and measure how far they travel. This allows us to quantitatively compare the efficacy of each phenotype, i.e. a *fitness function* (flight length).

In order to develop a population of planes achieving longer and longer flight lengths, we could then preferentially select those planes which flew the furthest and make modifications to the ways in which they were folded. This might include making minor random *mutations* to individual folds, or combining the folds from two winning planes. In biological reproduction, the latter is known as *crossover*; the genetic material from parents is mixed, as in sexual reproduction. We could then make a set of new planes which implemented these variations – the *replacement scheme* – either mixing them with solutions from the first round, or creating an entirely new population before launching them all in the air again.

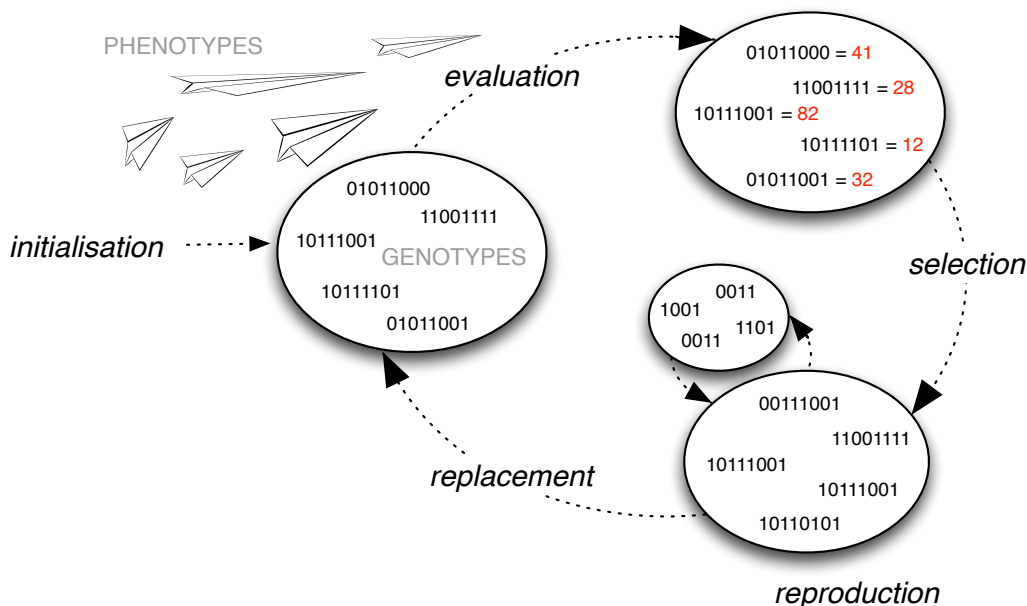


Figure 3: Outline of a genetic algorithm for evolving paper airplane designs.

Iterating this process can, in theory, lead to functional paper airplanes. In practice, there are many factors that affect how successful the outcome is. The encoding scheme (how we represent the phenotype as a genotype), genetic operators (mutation and reproduction schemes) and the fitness function together define what we call a *fitness landscape* across the genotype

space. Fitter solutions sit at the tops of hills, and the less fit solutions down in valleys. A well-designed GA acts to guide a population of candidate solutions across the fitness landscape towards higher ground, until an individual or percentage of the population reaches a pre-specified value for an optimal solution.

A fitness landscape should be smooth, meaning that very similar genotypes result in very similar fitness scores (two similar planes are likely to fly similar distances): the more it looks like Mount Fuji, and less like the Manhattan skyline, the easier it is to make small and gradual steps upwards towards fitter solutions. Too many ‘foothills’ and the GA may never arrive at the best designs, but get stuck at *local optima*. A vast literature within the field of evolutionary computation discusses how the efficacy of search (navigating this landscape) can be improved by different approaches to designing population structure (Collins and Jefferson 1991; Husbands 1993), selection method (e.g. Eiben and Smith 2003; Mitchell 1998), replacement schemes and other factors. In practice, the best choice of genetic representation, operators, and evaluation functions are often problem-specific and can be best illustrated by example.

2.2. Using Evolutionary Algorithms in Music

Evolutionary algorithms (EAs) have been applied in a wide range of musical contexts and applied at levels of the compositional process from the harmonisation of Bach Chorales to evolution of bebop improvisers, (see e.g. Burton and Vladimirova 1999b; Miranda and Biles 2007). Here we focus on a few specific examples in order to illustrate the key design issues. We first consider the evolution of synthesis algorithms, where evolution offers a practical solution to the problem of searching for new sounds in an often unintuitive space.

Take a software synthesiser such as Native Instrument’s *FM8*. It boasts well over 1,000 parameters, representing a vast, high-dimensional search space. The synth’s 960 named presets offer a means for users to intuitively access sounds, but this does not help in the discovery of novel sounds, which must be done by manual parameter tweaking. Although the default approach to controlling digital musical instruments, it is not necessarily efficient. Assuming a musically sensible design, evolutionary computation offers a practical alternative.

MutaSynth (Dahlstedt 2001) and *Synthbot* (Yee-King and Roth 2008) are two of many projects exploring the application of artificial evolution to syn-

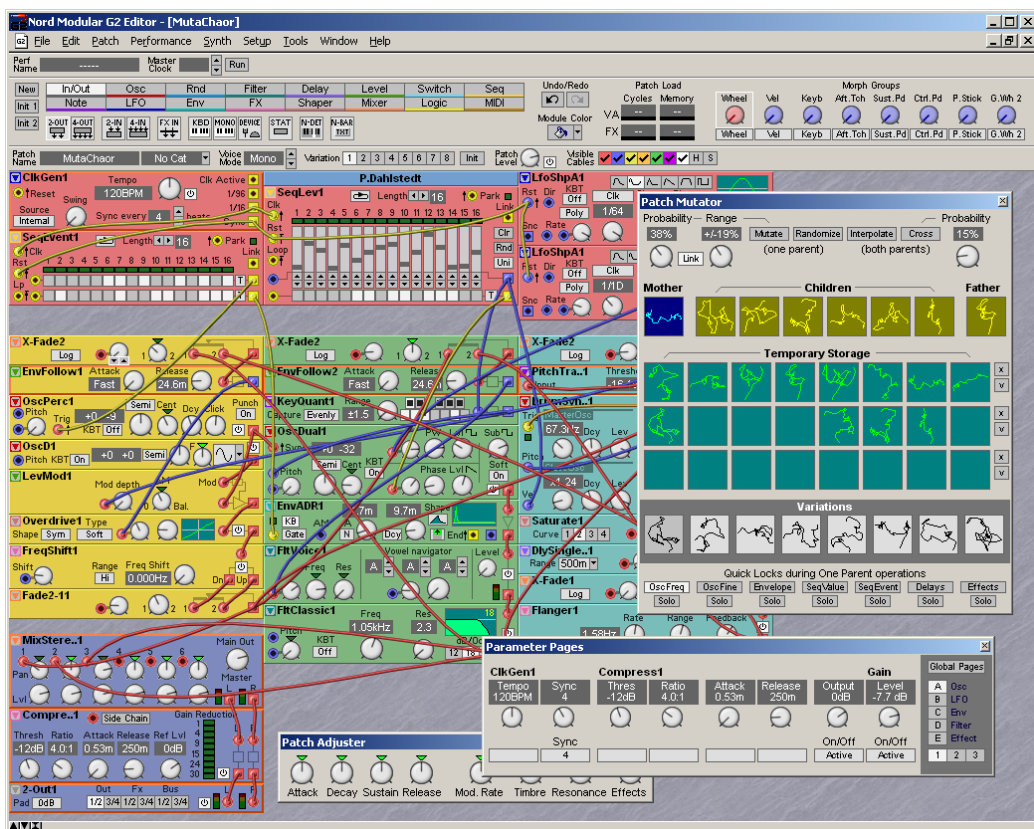


Figure 4: The Nord Modular editor with Patch Mutator on the right hand side

thesis. *MutaSynth* was designed as a general purpose tool for evolving programmable hardware synthesizers and later integrated into and distributed with Clavia’s popular Nord Modular G2 series of programmable hardware synths as *PatchMutator*, shown in Figure 4. Any of the control knobs available to the user can be encoded to form part of the genotype for the system’s evolutionary search. The sonic results of eight variations of the given parameter set are presented in a grid to the user, who auditions them, and picks their favourites. The system then creates further variations, which are presented to the user *ad infinitum*, until a desired sound is achieved.

Synthbot similarly enables an automatic search of a synthesiser parameter space, but rather than requiring feedback from the user, it is designed to search automatically for a match to a given target sound, which the user provides as input.

2.3. Approaches to GA Design

MutaSynth and *Synthbot* are used here to illustrate some of the issues around designing genetic representations, operators and fitness functions, and how these influence the fitness landscapes. Imagine, for example, designing a GA to search for harmonious sounds in the space of all possible Frequency Modulation (FM) synthesis parameters (Chowning and Bristow 1987).

Encoding schemes. The genome might include a representation of the synthesis graph, specifying the configuration of modulator and carrier oscillators, the frequencies and amplitudes of the component oscillators and the modulation depths. Just as carefully designing mappings from parameters to the knobs and sliders on a graphical user interface can increase usability, so applying some domain-specific knowledge in the design of the scheme can greatly enhance the power of search for a pre-specified task. The search for pleasing harmonic sounds, for example, could be expedited by encoding the modulation frequency *relative* to the carrier frequency, rather than as an absolute value (as we know that the ratio between these values dictates the harmonicity of the sound). A successful representation scheme then, will shape the fitness landscape in musically-meaningful ways. We can think of this in terms of stretching or magnifying areas of greater potential interest in the fitness landscape. In *MutaSynth*, for example, Dahlstedt uses non-linear (exponential, logarithmic or cubic) mappings from genotype to phenotype in order to make the most musically useful values more probable, while allowing for the possibility of more extreme values.

