An Introduction to Musical Metacreation

PHILIPPE PASQUIER, School of Interactive Art + Technology, Simon Fraser University
ARNE EIGENFELDT, School for the Contemporary Arts, Simon Fraser University
OLIVER BOWN, Department of Art and Design, University of New South Wales
SHLOMO DUBNOV, Department of Computer Science, University of California at San Diego

Musical metacreation (MuMe), also known as musical computational creativity, is a subfield of computational creativity that focuses on endowing machines with the ability to achieve creative musical tasks, such as composition, interpretation, improvisation, accompaniment, mixing, etc. It covers all dimensions of the theory and practice of computational generative music systems, ranging from purely artistic approaches to purely scientific ones, inclusive of discourses relevant to this topic from the humanities. MuMe systems range from purely generative ones to a variety of interactive systems, such as those for computer-assisted composition and computer-assisted sound design. In order to better appreciate the many dimensions of this interdisciplinary domain and see how it overlaps and differs from research in computer music, this introduction provides a general entry point. After defining and introducing the domain, its context, and some of its terminology, we reflect on some challenges and opportunities for the field as a whole.

CCS Concepts: • **Computing methodologies** → **Artificial intelligence**; **Machine learning**; • **Applied computing** → **Sound and music computing**

Additional Key Words and Phrases: Artificial intelligence, machine learning, creative computing, computational creativity, generative art, computer music

ACM Reference Format:
Philippe Pasquier, Arne Eigenfeldt, Oliver Bown, and Shlomo Dubnov. 2016. An introduction to musical metacreation. Comput. Entertain. 14, 2, Article 2 (December 2016), 14 pages.
DOI: http://dx.doi.org/10.1145/2930672

1. INTRODUCTION

In 1960, Herbert Simon defined Artificial Intelligence (AI) as "the science of having machines solve problems that do require intelligence when solved by humans" [Simon 1960]. Since then, AI has had some tremendous successes at rational problem solving. A wide variety of algorithms have been designed and studied that allow us to explore increasingly complex search spaces in order to attain optimal or near-optimal solutions to abstract and concrete problems, such as flying planes; regulating nuclear plants; designing electric circuits; automating negotiation; diagnosing diseases; and playing chess, Go, or Jeopardy. While the list of tasks addressed at human-competitive levels by AI systems is growing steadily, some are looking more carefully at an old question: Can machines be creative?

Authors' addresses: P. Pasquier (corresponding author), School of Interactive Art + Technology, Simon Fraser University, Vancouver, Canada; email: pasquier@sfu.ca; A. Eigenfeldt, School for the Contemporary Arts, Simon Fraser University, Vancouver, Canada; email: eigenfel@sfu.ca; O. Bown, Department of Art and Design, University of New South Wales, Sydney, Australia; email: ollie@icarus.nu; S. Dubnov, Department of Computer Science, University of California at San Diego, San Diego, USA; email: sdubnov@ucsd.edu.

DOI: http://dx.doi.org/10.1145/2930672

2. COMPUTATIONAL CREATIVITY

Computational creativity, or metacreation, is a fast-growing field that addresses that question by exploring the automation of creative processes, with the aim to endow machines with creative behaviors. As a field, it investigates:

—*Creativity as It Is.* It is striving to understand and simulate human creativity. What is creativity? If we can understand this complex notion, can we simulate it? As such, when such a cognitive modeling approach is chosen, it is also part of cognitive science.
—*Creativity as It Could Be*. The field is also devoted to exploring processes of which we know humans alone to be incapable. The outcomes of those artificial processes might nonetheless be considered novel and valuable, that is creative.

2.1. How Does Computational Creativity Differ from Traditional AI?

Computational creativity differs from the rest of rational problem solving AI, in that it frequently addresses problems for which the notion of optimality of a solution is *a priori* ill-defined. These problems include those for which there are no definitive yes or no answers, no clearly defined win or lose states, no clearly defined goal states, no well-defined Pareto dominance, no obvious and complete objective functions, no clear utility functions, and ill-defined preference relations. In other words, it addresses a class of problems for which the usual guides and the search for solution end-points are not clearly defined.

We define computational creativity around the notion of creative tasks for which there is no clear "best" outcomes or optimal solutions. We cannot define the optimal choreography, composition, or musical interpretation. We can state our *preferences* for the best drawing, painting, narrative, poetry, joke, recipe, or level for a video game; however, there is no method to determine an objective answer to these ideals. A multitude of tasks that do not fall *a priori* under the paradigm of rational problem solving constitute this broad category, across a multitude of domains.