Representing the synthesis graph, rather than the parameters of fixed graphs, allows for the evolution of different configurations, as in Garcia's use of genetic programming (GP) (Garcia-Almanza and Tsang 2006). In GP (Koza 1992), another type of EA, the genotype takes the form of a mathematical function, which itself comprises a set of nested sub-functions, represented as a tree. Genotypes in this form can be mutated and crossed over in ways that grow or shrink the overall size of the function, and can result in dramatic new designs in a single operation. The rules for the genetic operations can be designed such that the resulting functions will always be mathematically valid. This provides a potentially powerful mechanism to evolve code, but is not without issues. One challenge is that pruning and grafting of GP subtrees can lead to dramatic discontinuities in the fitness landscape, obstructing intuitive evolution.

Genetic operators. Variation in the population is introduced by the genetic operators (mutation and cross-over). In our paper airplane example, we randomly modified the fold lines. Domain specific knowledge can be beneficial here too. In *MutaSynth*, these standard operators are mixed with a morphing process by which offspring can be created in a more intuitive way by interpolating genotype parameters between two parent genotypes, providing a continuous ‘cross-fade’ between two points in the phenotypic space of evolved sounds.

Fixed fitness functions. The design of fitness functions suitable for creative application of EAs has been the topic of much research in the field. *Synthbot* aims to evolve parameters of a synth to match a given target sound: a *fixed* similarity measure is provided by a perceptually motivated timbral measure, Mel-Frequency Cepstral Coefficients (MFCCs). This works well for optimisation tasks, however when a more general ‘aesthetic’ quality is desired, it can be difficult to formalise the desired target, despite efforts to measure aesthetic quality in music and visual art (e.g. Birkhoff 1933; Romero and Machado 2007).

Interactive genetic algorithms. Dawkins tantalised a generation of artists with the creative power of evolution in his simple *Biomorphs* program (Dawkins 1986), which he used to illustrate how easily human selection could lead to a diverse array of lifelike forms. This use of human aesthetic judgement in place of a formalised fitness function (as used in *MutaSynth* above) has been explored extensively in the application of EAs in visual art (Todd and Latham 1999; Sims 1991) and is known as an interactive genetic algorithm (IGA). As a stochastic search method, EAs typically require large populations to be explored over many generations. A key issue for this method then, is the amount of time required to perform the fitness judgments, described as the *fitness bottleneck* (Biles 1994). This is an issue for music in particular, as its fundamentally temporal nature obviates presentation of multiple individuals in parallel. User fatigue becomes a significant constraining factor, and various approaches to overcoming this have been explored.

Distributed IGAs. One natural response to the fitness bottleneck is to distribute the IGA amongst multiple users via the internet. Distributing the selection process places a lighter burden on a single user, but presents a challenge in organising how multiple users might work together to produce evolved outputs: differing preferences may result in serendipitously creative

outcomes or result in a directionless tug-of-war. The potential population size of candidate solutions also becomes overwhelming. This approach, which has been applied more extensively in the visual arts, has been used in Draves' *Electric Sheep* project (Draves 2005) (an animated screensaver) and Secretan's *PicBreeder* (Secretan et al. 2008), both of which show some promise for arriving at complex structures that wouldn't have come about through a traditional human design process.

Artificial critics. Another alternative is to either fixed or interactive fitness functions is train a machine learning system such as a neural network to perform fitness assignment. Early attempts suggested that this is merely deferring the problem: if a fitness function is hard to formalise for something as ineffable as a good jazz solo, then it will also be hard to train a machine learning system to take the place of that fitness function (Biles 1994). However, as machine learning achieves more impressive results this may still prove to be a fruitful approach in the future.

Coevolutionary approaches. Rather than using evolution as a *convergent* optimisation tool or interactive search mechanism, the power of evolution as a *divergent* engine for generating novelty and diversity has also been explored by many musicians. Taking further inspiration from the natural world, *coevolutionary* approaches have been explored in which populations of solutions are evaluated by populations of critics, which are themselves evolving. By virtue of this dynamic coupling, coevolution can also produce diversity within a population. Synchronic diversity can be generated through sexual selection leading to speciation – splitting the population into subpopulations of individuals with distinct traits and preferences (see e.g. Todd and Latham 1991). Coevolution can also amplify the diversity of novel forms over time, causing rapid evolution of traits as in predator-prey 'arms race' models (e.g. Futayama and Slatkin 1983).

Taking inspiration from a model of the evolution of bird song (Todd and Werner 1998), this coevolutionary approach has become a popular paradigm within computational creativity research, both for the generation of music and art, and as a modelling tool in the simulation of creative societies (Dahlstedt and Nordahl 2001). These interactions between critic and composer in the coevolutionary approach are seen as a proto-social behaviour (Miranda 2002a) and discussed in the context of multi-agent systems in Section 3.

Broadening the metaphor from the Darwinian evolution of isolated genotypes to the digital specification of entire ecosystems, the *computational*

ecosystems discussed in Section 5 develop the standard evolutionary algorithm by embedding the population in an environment where they interact with each other and with other environmental elements and spatial constraints. In these cases ‘fitness’ is no longer an explicitly pre-specified function as in standard EC, but defined *implicitly* in the interaction between individuals and their shared environment.

2.4. *Variations on a Theme*

The examples above aim to give a flavour of the different approaches to the design of EAs in musical applications. As is clear, many aspects of music making don’t fit neatly into an engineering optimisation framework and there are many interesting and promising cases where standard EAs have been ‘hacked’ for idiosyncratic aims. A common theme is the adaptation of classical EA components for the generation of diversity and variation. For example Waschka II and Magnus both take advantage of the structural changes which take place in the population through time (Waschka II 2001; Magnus 2006). Rather than employing EAs as a search mechanism and listening to the ‘winning’ individual at the final generation, the evolutionary process itself is sonified, conveying the changes that occur in the population across generations. Kiefer describes a GP-like variant designed for live performance situations that can be used to generate and interactively explore synthesis graphs on the fly (Kiefer 2014). Here again, no fitness function is specified but the representation scheme is adopted as a means for a user to rapidly search through a vast space of possibilities.

2.5. *Evolution and Usability*

This highlights the importance of human-computer interface (HCI) design aspects of evolutionary approaches. That IGAs present unique HCI issues has long been recognised (Todd and Latham 1999); novel interfaces which allow users to mix their own multi-objective fitness functions for building structures have been explored (Bentley and O’Reilly 2001). Dahlstedt’s *MutaSynth* has also has a strong user-centred focus, aiming to understand and tackle the common limitations of IGAs in functional end-user-software. For example, *MutaSynth* is designed to be operated with one hand so that a user can play a keyboard and constantly audition sounds whilst controlling the algorithm. Dahlstedt also experimented with the visualisation of genotypes, addressing issues of recall and rapid evaluation: “This visual representation was not a faithful representation of the actual sound. Rather, it was derived

from the parameter values of the synthesiser, which were used as length and angle values for a multi-segment line, scaled to fit the window. A small change to a parameter value would cause a small change to the visual representation.” (Dahlstedt 2001, p.90). The visual representation of sounds also aided users in recalling and organising the sounds they had discovered. Also with *MutaSynth*, since the ‘genes’ are also the synthesis parameters available to the user, the user can get in and tweak any evolved sounds, committing their modifications to the evolving population.

Through the commercial availability of his software, Dahlstedt was also able to gain user feedback. For example, users reported that the software supported working with increasingly complex patches where the relationship between individual synthesis parameters and good sounds became increasingly obscured. It has also been suggested that IGAs are most applicable in areas where the creator lacks either mastery or a strong conceptual model of the type of entity being created (Takagi 2001), which supports the use of EAs in synthesiser programming.

Many of these recent developments – and also much of the frustration with lack of progress, despite powerful tools – in the use of evolutionary computation in creative tasks, suggests that this focus on user workflows and basic interaction design, incorporating evolution in sensible ways into existing practice, is where fruitful advances can be achieved. The need for computational creativity to embrace interaction design has been emphasised recently (McDermott et al. 2013; Bown 2014).

3. Multi-Agent Compositions

A diverse group of researchers in mathematics, physics, and several branches of biology view self-organisation alongside natural selection as a complementary mechanism of evolution (Nicolis et al. 1977; Kauffman 1993; Camazine et al. 2001; Bak et al. 1988). Self-organisation refers to the process whereby an observed complex macro-level structure emerges from a series of local interactions between relatively simple agents: insect swarms, animal markings and even high level neural maps (Kaschube et al. 2010) and gross physical movement (Kelso 1995) arise through local interactions in the absence of top-down control. Self-organising phenomena have been extensively studied across disciplines such as ALife (Langton 1989) artificial chemistry (Dittrich et al. 2001), ecology, sociology and neuroscience using agent-based simulations. The aesthetic potential of such emergent behaviour has been

explored across many artistic domains, including sound and musical composition (Beyls et al. 2015; Miranda 1995; Blackwell 2003).

3.1. *Agent-Based Modelling*

Agent-based simulations are used to explain high level complex structures in terms of the interaction of a collection of simple agents in a shared environment and as such are a key tool for understanding self-organisation and other emergent phenomena. Important early work in this area includes Von Neumann’s description of self-replicating machines: devices capable of creating copies of themselves by precisely following a set of detailed instructions (Neumann 1966). These ideas were developed by Ulam, as a collection of cells on a grid, and evolved into what we now know of as cellular automata (CA). A CA can be conceived of as a regular N-dimensional grid of cells. In common with most agent-based modelling systems, a set of ‘agents’ are situated in a shared environment; each can exist in a finite number of internal states and act according to update rules which are in turn contingent upon the states of surrounding agents – their ‘neighbours’. For certain rule sets, an astonishingly complex range of global dynamics emerge from simple, local interactions. CAs have been explored in a wide variety of contexts to model phenomena across ecology (Hogeweg 1988), biology (Ermentrout and Edelstein-Keshet 1993) and sociology (Epstein 1996).

This dynamic complexity has also been extensively explored in compositional processes. Xenakis used CAs in the mid-1980s to create the “complex evolution of orchestral clusters” in *Horos* (Hoffmann 2002, p.122), and various computer scientists and composers have followed suit (e.g. Beyls 1991; Millen 1990; Miranda 1995; Burraston et al. 2004, for a historical and technical review).