In a nutshell, *computational creativity* is the science of machines addressing creative tasks. Humans perform these tasks every day; although they do not have clearly defined solutions, they are nevertheless clearly defined areas of activity with well-known associated behaviors.

2.2. Why Is Computational Creativity Important?

There are a number of reasons why research and progress in this area matters. First, it is fundamental research on creativity and creative processes. As with intelligence, the nature of creativity remains the topic of much debate, yet it is one of the essential traits of humanity that scholars have been trying to understand and define for millennia. Because creativity is a complex phenomenon, the computational implementation and simulation of the various theories and models is an ideal method of inquiry. As with any complex phenomenon, such as the weather, the nervous system, or capital markets, a powerful way to deploy, test, and eventually validate models of these systems is through computational simulation.

Second, our use of computers is no longer limited to rational problem solving[1]; instead, they have become ubiquitous within personal communication, entertainment, art, and culture. With the computerization of society in the developed world, more and more people are surrounded by an increasing number of processors that are used

[1]"A system is rational if it does the right thing", which assumes well-defined problems to which there are correct, and optimal or near optimal solutions [Russell and Norvig 2009]. Rational problems are very important and cover a wide range of tasks, from playing Go, and designing electronic circuits, to general calculation, such as basic algebra.

for creative tasks: composing messages, editing photographs and movies, designing graphics, computer entertainment, as well as within the more traditional fine arts. This increasing volume of creative tasks calls for more attention to be devoted to creative processes, and their eventual automation.

Third, there is a need for generative systems in the creative industries. The move from linear to nonlinear media entailed an explosion in the number of assets required. Within a feature animation movie, for example, every second of the film will be painstakingly edited by a creative team. However, unlike a film, the elements within online multiplayer computer games are not fixed; new elements, from the smallest of objects to entire worlds, replete with sound and music unique to the situation, are required to fulfill the demands of millions of players, each playing hours per week. This is a problem currently faced by the video game industry, but is one that increasingly will appear in any nonlinear creative application.

Finally, current creative software is fundamentally inert. AI and machine learning methodologies have seldom made their way into the creative software industry. Software is often creatively brittle and doesn't afford flexible interactive creative search. The software will require the user to click the same buttons, select the same menu items, and repeat the same mundane tasks sequentially in order to complete similar processes, processes that have often been done in similar tasks in the past, all of which interfere with artistic flow. Providing creative suggestions through autonomous and learned behaviors, and more generally bringing greater automation and artificial intelligence to creative software, is an important goal of applied computational creativity.

2.3. Defining Creativity

Despite the best efforts of philosophers, psychologists, cognitive scientists, and educators, there is no clearly accepted definition of creativity [Peter 2009]. A fairly modest and uncontroversial approach was formulated by creativity philosopher Margaret Boden [2004]. She defines creativity as the ability to come up with ideas or artifacts that are original, and valuable: original in that they are new and valuable in any sense one can reify this notion, monetary or otherwise.

Boden goes on to delineate two kind of creativity:

—*P-creativity*. Psychological creativity (novel and valuable for the individual), a.k.a *mundane or everyday creativity*;
—*H-creativity*. Historical creativity (novel and valuable for the group, i.e., humanity), a.k.a *eminent creativity*.

Perhaps more relevant for our purpose, she also defines three types of creativity:

—*Exploratory Creativity.* In which one is producing novel and valuable instances in an existing creative space. Given all the pieces of heavy metal music, for example, can one find one that is unique (i.e., novel) and considered interesting (i.e., valuable), yet clearly remaining within that defined space?
—*Combinatorial Creativity*. This type of creativity involves bringing creative spaces together in a new ways. A cellphone that is also a camera is a simple example of combinatorial creativity that was novel a decade ago, and still remains valuable today. Conceptual blending and metaphors are also examples of combinatorial creativity.
—*Transformational Creativity.* This occurs when the creative space itself is modified: for example, the creation of a new style of art, music, fashion, heretofore never imagined.

The first two notions can, and have been, automated through computational means, and many fine examples exist. The third is less well understood, and is complicated and somewhat rare; the conditions for this to occur, and to be recognized, are beyond the current conceptions of computational creativity.