Beyond the more abstract example of CAs, many models look at population behaviour in animals and humans. Thomas Schelling’s study of patterns of segregation in urban geography was one of the first agent-based models to have significant impact in the social sciences (Schelling 1971). Through the 1980s, agent-based models were developed across a wide variety of domains to explore game theory (Axelrod and Hamilton 1981), evolutionary processes (Hinton and Nowlan 1987) and as generative tools in computer graphics (Reynolds 1987). Reynolds’ *boids* model demonstrated that the phenomenon of flocking and swarming behaviour in bird, fish and insect species could be achieved in a simple computer model where each agent adjusted its

update vector in relation to near neighbours according to three rules: cohesion, alignment and separation (Reynolds 1987). Although making no claims to biological plausibility, the model demonstrates that coherent and robust flocking can arise through a process of self-organisation.

3.2. Emergence in Human and Artificial Music

Besides their more abstract pattern-generating properties, multi-agent models appeal to musicians through their potential to explore new artificial notions of ensemble behaviour. Self-organising processes had been explored by post-war experimental composers and pioneers of free improvisation. The text-score of Cornelius Cardew's *Paragraph V* of *The Great Learning* (Cardew 1971), for example, instructs singers to start on a random note and proceed each breath by picking a note sung by their neighbours. From a random beginning, a unified harmony emerges. The concept provokes algorithmic composers to this day: what are the timbral, harmonic, rhythmic and structural possibilities of self-organisation? How can humans and machines interact within such a framework?

Exploring such questions, Blackwell and Young (Blackwell 2003) used swarming as a model of free improvised ensemble performance. They considered a symbolic (MIDI) version in which particles swarm around a 3D (pitch, density, volume) space, and a granular audio version, in which particles swarm in a space of grain parameters.

The classic swarm algorithm is made interactive: the signal from an improvising musician is analysed and used to define points in the swarm space which act as attractors to the particles, creating a cycle of interaction between the human and the multiple swarm particles. The nature of the swarming algorithm means that the particles are loosely coupled in their movement: the resulting trajectories are likely to be similar, generally moving in parallel, providing a form of coordination that may result in counterpoint, repetition with interesting variation, canons or novel harmonisation. Swarm models have also been used in a more literal manner, using 3D audio to create the effect of a multiplicitous swarming of sounds around the listener (Kim-Boyle 2005).

In general, the self-organising dynamics of agent-based systems have been applied to good effect where the time-evolution of the system state is mapped in a sonically meaningful way. But how meaningful can such abstract systems be as models of human musicality? Does this kind of self-organising system capture aspects of human musical intelligence? On the one hand, various

creatively productive musicological accounts have focused on the dynamic surface structure of music (Toch 1977), metaphors from Newtonian physics such as gravity, inertia and momentum (Larson 1997) in describing how melodies play out, interact and are perceived. On the other hand, such models are seen to contribute only a limited amount to our understanding of human musical behaviour. Nevertheless, such experiments help explore the nature of music in an in-between space where the computer behaviour is neither that of a mere static object, nor of a sophisticated cognitive nature. In that sense, it may be reasonable to think of modelling musical behaviour without modelling human-level cognition, a topic discussed in Section 4.

3.3. Modelling Musical Communities

Another class of models looks at human interaction and self-organisation in social behaviour. Miranda’s models of musical interaction in virtual agents tackle both creative and empirical questions in an interesting synthesis of methodologies (Miranda 1999, 2000, 2002a,b). Following approaches in ALife and in the evolutionary models of language made famous by Kirby and Hurford (1997), Miranda presents agent-based modelling as a means to study the evolution of musical behaviour in human and artificial societies.

In language evolution, Kirby and Hurford challenged the expectation that we should seek adaptive functional explanations for the features of language, using a proof-of-concept model (e.g. Kirby and Hurford 1997). They showed that certain aspects of language structure could emerge from a process of iterative learning, i.e., through self-organisation, without the need to implicate evolved functions. Iterative learning is a process whereby one agent produces some form of output and another agent is presented with that output and learns its structure – typically using artificial unsupervised learning. The information content is just passed from one individual to the next – basically Chinese Whispers – and we can look at the long-term coevolution of content and learnt cognitive structure. Over successive iterations, whatever is most ‘copyable’ is what gets copied; the iterative learning shapes the structure and usage of the artificial language, leading Kirby and Hurford to comment that “it appears that languages adapt to aid their own survival over time” (Kirby and Hurford 1997, p.2). Miranda adapted such models to musical contexts (Miranda et al. 2003), and developed musical works using a similar scheme in which a population of agents first establishes a common vocabulary and then performs with it.

Models of human performance interaction have also been derived from agent-based models of *game theory* – the study of interactive decision making amongst agents. In some scenarios there is a clear mutual benefit to joint cooperation, but in countless real world cases there is some individual benefit to being selfish – assuming you can get away with it – and the greatest reward comes from being selfish whilst somehow ensuring that your co-player chooses to act cooperatively (often referred to as ‘freeloading’). This scenario is known as the Prisoner’s Dilemma, imagining two criminals who are being separately interrogated, with respective, mutually dependent payoffs for confessing or refusing to confess. The dilemma is that if both you and your co-player choose to be selfish, then you are both worse off than if you had both cooperated. This simple form of analysis has proven to have great power in understanding individual behavioural strategies, and has been successfully applied to question of the biological evolution of social behaviour (Maynard Smith 1982).

Didovsky developed some of the earliest multi-agent game works during the ascendancy of ALife. *Lottery* (Didkovsky 1992) presents a model society competing for access to a limited resource determined by a lottery process. Performers join this virtual society in playing the game. Harrald (2005) similarly uses game theoretic dynamics to establish rules for indeterminacy in the tradition of the 1950s New York School, specifically the work of Christian Wolff, using “unpredictable chains of performance situations that could arise only through the act of performance” (Harrald 2005, p.69).

The latter uses the *iterated* Prisoner’s Dilemma. If you have a history of interaction with your co-player it is possible to develop strategies for cooperation. Axelrod and Hamilton showed that a very simple strategy, called *tit-for-tat*, would suffice to establish mutual cooperation, without being prone to exploitation (Axelrod and Hamilton 1981): begin with cooperation, then simply copy what your co-player did in the previous turn. This became the major paradigm for understanding reciprocal altruism in ecology (e.g. Wilkinson 1984). Harrald’s composition ENSEMBLE establishes Prisoner’s Dilemma competitions between players. The players are programmed with different strategies, such as *tit-for-tat*, and the results of games at each round are used to dictate musical following behaviour by assigning different behaviours to defectors and cooperators.

Murray-Rust and colleagues created a system of intelligent musical agents that interact using ‘musical acts’ (Murray-Rust et al. 2006), a concept based on speech act theory (Searle 1969). This sets out to devise a sensible proto-

col by which agents can communicate amongst each other, opening up the possibility of allowing richer emergent and creative behaviour in multi-agent music systems, that might reflect emergent behaviour in humans. Murray-Rust used Terry Riley’s celebrated minimalist work, *In C* (1964) as a case study for how the agents could be put to task. The score consists of a series of short phrases; each player may proceed through each at their own pace. In this implementation the players are required to stay within 2 or 3 patterns of each other, must listen to each other and occasionally drop out, and must aim to merge into full unison at least once or twice during the piece. The application of the musical acts to this task demonstrated how a system of communication between multiple agents could be practically devised and structured with a specific outcome in mind.

Eigenfelt has produced a number of works (Eigenfeldt 2007a,b, 2008) based around the design of multi-agent systems where the interacting agents collaboratively develop compositional content by listening and responding to each other. The behaviour of agents evolves over time, with agents ‘reflecting’ on their behaviour according to a number of pre-programmed personality traits.

3.4. Agent-Based Models at Different Scales

As well as their application in virtual systems running on a single machine, agent-based software models also have relevance to how we approach performance with networked multi-computer systems and multiple performers. *The League of Automatic Music Composers* and later *The Hub* (Gresham-Lancaster 1998) are widely cited as pioneering the practice of networked computer music performance, and thus inherently initiating forms of experimentation into the modelling of multiple musical agents. *The Hub* experimented with audience participation over the internet as early as 1989, creating works in which multiple human users and algorithms acted on a shared memory space to produce musical output, and in doing so encountered questions of how to structure massive multi-user musical constructions.

A new project by Eigenfeldt et al. (2015) attempts to create a very general purpose distributed multi-agent architecture, the *Musebot Ensemble*, that allows developers to easily work together and attempt to exploit the emergent properties of musical interaction by running their distinct musical agents together in the same environment.

Multi-agent strategies need not only model the network interactions between individuals. In AI a number of theories and methods are based on a

model of cognition in which competing hypotheses battle it out for recognition, as determined collectively or by some central attention system (Minsky 1988). This has proven to be an effective strategy in music information retrieval as well, specifically in the case of beat tracking (e.g Dixon 2000; Large and Palmer 2002), based on the idea that rhythm perception uses multiple resonant oscillators and determines which resulting oscillation gives the strongest signal. This can lead to multiple harmonic rhythmic oscillations that may explain our metrical perception. Wiggins (2012) has proposed a multi-agent model of musical creativity which frames an understanding of music information dynamics and expectation in terms of massively parallel mental computation.

4. Adaptive Collaborators

“Only the environment can design a brain.” (Ashby Aphorism 151)

Understanding and modelling adaptive behaviour was of primary concern to cyberneticians – indeed for Ashby the ‘problem’ of life boiled down to the problem of adaptation – and is a critical concept in contemporary cognitive science and AI. In general, a system that is adaptive is able to provide appropriate responses to situations it encounters in its environment. This may be as simple as an autonomous vehicle correcting its path in response to an obstacle, a kitten learning to avoid a fire, or as profound as the emergence of a new species under evolution by natural selection. In contrast to the founding assumption of symbolic AI, numerous early cybernetic accounts and devices (introduced below) powerfully illustrated that the appearance of what looks like intelligent behaviour to an observer does not necessarily require sophisticated internal mechanisms. Such ideas inspired the development of ALife approaches to autonomous robotics and continue to influence contemporary philosophy of life and mind (Di Paolo 2003; Froese et al. 2010).

Some of these core principles were illustrated by the *Machina Specularix*, built by neurophysiologist and robotician William Grey Walter (Walter 1950). Walter built two electro-mechanical ‘tortoise’ robots, named Elma and Elsie, in which the output of a pair of light sensors controlled the robot’s wheel motors. Putting two robots together (or in front of a mirror), each with a light on top, he achieved phototaxis (light-seeking behaviour), but also compelling ‘lifelike’ dances. From such work we have the seeds of the later

formalised principles of *situatedness* and *embodiment*. Situatedness refers to the fact that the relevant behaviour is only apparent once the agent is placed in the relevant, rich environmental context. Embodiment refers to the fact that the behaviour is contingent upon a physicality: a physical body that has myriad interactive affordances beyond those which may have been consciously implemented by the designer. The term ‘lifelike’ has been used widely in the ALife literature to express the notion that both organic and artificial systems may, specifically through their adaptive behaviour, take on the appearance of a living system, inviting ‘attributions of intentionality’.

Biologically-inspired adaptive systems provide opportunities to explore novel modes of interaction and decision-making in a musical context. Much of the work in intelligent music systems is implicitly geared towards modelling human musical behaviour. But as we have seen in other examples above, computers allow us to explore musical interaction in new ways, creating novel and possibly hybrid interaction scenarios that weave in exotic computational behaviours. Computer music makers have incorporated simple or complex adaptive behaviours into their software systems, drawing on a variety of models, and applied to a variety of compositional goals. These scenarios place the performing musician in a new relationship with the musical dynamics where they are not necessarily in direct control of the outcomes, but act as negotiators with the machine. The process becomes one of mutual adaptation.

4.1. Adaptive Behaviour and Musical Interaction

For musicians, the idea that simple, situated, biologically primitive behaviours might underlie complex pattern making in music is compelling. Some musicians working in this area have taken a very much ALife inspired view of music itself as a ‘dynamical complex of interacting situated embodied behaviours’ (Impett 2001). From this perspective, there is arguably greater potential for human and machine to take on equivalent roles in co-creation. Each contributes to a shared collaborative musical environment which in turn affords musical opportunities for the other.

This conversational model of musical interaction was explored in the 1950s by cybernetician Gordon Pask in a quirky experimental analogue audio-visual improvisation system *Musicolour* (see e.g. Pask 1971, for an overview). The system listened to a performing musician and responded with patterns of coloured lights. But by Pask’s design, the system would become bored if the input was always the same, and would adapt its behaviour to provoke a new

response. The system was not particularly sophisticated in its behaviour, and had limited scope, but began to chart how we might design interactive music systems that maintained a compelling engagement with a musician; players touring Music Halls with it in the 1950s reportedly engaged with it much like another musician (Haque 2007).

4.2. Behavioural Objects

In earlier work we have explored a number of biologically-inspired models which display adaptive behaviour, including various neural models and homeostasis. We proposed the term *behavioural object* to refer to the ways in which software can act as the focus of interaction. This extends the traditional model of interaction inherited from acoustic instrumental performance, which focuses solely on musical interaction at performance time (see Bown et al. 2009, for details) to include other types of interaction: the social interaction between communities of developers; the interaction between developer and software; and the musically significant interactions between software elements themselves, incorporating all of the biologically-inspired forms of inter-agent interaction discussed in this chapter. To frame these distinct forms of interaction, we distinguished two senses in which a behavioural object could have agency: performative agency (in performance time) and memetic agency (out of performance time, i.e., spanning multiple performances or acts of musical production). Performative agency refers to the ability of a software system to influence the outcome of a specific musical performance. For those primarily interested in human-computer collaborative improvisation systems, performative agency is directly synonymous with the quality of musical interaction it affords.

Eldridge (2005) explored an Ashbian model of homeostasis as a core organising mechanism for electro-acoustic improvisation. Ashby addressed a fundamental conundrum: how can a system (biological or mechanical) be at once state determined, and yet adapt to a changing environment and learn? (Ashby 1960). He proposed that the key mechanism underlying adaptive behaviour is homeostasis – the maintenance of key internal variables in the face of external perturbation. As a good cybernetician he built a physical device to demonstrate his theoretical notion of *ultrastability*: the homeostat. This was an electro-mechanical device, but the critical elements can be described in the abstract and consist of: a *physical system* which is capable of interaction with its *environment* such that it perturbs the value of *an essential variable* with a specified boundary of viability (e.g. body temperature, blood

pressure etc.). When this boundary is exceeded a *selector* is triggered that specifies a break, or change, in the system (e.g. starting to shiver when cold or sweat when hot) such that it *randomly* changes its organisation until a new set of parameters is arrived at under which the essential variable is brought back within limits. This conceptualisation of homeostasis and ultrastability has influenced advances in autonomous robotics and contemporary theories of cognition (Di Paolo 2003; Maturana and Varela 1987; Froese et al. 2010).

The device was simulated in a neural-network style model of interconnected nodes which displayed the key homeostatic and learning behaviour demonstrated by Ashby’s machine: from an initially random state, oscillatory dynamics emerge. Small perturbations cause temporary disturbance followed by a return to the initial state; larger perturbations push the system into a new trajectory where it settles into a different cyclic attractor. In Ashby’s original conception, the essential variables represented the nervous system in interaction with the environment; in an electro-acoustic performance setting, the model provides a conceptual vehicle and algorithmic means for collaborative human-computer interaction. The model was implemented in an audio-visual improvisation system where the homeostat received input from a visual display, the essential variables being used to ‘remix’ samples taken live from an acoustic improviser which in turn generated visuals, closing the loop. The homeostat acted to recompose earlier musical elements, influencing the performer’s subsequent improvisations to create an electro-acoustic dance-like experience (Eldridge 2005).

Bown took inspiration from the minimal cognition work of Beer (e.g. Beer 1997), in which generic network architectures were evolved from first principles to learn simple cognitive tasks in an elementary virtual world. The dynamics of continuous-time recurrent neural networks (CTRNNs) were evolved with fitness functions defined in highly abstract musical terms (Bown and Lexer 2006). With creative search rather than musical competency in mind, these fitness functions were designed to steer the evolved structures towards some degree of complexity and responsiveness of behaviour, rather than specific musical outcomes of rhythm and harmony. Fitness properties consisted of notions such as ‘always output cyclical patterns’, or ‘if the input changes, the output should change’ – compositional instructions which may or may not be evolve literally. Bown showed that it was relatively simple to evolve rich, complex, interactive behaviours that could then be applied in a variety of generative or interactive musical contexts. He also introduced a music-specific variation to the CTRNN algorithm, in the form of a sinusoidal

transfer function, that enabled individual nodes in the network to oscillate on their own, boosting the overall oscillatory behaviour of the network.

This work was followed up by Zamyatin (Bown 2011), an improvising agent with a novel fitness function in which a decision tree, which exhibits feedback by influencing an internal state array, in turn influences the decision tree's decisions. The decision tree was shown to have a number of traits that made it more practical in musical use than the CTRNN. For example, it outputted discrete rather than continuous data, which proved to be more convenient for mapping to event-based musical control.

4.3. Subsumption Architectures for Musical Agents

A significant strand of work approaches the study of musical interaction through the development of software performance systems, explicitly adopting the design philosophy of roboticist Rodney Brooks.

The 1980s saw a resurgence of interest in cybernetics which fuelled a new approach to autonomous robotics. Brooks was one of many who eschewed the representationalist approach of AI – or Good Old Fashioned AI (GOF AI) as it was dubbed – arguing that AI was too focused on advanced cognitive behaviour (such as the logic of playing chess) rather than fundamental intelligence, or minimal cognition (such as picking up and moving the chess pieces). Under the dictum “the world is it's own best model” (Brooks 1999, p.115), Brooks developed a subsumption architecture approach to robotics in which real time sensory information is coupled to action selection in an intimate, bottom-up fashion, rather than being guided by symbolic mental representations of the world (Brooks 1991). This approach, Brooks argued, enables agents to respond quickly and appropriately to changes in an unpredictable world, a task which is fundamental in the real world but which challenged the task-specific GOF AI robots of the time.

Brooks' approach was parsimonious and incremental, following observed biological behaviours. Starting with the basic behaviours (move forward, avoid obstacle) he built robots comprising sensors, motors and the simplest possible 'brain' (designed by hand). Once debugged and tested in the real world, additional behavioural layers were added, mimicking the phylogeny of real creatures. Layers worked in parallel and generated outputs which might be signals to inhibit or suppress other layers or issue commands to actuators. This deviated from the more common top-down designs produced on GOF AI principles, aligning instead with the emerging evolutionary perspective

of cognition as compartmentalised into efficient but narrow domain-specific competencies (e.g. Barkow et al. 1992).

Bryson (1995) was the first to apply such techniques to music performance, which she also treated as an empirical investigation into what level of cognitive sophistication music operates at (many of the systems described here implicitly pose this question, but Bryson explicitly noted the epistemological potential of this approach). Her reactive accompanist modelled the capacity to derive chord structure from a melody (Bryson 1995, p.2) and followed Brooks' subsumption architecture principle. Bryson showed that the reactive accompanist could perform to a relative degree of sophistication using this principle, demonstrating the efficacy of subsumption and action selection as part of a software design method in creative computer music.

The subsumption approach to handling input and output is highly suggestive of the tight coupling between listening and playing (input and output) of free improvisers, who exhibit a robust, flexible approach to dealing with unpredictable changes in the sonic environment (Clarke 2005; Sudnow 1978). Inspired by this observation, Linson developed an artificial improvising agent based on subsumption architecture, *Odessa*, in order to carry out research into the complex dynamics of musical improvisation as a 'situated psychosocial and embodied cognitive practice' (Linson et al. 2015). His design follows Brooksian parsimony at every level. Just three competing behaviours are implemented: the ability to spontaneously produce output – 'Play' – to respond to musical input – 'Adapt' – and to disregard input, introducing silence and initiate endings – 'Diverge'. Within these, the simplest conceivable algorithms are deployed. Despite this, Linson demonstrated that complex interactive behaviour, subject to evaluation by experts who improvised with it, can emerge from interactions between modules, supporting the idea that in-the-moment inferences, based on behavioural cues, perceived in real time, can lead to the attribution of intentional agency in musical machines. Linson's work addresses our understanding agency and autonomy in interactive computer music, a core theme which is revisited at the close of this chapter.