An alternative distinction [Bown 2012] is between *generative* and *adaptive* creativity. Generative creativity actually extends well beyond the scope of human action, and includes any process whereby new things come into being. For example, natural evolution is a profoundly creative force, but its nature is generative; there is no beneficiary to its unfolding. Adaptive creativity is the more common sense of creativity as an intentional cognitive act: adaptive in the sense that it improves the actor's situation. The distinction is important because there are many human processes, usually occurring on the social level, where there is no apparent overall value to the creative product, such as in entire movements in art or music. We would call these generative. Thus, creative outcomes can be considered the combination of both generative and adaptive creative processes.

3. DEFINING MUSICAL METACREATION

With these basic notions in place, we can define *musical*[2] *metacreation* as a subfield of computational creativity that addresses music-related creative tasks. While computer music and related fields address a variety of problems related to music – including music perception, music recognition, music classification, music representation, and music cognition – MuMe has its own set of specific dimensions, presented in the following subsections.

3.1. Domain of the System and Problem Addressed

Due to its structured nature [Damschroder and Williams 1990], music has as rich history of generative approaches devoted to its creation [Hedges 1978; Sloboda 1988; Englert 1989; Dorin 2001; Nierhaus 2009; Edwards 2011]. Musical metacreation is the natural evolution of these procedures, and a wide variety of research has been undertaken to investigate automated musical creation tasks often already achieved through the use of computer software. As such, the first characteristic of a MuMe system is the problem it addresses; in other words, what does it generate?

While the diversity of problems addressed by MuMe systems is too long to be exhaustive here, groups or families of systems can be organised around recurrent "canonical" problems:

—*Composition*. The system must produce novel music compositions, either through algorithmic means, or using knowledge derived from a corpus of existing compositions. The output can be any symbolic representation, such as a score in MIDI format or other notation, of a pattern, part, piece, or group of musical pieces;
—*Interpretation*. Given a composition (in any format), the interpretation consists in producing an audio rendering of the composition. A typical example is a musician or an orchestra performing from a score;
—*Improvisation*. The creation and real-time performance of new musical material spontaneously, to varying degrees, and often enacted in groups. While it can be prepared and rehearsed, it is assumed that improvisation occurs in real-time.
—*Accompaniment*. One or several musical parts (e.g., melody, harmonic progression, rhythmic accompaniment) is provided, and other parts need to be either composed or interpreted (or both).

In addition to such classic musical tasks, specific tasks and problems have been devised by the MuMe community. An example of such a task or sub-problem is *style imitation*

[2]Here, and in the following, music is meant in an inclusive sense of organized sound, as well as peripheral domains such as sound design, sound synthesis, and any creative task involving sound-related concepts.

[Dubnov et al. 2003], which can be broadly defined as such: Given a corpus C = {C_1, ... C_n} representative of style S, generate new instances that would be classified as belonging to S by an unbiased observer (typically a set of human subjects).

Another example is *continuation* that has been brought forward through the work on style imitation of Pachet [2003], and consists in having a musician play or improvise, and the system taking over once the musician stops.

The basic problems and tasks introduced above can be refined and broken down into a multitude of sub-problems and sub-domains. Each MuMe system has its own domain and task. For example, systems may limit themselves to addressing the generation of:

—Harmonic progressions [Eigenfeldt and Pasquier 2010; Whorley et al. 2010; Groves 2013; Manaris et al. 2013; Pachet and Roy 2014];
—Rhythm generation [Eigenfeldt 2008; Chordia and Rae 2010];
—Melodic generation [Bosley et al. 2010; Sarwate and Fiebrink 2013];
—Orchestration [Handelman et al. 2012];
—Harmonization [Pachet and Roy 2001; Simon et al. 2008; Pachet and Roy 2014];
—Affective interpretation [Kirke and Miranda 2009];
—Affective composition [Birchfield 2003; Wallis et al. 2011; Eigenfeldt et al. 2015];
—Automatic mixing [Reiss and Perez Gonzalez 2008; Reiss 2011];
—Soundscape composition [Eigenfeldt and Pasquier 2011; Thorogood et al. 2012].

In doing so, some generate simple patterns, phrases, or sequences, while other systems may incorporate all aspects in their generation of full pieces, thereby accounting for the multilevel nature of musical constructs.

Some systems not only address a musical task, but might be “multidisciplinary” or multimedia in nature, such as generating music for film [e.g., Sorenson and Brown 2008], installation [e.g., Schedel and Rootberg 2009; Eigenfeldt et al. 2014], or games [Collins 2008; Brown 2012].