5. Creative Ecosystems

The term ecosystem was first used in Arthur George Tansley's paper on vegetational concepts, to pronounce his conviction that organisms cannot, fundamentally, be considered in separation from "the environment of the

biome – the habitat factors in the widest sense ... with which they form one physical system” (Tansley 1935, p.299). Tansley wished to comprehend not only communities of organisms, but also to bring the complex interactions of biotic and abiotic factors surrounding them into focus. By giving a name to this tightly coupled collection of biotic and abiotic organisms and processes, he sought to establish the ecosystem as “a recognisable self-contained entity” (Tansley 1935, p.228).

The concept has gained popularity in recent years as a metaphorical cornerstone in a range of musical contexts: as a framework for understanding the influences and elements of *performance ecosystems* in theory and practice (Bowers 2002; Waters 2007); as an approach to composition based on acoustic and adaptive systemic feedback in *audible ecosystems* (Di Scipio 2003); and as a vehicle for creative discovery, in *computationally creative ecosystems* (McCormack 2012). Although differing in basic materials, all three approaches take Tansley’s original conception as a central organising principle and theorise or construct situations in which organisms (human musicians in the world, agents in simulation, or some electro-acoustic hybrid) and their environments are irrevocably coupled via feedback processes, such that any musical agency of the system can only sensibly be understood as originating from the system as a whole. Systemically, concern shifts from understanding embodied, situated agent behaviours *in* an environment, to conceiving of the environment and organisms within it as a super-entity.

5.1. Audible Ecosystems

Di Scipio’s *Audible Eco-Systemic Interfaces*, is a series of works which integrates aesthetic, philosophical and technical aspects of ecosystemic thinking. Technically, Di Scipio creates software components that monitor, adapt to and transform their ambient acoustic environment. For example a signal processing module might be set up to automatically alter its internal configuration according to changes in the input sound, in turn altering the acoustic environment (see Figure 5). Algorithmically the idea is fairly simple and could be described as ‘adaptive audio feedback’. But in doing so, sound itself determines the conditions and boundaries for its own transformations. This has a practical creative value as the real world has a dynamic richness that is hard to reproduce in software; conceptually it establishes the same conditions in which the human agent is situated in a music performance context.

Philosophically, his approach references autopoiesis (literally ‘self’ ‘creation’). The term was coined in the 1960s by the biologists Humberto Maturana and Francisco Varela (see Varela et al. 1974, for the first English publication) to convey their conviction that the root of biological autonomy lies in the circular organisation of living organisms, a feature they saw to be both necessary and sufficient for life. The use of sonically situated adaptive processes can be seen as a *structural coupling* of software and environment. This is a critical construct in autopoietic theory used to describe a process of engagement which effects a “... history of recurrent interactions leading to the structural congruence between two (or more) systems” (Maturana and Varela 1987, p.75). Di Scipio’s systems cannot be seen to literally be ‘alive’, but reference to autopoietic theory promotes a fresh perspective on agency and hence interaction in a systemic sonic practice and leads to a productive framing of (non-biological) autonomy in terms of self-determination. This breed of ‘self-determined’ acoustic ecosystem has grown into a genre in its own right within electroacoustic composition (Eldridge 2008; Bown 2009; Green 2011; Eldridge 2013; Sanfilippo 2013).

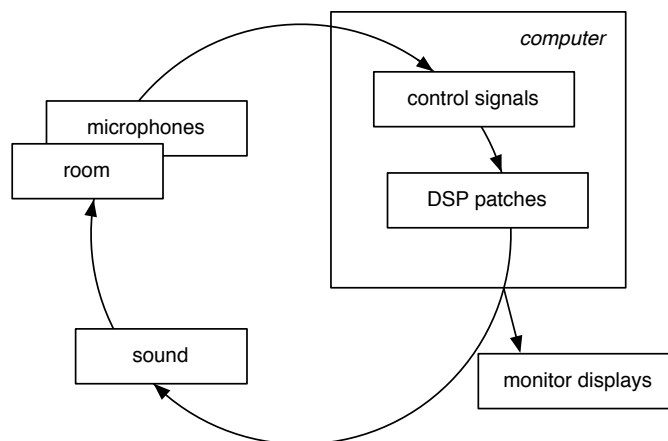


Figure 5: Basic design of Di Scipio’s Audible Eco-systemic Interfaces

As with other practitioners described in this section, Di Scipio takes as his starting point a critique of the notion that ‘basic’ interactive music systems place the autonomy of the performer at their centre: “agent acts, computer re-acts ... The only source of dynamical behaviour lies in the performer’s ears and mind” (Di Scipio 2003, pp. 270). By adopting an ecosystemic approach, where the software system can directly respond to and alter its environment,

the works demand consideration of the different notions of ‘interaction’ as applied to human performers and software components differentially and in combination.

A number of other cybernetically-minded practitioners at around the same time shared similar sentiments in critiquing interactive art and creative systems across sonic, visual and kinematic arts including Simon Penny, Usman Hacque, Bert Bongers and Jon McCormack. The term ‘interaction’ becomes a battleground, upon which a potentially richer and more organic set of interactions with machines is at stake, compared to the standard material of HCI. Pask’s conversation model of interaction (outlined above in Section 4.1), for example, continues to be cited as a goal for what constitutes meaningful interaction.

Other ecosystemic approaches are more focused on the construction of narrative than on emergent complexity. Billed as “a musical composition that grows in the same way as a forest ecosystem”, *Living Symphonies*, is a recent work by Jones and Bulley which toured UK forests in 2014. The work comprises a spatial multi-agent system where agents represent organisms of different species of flora and fauna – beetles, birds, mammals, fungus and other plants. The population densities of species are parameterised by data collected from each site, reflecting the local biodiversity and community structure and further influenced by local weather conditions (wind speed and direction, sunlight, rain fall) streamed from an onsite weather station. Different species are sonified by distinct musical motifs, composed from fragments of acoustic instrumental recordings and generated according to the current simulation state. The simulation runs at quasi real time such that the movement of agents across the simulated world are mapped across the real world, heard as movement across speakers. The audience experience the real-time dynamics of a simulated ecosystem which portrays the changing activity of the forest.

5.2. Computational Ecosystems for Creative Discovery

The prospect that an ecosystemic approach to algorithm design and implementation might afford a richer, creatively productive means of interaction motivates a revised view of human-machine co-creativity, described by McCormack in the language of computational creative discovery: ‘where computer processes assist in enhancing human creativity or may autonomously exhibit creative behaviour independently’ (McCormack 2007, p.1). The intention is to develop ways of working with technology that achieve creative

possibilities unattainable from any existing (software) tools or methods. McCormack has looked at extending the EAs and agent-based systems described above, incorporating adaptive behaviour at the ecosystem level. Thus rather than thinking of evolution as applying to the optimisation of a specific task (or satisfaction of a specific human preference goal) as in many of the approaches outlined in Section 2 it is treated as more of a free-for-all in which self-organisation may also play a significant role in the process, evoking Darwin’s image of the “tangled bank, clothed with many plants of many kinds” (Darwin 1861, p.524).

In works such as the audio-visual installation *Eden*, McCormack develops upon earlier adaptive agent-based models (Holland 1999) and ALife art which deployed IGAs in exhibition contexts (e.g. Sommerer and Mignonneau 1997) to create an evolutionary agent-based system in which physical presence of audience members (monitored via IR cameras) becomes a resource for a population of virtual agents, whose behaviour then seeks to best maintain their attention (McCormack 2001). Later, McCormack and colleagues engaged in a more systematic study of ecosystemic approaches to algorithmic art. The standard EA is extended from modelling isolated genotypes to the digital specification of an environment and agent’s interactions with it and each other; rather than explicitly or interactively defining a fitness function, evolutionary pressure is exerted via competition for an abstract resource, which may in turn alter the environmental structure. Within this framework, various evolutionary processes and principles were explored, including resource recycling, mutualism, evolution of generalism and specialism and niche construction.

That heterogeneity of resources in an environment can give rise to complex agent behaviour at the population level was established in early agent based models (Epstein 1996). Other evolutionary models consider what conditions are required to establish and sustain species diversity and inter-species interaction. In musical applications we have seen one example already in the competitive coevolution of ‘singers’ and ‘listeners’ (see Section 2.3). Organism interactions are not limited to direct competition however: the myriad of symbiotic relationships observed in the natural world provide rich inspiration for achieving synchronic diversity in creative system design. The sonic ecosystem *Filterscape* (Eldridge and Dorin 2009) was built to explore the conditions under which energy recycling can lead to the emergence of cooperative as well as competitive survival strategies, increasing the *behavioural* diversity of a generative system as well as environmental heterogeneity.

The ecological concept of niche was also explored as a means to investigate the evolution of generalism and specialism. A niche describes the role of an organism: the ways in which it interacts with and depends upon other elements in its environment (Elton 1927). The structure of individual niches and the relationships between them provide a means to understand competition and analyse the species composition of a community, as well as its stability. The *width* of a niche describes the degree of specialisation of an individual or species: a specialist has a narrow niche, it occupies a limited and particular habitat or set of activities; a generalist occupies a broad niche and can make use of a wider diversity of actions and habitats. Niche overlap provides a measure of competition between species. Drawing from early work in theoretical ecology, Eldridge and Dorin demonstrated that the trajectory of an evolving population through phenotype space could be controlled by explicitly manipulating the width of the niche, biasing the evolution of generalist or specialist survival strategies (Eldridge et al. 2008). This provided a means to support the generation of complexity, whilst providing a simple, intuitive ‘handle’ with which a user could exert some degree of influence over an otherwise intractable complex system.

McCormack and Bown subsequently implemented a form of niche *construction* – the process whereby organisms, through their activities and choices, modify their own and each other’s niches – as a means to increase environmental heterogeneity and so more complex structures in music and art. For example, the line-drawing agents of Annunziato and Pierucci (2000) were adapted to include a niche construction element within an evolutionary framework (McCormack and Bown 2009). Agents draw different types of lines (e.g. varying curvature or density) and have a preference for local environments which contain certain line structures. Line drawing tendencies and preferences are both encoded genetically. Coevolution of tendency and preference in localised areas acts to create local variation and so greater global diversity. Bown and McCormack also developed various sonic ecosystem models that made use of the acoustic environment as a niche, following Di Scipio, that could both be modified and have an influence on the evolution of agents (Bown and McCormack 2010). They then looked at general principles for designing creative ecosystems, including design-focused principles such as maximising the amount of knowledge you have about the system, and more specific issues such as understanding when and how boundaries might emerge in spatial networks of agents (Bown et al. 2011).

The experiments begin to show how ecological principle can be applied

in the development of software for creative discovery to increase complexity of structure or behaviour, whilst allowing some degree of user influence, in some cases. Like many of the projects described here, such work offers new perspectives on how we conceptualise creative collaboration as we increase the generative autonomy of software.

6. Looking Forward

The principles of cybernetics and ALife established during the 20th century continue to feed into experimental practice in algorithmic music, which is itself evolving new tools and methods. But a considerable challenge remains – and this is as much a challenge of HCI as of algorithm design – to find effective and usable ways to exploit these principles in compositional practice. Biologically-inspired computational creativity has found its way into commercial software in a small number of cases, but has not yet come of age.

Likewise the vision of virtual worlds autonomously spawning the kinds of rich complexity seen in natural systems remains an elusive goal in ALife research (Bedau et al. 2000). This may simply be because we have not arrived at the order of magnitude required, because there are still missing elements to our algorithmic toolkit, or perhaps because there is a more fundamental difference in kind between digital computers and the biological world. Some argue that if it is possible in nature it should be possible in virtual worlds, others assert that computation has nothing to do with what goes on in brains and other biological systems (Harvey 1997).

But although inspired by the sciences of the artificial, biologically-inspired algorithmic composition does not share the same aims and is not hampered by what some see as disappointments. The snapshot of practice we have given in this chapter is symptomatic of the creative fluidity found throughout algorithmic music: the algorithmic musician is seen weaving together equations from complex systems, biology, social behaviour and so on, to create new artistic works that employ nature in their behaviour, rather than merely representing nature. The changing technological environment also seems to be shifting towards one in which biologically-inspired and agent-based approaches may have new grounds for making relevant contributions; massively distributed computing networks, increasingly high-level APIs, the widespread scriptability of advanced software platforms, the connectivity of web-programming, and the rich ecosystems of the Internet of Things, all po-

tentially offer new creative environments for the application of the methods described in this chapter.

As this practice progresses, we have argued, there is no need to remain true to biology, social behaviour or human musicality. Indeed, perhaps the greatest value of a biologically-inspired approach is that it offers a radically different starting point for the coevolution of new forms of music making. In this sense we see two key themes for future exploration: a further advancement in understanding the role of software in the co-creation of music making, through the concepts of autonomy and agency, and the value of biological metaphors in supporting the creative development and cultural sharing of computer music.

Autonomy and Agency. Autonomy can be usefully understood as the ability of a system to control its own future state, rather than be controlled by external factors. This has proven challenging to measure, and it is arguably the case that qualitative, narrative approaches to understanding autonomy may be especially relevant in the application area of music. The most common narratives surrounding autonomy relate to two primary forms of adaptation: evolution and learning. Although it is hard to witness the autonomy of the slowest of organisms, including plants and trees, we understand that species are autonomous by virtue of their evolutionary inception; they need no help in surviving. Likewise, learning creates a situation in which the actions of an individual may be the result of an inductive analysis of their environment. We might therefore appreciate that learning is one way for a machine to acquire behaviour which has not been programmed directly, and recent advances in deep learning have driven this notion home to a wider public (Devlin 2015).

Closely related to autonomy, and more central to questions of creativity, is the notion of agency: the ability to have influence on the world, as in the production of an artwork. How and when we attribute creative agency to software systems is a question being asked by many of the researchers mentioned in this field. Autonomy and agency may not take the obvious forms associated with the simulation of human behaviour, and an important challenge is to come to terms with broader notions of creative software agency, for which biological and social science understandings of action and interaction will inevitably act as fertile sources of ideas, as well as philosophical accounts of intentionality (Dennett 1987) and expressivity (Linson 2013).

The poetics of biologically-inspired algorithms. Algorithmic music in general utilises a plethora of computer models derived from across the spectrum of scientific disciplines, but in many ways biological models are the most exotic and evocative, posing conceptual challenges for how we think about music, and linking the most hypermodern of activities to the most primitive. Biologically-inspired models offer an alluring narrative for audiences and computer music makers. Beside epistemological, technical and creative opportunities then, biologically-inspired models hold potential as vehicles for both cognitive and cultural engagement.

In the forward to *Visual Complexity* Lev Manovich highlights data art as a new medium for critical reflection on the world: “Figurative artists express their opinions about the world by choosing what they paint . . . Now artists can also talk about our world by choosing which data to visualize.” (Lima 2006). In an analogous way, the choice of model or system design in algorithmic composition becomes a vehicle for expression and comment – poetic, philosophical or political – about our world and our relationship to it (Eldridge 2012). In offering a familiar narrative frame, biological models and metaphors support engagement for both creator and audience. During the design and development of new works, such metaphors support rich ‘system stories’ (Whitelaw 2005) providing a cognitive scaffold for the coder. Similarly, for the audience, familiar narratives provide a ‘way in’ to algorithmic music, which can otherwise be less-than approachable to wider audiences (Garnett 2001; Stubbs 2009). At a time when technology develops as fast as our natural environment is threatened, biologically-inspired models offer a valuable vehicle for reflecting upon the relationships between the biological, cultural and technological worlds in which we live, and for sharing complex concepts and aesthetics with audiences.

Annunziato, M. and P. Pierucci (2000, August). Emergent relationships. In *Proceedings of the Alife VII*, Portland, Oregon.

Ascott, R. (1967). Behaviourist art and the cybernetic vision. *Cybernetica, Journal of the International Association for Cybernetics* 10(4), 25–56.

Ashby, W. (1956). *An Introduction to Cybernetics*. Chapman & Hall Ltd., London.

Ashby, W. R. Ashby’s card index of aphorisms. From Ashby’s notebooks. Available online at <http://rossashby.info/aphorisms.html>.

- Ashby, W. R. (1960). *Design for a Brain*. John Wiley & Sons. inc.
- Axelrod, R. and W. D. Hamilton (1981). The evolution of cooperation. *Science* 211(4489), 1390–1396.
- Bak, P., C. Tang, and K. Wiesenfeld (1988). Self-organized criticality. *Physical review A* 38(1), 364.
- Barkow, J. H., L. Cosmides, and J. Tooby (1992). *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. New York: OUP.
- Bedau, M., J. McCaskill, N. Packard, S. Rasmussen, C. Adami, D. Green, T. Ikegami, K. Kaneko, and T. Ray (2000). Open problems in artificial life. *Artificial Life* 6, 363–376.
- Beer, R. D. (1997). The dynamics of adaptive behavior: A research program. *Robotics and Autonomous Systems* 20, 257–289.
- Bentley, P. J. and U. O’Reilly (2001). Ten steps to make a perfect creative evolutionary design system. In *GECCO 2001 Workshop on Non-Routine Design with Evolutionary Systems*.
- Beyls, P. (1991). Chaos and creativity: The dynamic systems approach to musical composition. *Leonardo Music Journal* 1(1), 31–36.
- Beyls, P., G. Bernardes, and M. Caetano (2015). The emergence of complex behavior as an organizational paradigm for concatenative sound synthesis. In *Proceedings of the third conference on Computation, Communication, Aesthetics and X*, Glasgow, UK.
- Biles, J. A. (1994). GenJam: a genetic algorithm for generating jazz solos. In *Proceedings of the 1994 International Computer Music Conference*, pp. 131–137. San Francisco: ICMA.
- Birkhoff, G. D. (1933). *Aesthetic measure*. Cambridge, Mass.
- Blackwell, T. (2003). Swarm music: Improvised music with multi-swarms. In *Proceedings of the AISB Symposium on Artificial Intelligence and Creativity in Arts and Science*, pp. 41–49.

- Bowers, J. (2002). Improvising machines: Ethnographically informed design for improvised electro-acoustic music. Master's thesis, School of Music, University of East Anglia, Norwich, UK.
- Bown, O. (2009). A framework for eco-system-based generative music. In *Proceedings of the SMC*, pp. 195–200.
- Bown, O. (2011). Experiments in modular design for the creative composition of live algorithms. *Computer Music Journal* 35(3), 73–85.
- Bown, O. (2014). Empirically grounding the evaluation of creative systems: incorporating interaction design. In *Proceedings of the Fifth International Conference on Computational Creativity*.
- Bown, O., A. Eldridge, and J. McCormack (2009). Understanding interaction in contemporary digital music: from instruments to behavioural objects. *Organised Sound* 14(02), 188–196.
- Bown, O. and S. Lexer (2006). Continuous-time recurrent neural networks for generative and interactive musical performance. In F. Rothlauf and J. Branke (Eds.), *Applications of Evolutionary Computing, EvoWorkshops 2006 Proceedings*.
- Bown, O. and J. McCormack (2010). Taming nature: tapping the creative potential of ecosystem models in the arts. *Digital Creativity* 21(4), 215–231.
- Bown, O., J. McCormack, and T. Kowaliw (2011). Ecosystemic methods for creative domains: Niche construction and boundary formation. In *Symposium on Artificial Life*, pp. 132–139. IEEE.
- Brooks, R. A. (1991). How to build complete creatures rather than isolated cognitive simulators. In K. VanLehn (Ed.), *Architectures for intelligence: Carnegie Mellon Symposia on Cognition*, pp. 225–239. L. Erlbaum Associates.
- Brooks, R. A. (1999). *Cambrian Intelligence: The Early History of the New AI*. MIT Press Cambridge, MA.
- Bryson, J. (1995). The reactive accompanist: Adaptation and behavior decomposition in a music system. In *The Biology and Technology of Intelligent Autonomous Agents*, pp. 365–376. Springer.