One of the pervasive issues with MuMe systems is the question of specifying the creative interaction between the musician or user and the system. The research on metacreative musical interfaces draws upon the research on musical interfaces at large within the NIME (New Interfaces for Musical Expression) community. Questions of a system’s autonomy – specifically in relation to interaction – will be discussed below.

3.2. Levels of Autonomy of the System

MuMe software can be classified on a continuum that ranges between two extremes: from systems with no autonomy to completely autonomous systems. At one end of the spectrum, we find systems that have no or very little automation, or autonomy, e.g., the typical digital audio workstation (DAW), which historically has been the equivalent of a musical typewriter: a general purpose music system in which every single action requires human intervention. We do note that some contemporary DAWs include features that could be regarded as creativity enhancers, and automate some tasks for the user, including automatic rhythm alignment or audio transcription.

At the other extreme, there are completely generative systems. The user launches – and potentially configures – the system, and the musical task to which it is dedicated unfolds – e.g., *EMI* [Cope 1996]. Note that purely generative systems can be interactive – e.g., *Voyager* [Lewis 2000]. Eventually, all graduations of this continuum, can be envisioned through a range of computer-assisted creativity tools, such as those devoted to computer assisted composition tools [e.g., Truax 1985; Cope 1997; Assayag et al. 1999; Maxwell et al. 2012].

The autonomy [Bradshaw et al. 2013] of a system is defined by both its level of self-sufficiency and its level of self-directedness, and the balance between these dimensions is one of the main reasons for the diversity of MuMe systems. A full taxonomy of MuMe systems according to these has yet to be decided upon, although one has been proposed [Eigenfeldt et al. 2013]. This issue is also debated in the computational creativity community, where some distinguish merely generative systems, generative systems with feedback, and generative and reflective systems [Agres et al. 2016]. More generally, and to illustrate the complexity of the subject, such typologies are still debated in the AI community, with various types of agents ranging from reactive to cognitive being defined by the various dimensions of autonomy and intelligence being considered [Wooldridge 2009].

3.3. Origin of the System's 'Knowledge'

Early algorithmic composition systems [e.g., Koenig 1983] relied upon the user to assign values, dynamic or otherwise, to various parameters prior to generation, thereby relying upon the user's musical knowledge and/or aesthetic judgment. Later systems [Cope 1996] have extracted such knowledge from a provided corpus of musical works or excerpts. We therefore make a distinction between non-corpus-based systems that generate their output without being exposed to musical information as an input, and corpus-based systems that have been exposed to music, either through symbolic notation (most often MIDI data) or from audio data.

3.4. The Nature of the Inputs/Outputs of the System

These are also varying in type:

—*Symbolic*. Most MuMe systems have input and/or output that are of a discrete symbolic nature, such as musical notation, MIDI files, or MIDI. Examples include Cope's *EMI* (MIDI file input, and MIDI file output), Lewis' *Voyager* (MIDI input and output) and *Kinetic Engine* (no input, MIDI output) [Eigenfeldt 2008];
—*Audio*. Some systems use audio signals as input and/or generate audio directly. Examples include Young's *piano_prosthesis* and *oboe_prosthesis* [2008] or the *Audio Metaphor* soundscape generation systems [Thorogood et al. 2012];
—*Hybrid*. Some systems accept both types of inputs and outputs. Examples include Sony's *Continuator* [Pachet 2003], and UCSD/IRCAM's *Audio Oracle* [Dubnov et al. 2003].

It is worth noticing that these inputs/outputs can be specific non-directly musical representations. For examples, systems have been conceived that generate presets for a musical instrument [Macret and Pasquier 2013], generate sound making programs [Garcia 2001; Macret and Pasquier 2014], or synthesizer architectures [Yee-King 2011].

3.5. The System's Relationship to Time

MuMe systems can be differentiated in how they operate in time:

—*Online*. The music is generated live in performance, in real time. Such systems may react to live input or be interactive in some way, or may be purely generative.
—*Offline*. While the generation itself can occur slower or faster than real time, the actual generation does not occur in real time.

A class of systems that typically take both inputs and outputs and are online are musical agents. The domain of musical agents has been very dynamic and the exploration of a vast variety of mono- and multi-agent architectures are being explored (e.g., Wulfhorst et al. 2001; Dahlstedt and McBurney 2006; Murray-Rust et al. 2006; Beyls 2007; Eigenfeldt 2008; Collins 2008a; Whalley 2009; Eigenfeldt and Pasquier 2011; Martin et al. 2011).