- Burnham, J. (1968). *Beyond Modern Sculpture: The Effects of Science and Technology on the Sculpture of this Century*. G. Braziller.
- Burraston, D., E. Edmonds, D. Livingstone, and E. Miranda (2004). Cellular automata in MIDI based computer music. In *Proceedings of the 2004 International Computer Music Conference*.
- Burton, A. R. and T. Vladimirova (1999a). Generation of musical sequences with genetic techniques. *Computer Music Journal* 23(4), 59–73.
- Burton, D. and T. Vladimirova (1999b). Generation of musical sequences with genetic technique. *Computer Music Journal* 23(4), 59–73.
- Camazine, S., J.-L. Deneubourg, N. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau (2001). *Self-Organization in Biological Systems*. Princeton University Press.
- Cardew, C. (1971). The Great Learning Paragraphs 2 and 7. Deutsche Grammophon/Universal Classics 471 572.
- Chowning, J. and D. Bristow (1987). *FM Theory and Applications: By Musicians for Musicians*. Hal Leonard Corp.
- Clarke, E. F. (2005). *Ways of listening: An ecological approach to the perception of musical meaning*. Oxford University Press.
- Collins, R. J. and D. R. Jefferson (1991). Selection in massively parallel genetic algorithms. In *Proceedings of the Fourth International Conference on Genetic Algorithms*, pp. 249–256. Morgan Kaufmann.
- Cope, D. (1996). *Experiments in Musical Intelligence*. Madison, WI: A-R Editions.
- Dahlstedt, P. (2001). A mutasynth in parameter space: interactive composition through evolution. *Organised Sound* 6(02), 121–124.
- Dahlstedt, P. and M. G. Nordahl (2001). Living melodies: Coevolution of sonic communication. *Leonardo* 34(3), 243–248.
- Danforth, C. M. (2013). Chaos in an atmosphere hanging on a wall. Accessed: 2015-08-20.

- Darwin, C. (1861). *On the Origin of Species by Means of Natural Selection, or the preservation of favoured races in the struggle for life*. Murray.
- Dawkins, R. (1986). *The Blind Watchmaker: Why the Evidence of Evolution Reveals a Universe Without Design*. London: Penguin.
- DeLanda, M. (2002). Deleuze and the use of the genetic algorithm in architecture. Available online at <http://www.cddc.vt.edu/host/delanda/pages/algorithm.htm>. Accessed 20.08.2015).
- Dennett, D. (1987). *The Intentional Stance*. Cambridge: MIT Press.
- Devlin, H. (2015, February). Google develops computer program capable of learning tasks independently. *The Guardian*.
- Di Paolo, E. A. (2003). Organismically-inspired robotics: homeostatic adaptation and teleology beyond the closed sensorimotor loop. In K. Murase and T. Asakura (Eds.), *Dynamical systems approach to embodiment and sociality*, pp. 19–42. Adelaide: Advanced Knowledge International.
- Di Scipio, A. (2003). Sound is the interface: From interactive to ecosystemic signal processing. *Organised Sound* 8(3), 269–277.
- Didkovsky, N. (1992). “Lottery”: Toward a unified rational strategy for cooperative music-making. *Leonardo Music Journal* 2(1), 3–12.
- Dittrich, P., J. C. Ziegler, and W. Banzhaf (2001). Artificial chemistries – a review. *Artificial life* 7(3), 225–275.
- Dixon, S. (2000). A lightweight multi-agent musical beat tracking system. In *Pacific Rim International Conference on Artificial Intelligence*, pp. 778–788.
- Draves, S. (2005). The electric sheep screen-saver: A case study in aesthetic evolution. In *Applications of Evolutionary Computing*, pp. 458–467. Springer.
- Eiben, A. E. and J. E. Smith (2003). *Introduction to Evolutionary Computing*. Springer Science & Business Media.

- Eigenfeldt, A. (2007a). The creation of evolutionary rhythms within a multi-agent networked drum ensemble. In *Proceedings of the 2007 International Computer Music Conference*, Copenhagen, Denmark, pp. 267–270.
- Eigenfeldt, A. (2007b). Drum circle: Intelligent agents in Max/MSP. In *Proceedings of the 2007 International Computer Music Conference*, Copenhagen, Denmark, pp. 9–12.
- Eigenfeldt, A. (2008). Emergent rhythms through multi-agency in Max/MSP. In R. Kronland-Martinet, S. Ystad, and K. Jensen (Eds.), *Proceedings of the 4th International Computer Music Modeling and Retrieval Symposium, CMMR 2007*, pp. 368–379.
- Eigenfeldt, A., O. Bown, and B. Carey (2015, June). Collaborative composition with creative systems: Reflections on the first musebot ensemble. In *Proceedings of the Sixth International Conference on Computational Creativity*, Park City, Utah.
- Eldridge, A. (2005). Cyborg dancing: generative systems for man-machine musical improvisation. In T. Innocent (Ed.), *Proceedings of Third Iteration.*, Melbourne, Australia, pp. 129–142. Centre for Electronic Media Art (CEMA).
- Eldridge, A. (2008). Self-directed feedback loops - OSX application. Available online at <http://ecila.org/>.
- Eldridge, A. (2012). The poietic power of generative sound art for an information society. In M. Matija, M. Kaltenbruger, and M. Cigar (Eds.), *Proceedings of the 2012 International Computer Music Conference*, Ljubljana, Slovenia.
- Eldridge, A. (2013). Improvising with self-observing systems: a duet for cellist and adaptive delay network. In M. Verdicchio and M. Carvalhais (Eds.), *Proceedings of the first conference on Computation, Communication, Aesthetics and X*, pp. 297 – 301. San Francisco: ICMA.
- Eldridge, A. and A. Dorin (2009). Filterscape: Energy recycling in a creative ecosystem. In *Applications of Evolutionary Computing*, pp. 508–517. Springer.

- Eldridge, A., A. Dorin, and J. McCormack (2008). Manipulating artificial ecosystems. In *Applications of Evolutionary Computing*, pp. 392–401. Springer.
- Elton, C. S. (1927). *Animal Ecology*. Sidgwick and Jackson, London.
- Epstein, J. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. Washington: Brookings Institution.
- Ermentrout, G. B. and L. Edelstein-Keshet (1993). Cellular automata approaches to biological modeling. *Journal of theoretical Biology* 160(1), 97–133.
- Froese, T. et al. (2010). Life after Ashby: ultrastability and the autopoietic foundations of biological autonomy. *Cybernetics & Human Knowing* 17(4), 7–49.
- Futayama, D. and M. Slatkin (1983). *Coevolution*. Sunderland: Sinauer.
- Garcia-Almanza, A. L. and E. P. Tsang (2006). Simplifying decision trees learned by genetic programming. In *Evolutionary Computation, 2006. CEC 2006. IEEE Congress on*, pp. 2142–2148. IEEE.
- Garnett, G. (2001). The aesthetics of interactive computer music. *Computer Music Journal* 25(1), 21–33.
- Green, O. (2011). Agility and playfulness: Technology and skill in the performance ecosystem. *Organised Sound* 16(02), 134–144.
- Gresham-Lancaster, S. (1998). The aesthetics and history of the hub: The effects of changing technology on network computer music. *Leonardo Music Journal* 8, 39–44.
- Haque, U. (2007). The architectural relevance of Gordon Pask. *Architectural Design* 77, 54 – 61.
- Harrald, L. (2005). Fight or flight: Towards the modelling of emergent ensemble dynamics. In *Australasian Computer Music Conference (2005: Brisbane, Australia)*.

- Harvey, I. (1997). Cognition is not computation; evolution is not optimisation. In *Artificial Neural Networks – ICANN’97 LNCS*, Volume 1327, pp. 685–690. Springer.
- Hinton, G. E. and S. J. Nowlan (1987). How learning can guide evolution. *Complex Systems* 1, 495–502.
- Hoffmann, P. (2002). Towards an “automated art”: Algorithmic processes in xenakis’ compositions. *Contemporary Music Review* 21(2-3), 121–131.
- Hogeweg, P. (1988). Cellular automata as a paradigm for ecological modeling. *Applied mathematics and computation* 27(1), 81–100.
- Holland, J. (1975). *Adaptation in Natural and Artificial Systems*. Ann Arbor: University of Michigan Press.
- Holland, J. H. (1999). Echoing emergence: Objectives, rough definitions, and speculations for echo-class models. In *Complexity*, pp. 309–342. Perseus Books.
- Husbands, P. (1993). An ecosystems model for integrated production planning. *International Journal of Computer Integrated Manufacturing* 6(1-2), 74–86.
- Husbands, P., P. Copley, A. Eldridge, and J. Mandelis (2007). An introduction to evolutionary computing for musicians. In E. R. Miranda and A. Biles (Eds.), *Evolutionary Computer Music*. Springer-Verlag.
- Impett, J. (2001). Interaction, simulation and invention: a model of interactive music. In *Proceedings of ALMMA 2001 Workshop on Artificial Models for Musical Applications*, Cosenza, Italy, pp. 108–119.
- Kaschube, M., M. Schnabel, S. Löwel, D. M. Coppola, L. E. White, and F. Wolf (2010). Universality in the evolution of orientation columns in the visual cortex. *science* 330(6007), 1113–1116.
- Kauffman, S. (1993). *The Origins of Order: Self-Organization and Selection in Evolution*. New York: OUP.
- Kelso, J. (1995). *Dynamic Patterns: The Self-Organization of Brain and Behavior*. The MIT Press/A Bradford Book.

- Kiefer, C. (2014). Interacting with text and music: Exploring tangible augmentations to the live coding interface. In *Proceedings of the International Conference for Life Interfaces*.
- Kim-Boyle, D. (2005). Sound spatialization with particle systems. In *Proceedings of the 8th international conference on digital audio effects (DAFX-05), Madrid, Spain*, pp. 65–68. Citeseer.
- Kirby, S. and J. Hurford (1997). Learning, culture and evolution in the origin of linguistic constraints. In *Fourth European Conference on Artificial Life*, pp. 502–521. MIT Press.
- Koza, J. R. (1992). *Genetic programming: on the programming of computers by means of natural selection*, Volume 1. MIT press.
- Langton, C. (1989). Artificial life. In C. Langton (Ed.), *Artificial Life. Santa Fe Institute Studies in the Sciences of Complexity*, pp. 2–3. Addison-Wesley.
- Large, E. W. and C. Palmer (2002). Perceiving temporal regularity in music. *Cognitive Science* 26, 1–37.
- Larson, S. (1997). Musical forces and melodic patterns. *Theory and practice* 22, 55–71.
- Lima, M. (2006). *Visual complexity: Mapping Patterns of Information*. Princeton Architectural Press.
- Linson, A. (2013). The expressive stance: Intentionality, expression and machine art. *International Journal of Machine Consciousness* 5(2), 195 – 216.
- Linson, A., C. Dobbyn, G. Lewis, and R. Laney (2015). A subsumption agent for collaborative free improvisation. *Computer Music Journal* 39(4).
- Lorenz, E. N. (1963). Deterministic nonperiodic flow. *Journal of the atmospheric sciences* 20(2), 130–141.
- Magnus, C. (2006). Evolutionary musique concrete. In *Applications of Evolutionary Computing*, pp. 688–695. Springer.

- Maturana, H. R. and F. J. Varela (1987). *The Tree of Knowledge: The Biological Roots of Human Understanding*. New Science Library/Shambhala Publications.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. New York: Cambridge University Press.
- McCormack, J. (1993). Interactive evolution of L-system grammars for computer graphics modelling. In *Complex Systems: From Biology to Computation*, pp. 118–130. ISO Press, Amsterdam.
- McCormack, J. (2001). Eden: An evolutionary sonic ecosystem. In *Advances in Artificial Life*, pp. 133–142. Springer.
- McCormack, J. (2007). Artificial ecosystems for creative discovery. In *Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation*, pp. 301–307. ACM.
- McCormack, J. (2008). Facing the future: Evolutionary possibilities for human-machine creativity. In *The Art of Artificial Evolution*, pp. 417–451. Springer.
- McCormack, J. (2012). Creative ecosystems. In *Computers and Creativity*, pp. 39–60. Springer.
- McCormack, J. and O. Bown (2009). Life’s what you make: Niche construction and evolutionary art. In *Applications of Evolutionary Computing*, pp. 528–537. Springer.
- McCormack, J. and A. Dorin (2001). Art, emergence and the computational sublime. In *Proceedings of Second Iteration: A Conference on Generative Systems in the Electronic Arts*. Melbourne: CEMA, pp. 67–81. Citeseer.
- McDermott, J., D. Sherry, and U.-M. O’Reilly (2013). Evolutionary and generative music informs music hci—and vice versa. In *Music and human-computer interaction*, pp. 223–240. Springer.
- Mendel, G. (1866). *Versuche über Pflanzenhybriden*, Volume 44. Brnn: Im Verlage des Vereines.
- Millen, D. (1990). Cellular automata music. In *Proceedings of the 1990 International Computer Music Conference*, pp. 314–316.

- Minsky, M. (1988). *Society of Mind*. Simon and Schuster.
- Miranda, E. (1999). Towards an Artificial Life approach to the origins of music. In *Proceedings of the XIX Annual Congress of the Brazilian Computing Society / VI Brazilian Symposium on Computer Music*, Rio de Janeiro. Brazilian Computing Society.
- Miranda, E. (2000). On the music of emergent behaviour: what can evolutionary computation bring to the musician? In *Proceedings of Genetic and Evolutionary Computation Conference*, Las Vegas, USA. University of Central Florida.
- Miranda, E. (2002a). Emergent sound repertoires in virtual societies. *Computer music Journal* 26(2), 77–90.
- Miranda, E. R. (1995). Granular synthesis of sounds by means of a cellular automaton. *Leonardo* 28(4), 297–300.
- Miranda, E. R. (2002b, October). Sounds of artificial life. In *Proceedings of the ACM SIGCHI International Conference - Creativity and Cognition 4*. Loughborough University ACM.
- Miranda, E. R. and A. Biles (2007). *Evolutionary Computer Music*. Springer.
- Miranda, E. R., S. Kirby, and P. Todd (2003). On computational models of the evolution of music: From the origins of musical taste to the emergence of grammars. *Contemporary Music Review* 22(3), 91–111.
- Mitchell, M. (1998). *An Introduction to Genetic Algorithms*. MIT press.
- Murray-Rust, D., A. Smaill, and M. Edwards (2006). MAMA: An architecture for interactive musical agents. *Frontiers in Artificial Intelligence and Applications* 141, 36.
- Neumann, J. V. (1966). *Theory of Self-Reproducing Automata*. Champaign, IL, USA: University of Illinois Press.
- Nicolis, G., I. Prigogine, et al. (1977). *Self-organization in nonequilibrium systems*, Volume 191977. Wiley, New York.
- Pask, G. (1971). A comment, a case history and a plan. In J. Reichardt (Ed.), *Cybernetics, Art and Ideas*. Studio Vista ltd.

- Reichardt, J. (1971). *Cybernetics, Art and Ideas*. Studio Vista Ltd.
- Reynolds, C. W. (1987). Flocks, herds and schools: A distributed behavioural model. In *Computer Graphics (SIGGRAPH '87 Conference Proceedings)*, Volume 21, pp. 25–34.
- Romero, J. J. and P. Machado (2007). *The Art of Artificial Evolution: a Handbook on Evolutionary Art and Music*. Springer.
- Sanfilippo, D. (2013). Turning perturbation into emergent sound, and sound into perturbation. *Interference 4*.
- Schelling, T. C. (1971). Dynamic models of segregation. *Journal of mathematical sociology 1*(2), 143–186.
- Searle, J. R. (1969). *Speech acts: An essay in the philosophy of language*, Volume 626. Cambridge university press.
- Secretan, J., N. Beato, D. B. D Ambrosio, A. Rodriguez, A. Campbell, and K. O. Stanley (2008). Picbreeder: evolving pictures collaboratively online. In *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, New York, NY, USA, pp. 1759–1768. ACM.
- Shiffman, D., S. Fry, and Z. Marsh (2012). *The nature of code*. D. Shiffman.
- Sims, K. (1991, July). Artificial evolution for computer graphics. In *Computer Graphics (Siggraph '91 proceedings)*, pp. 319–328.
- Sommerer, C. and L. Mignonneau (1997). A-volve: an evolutionary artificial life environment. In *Artificial Life V: Proceedings of the fifth International Workshop on the Synthesis and Simulation of Living Systems*, Nara-shi, Japan, pp. 167–175. MIT Press.
- Stubbs, D. (2009). *Fear of music: Why people get Rothko but don't get Stockhausen*. John Hunt Publishing.
- Sudnow, D. (1978). *Ways of the hand: The organization of improvised conduct*. MIT Press.

- Takagi, H. (2001). Interactive evolutionary computation: Fusion of the capabilities of ec optimization and human evaluation. *Proceedings of the IEEE* 89(9), 1275–1296.
- Tansley, A. (1935). The use and abuse of vegetational concepts and terms. *Ecology* 16, 284–307.
- Toch, E. (1977). *The shaping forces in music: an inquiry into the nature of harmony, melody, counterpoint, form*. Courier Corporation.
- Todd, P. and G. Werner (1998). Frankensteinian methods for evolutionary music composition. In N. Griffith and P. Todd (Eds.), *Musical Networks: Parallel Distributed Perception and Performance*, pp. 173–190. Cambridge, MA: MIT Press.
- Todd, S. and W. Latham (1991). Mutator, a subjective human interface for evolution of computer sculptures. Technical report, IBM, United Kingdom Scientific Center Report.
- Todd, S. and W. Latham (1999). The mutation and growth of art by computers. In P. Bentley (Ed.), *Evolutionary Design by Computers* (2nd ed.), Chapter 9, pp. 221–250. Morgan Kaufmann.
- Varela, F. G., H. R. Maturana, and R. Uribe (1974). Autopoiesis: the organization of living systems, its characterization and a model. *Biosystems* 5(4), 187–196.
- Von Bertalanffy, L. (1950). An outline of general system theory. *British Journal for the Philosophy of Science* 1, 134 – 165.
- Wallace, A. (1858). On the tendency of varieties to depart indefinitely from the original type. In *Proceedings of the Linnean Society of London*, Volume 3, pp. 53 – 62.
- Walter, W. G. (1950). Imitation of life. *Scientific American* 182(5), 42–45.
- Waschka II, R. (2001). *Artificial Life Models for Musical Applications*, Chapter Theories of Evolutionary Algorithms and a 'New Simplicity' Opera: Making Sappho's Breath, pp. 79–86. Cosenza, Italy: Editoriale Bios.
- Waters, S. (2007). Performance ecosystems: Ecological approaches to musical interaction. In *Proceedings of Electro Acoustic Studies Network*.

- Whitelaw, M. (2004). *Metacreation: Art and Artificial Life*. Cambridge, Mass.: MIT Press.
- Whitelaw, M. (2005). System stories and model worlds: A critical approach to generative art. In *Readme 100: Temporary Software Art Factory*, Norderstedt, pp. 135–154. BoD.
- Wiggins, G. A. (2012). The mind’s chorus: Creativity before consciousness. *Cognitive Computation* 4(3), 306–319.
- Wilkinson, G. S. (1984). Reciprocal food sharing in the vampire bat. *Nature* 308(5955), 181–184.
- Xenakis, I. (1971). Free stochastic music. In J. Reichardt (Ed.), *Cybernetics, Art and Ideas*, pp. 124–142. London: Studio Vista Ltd.
- Yee-King, M. and M. Roth (2008). Synthbot: An unsupervised software synthesizer programmer. In *Proceedings of the 2008 International Computer Music Conference*, Belfast, Ireland.