3.6. Generality of the System

Another dimension to consider regarding the nature of MuMe systems is the continuum between systems that are specific and systems that are generic. Many MuMe systems are created by musician/composers, with artistic intentions in mind; as such, they will tend to reflect at some level, purposefully or otherwise, the aesthetic of the designer. Such systems will tend to produce – ideally musically successful – variations of the same composition rather than creating completely novel outputs with each run. The sacrifice in generality may be appropriate for this context, in that the specificity that can favor successful music is increased [Truax 1980].

On the other end of this spectrum, some systems are generic. For example, corpus-based style-imitation systems [e.g. Conklin 2003] attempt to derive all musical knowledge from the corpus provided; and as such, they have the potential to be aesthetically agnostic. Along this continuum, a variety of levels of generality are possible, with system using heuristics or encoding certain musical knowledge with various levels of specificity.

This continuum of generality is also reflected by the MuMe community, which spans purely artistic work to purely scientific, going through all the possible hybrids and combinations in between.

3.7. The System's Inner Working

A consequence of the diversity of approach described above, is that the inner workings of the system varies from messy vernacular code full of heuristic and ad hoc approaches and algorithms, to the application of reproducible well-described and documented algorithms.

An entire book is necessary to simply list the main algorithmic approaches to musical composition [Nierhaus 2009]. In general, any process or algorithms that can be used in generative ways that one can think of has been tried for a musical task, be there from AI, machine learning, artificial life, pattern recognition, optimisation, or regular computer science and algorithmics. The whole gamut from simple randomness and basic probability theory to the most sophisticated deep learning algorithms or cognitive agent architecture has or is being explored by the hundreds of MuMe researchers, practitioners and enthusiasts around the world.

4. EVALUATION OF MUME SYSTEMS

Since MuMe systems explore musically creative tasks; they can, in principle, be evaluated in ways similar to their human-acheived counterparts. Therefore, evaluation of MuMe outputs can be done by/through:

(1) *Authors*. The artists, designers, computer scientists, and more generally creators of a MuMe system are usually the first to judge the success of their system's output;
(2) *Users, Peers, and Experts*. Other composers, musicians, sound designers, and researchers using the system and providing critical feedback;
(3) *Audience*. Many creative tasks are designed to produce artifacts that will be presented to a larger audience. It is possible to measure the popularity of a given system's output in very direct ways: is an audience willing to pay to experience the artifact, either through concert tickets sales or album purchases/downloads?
(4) *Press and Media Coverage*. The interest shown by critics, journalists, concert reviews, and album reviews;
(5) *Peer Reviewers, Curators and Jurors*. How the system is evaluated through peer-reviewed academic papers, inclusion in peer-reviewed concerts/festivals, the awarding of academic grants;

(6) *Theoretical and Analytic Measures*. Peer-review of academic papers describing the system's processes;
(7) *Empirical Studies*. Specific qualitative or quantitative user/audience studies on the system's output.

Musical metacreation can be seen as a scientific inquiry, an artistic inquiry, or both. The various types of evaluation listed above do cover both spectra, and the appropriate one must be chosen accordingly based upon the intentions of the authors. Different methods apply in each case, but both demand the same degree of rigour and clarity about the work's objectives and outcomes. In some of these cases, such as (4), great care must be taken to ensure that a critical analysis of the evaluation was conducted; did it only impress critics because of the exciting idea behind the project, or because of the quality of the system's original output?

Evaluations of artworks are difficult, and involve the complex and subjective notion of aesthetics [Dickie 1985]. However, we can attempt to address the following aspects of a MuMe system:

—*Quality*. Does the system achieve the task it sets out to achieve? Is it any good at it? Is the system human-competitive at the task?
—*Creativity*. Is the system demonstrating creativity? What type? How does the system fare against humans?
—*Believability*. For some research projects, the process implemented is meant to behave in ways that are anthropomorphic (or as a simulation of some computational model of a natural agent). How successful at such imitation is the system?
—*Complexity*. The algorithmic properties of the system are often of interest. What class of complexity is the creative task defined within? Does the system produce a richness of behavior?
—*Robustness*. Can the system operate in a robust fashion on all input, or is it meant to react to a small set of input data?
—*Reliability*. Can the system guarantee a reliable quality of output independently of the input? Has the output of the system been "cherry-picked"?

4.1. Theoretical Aspects of Evaluation

There are a variety of theoretical reasons as to why evaluation is difficult. By the very nature of an artistically creative task, we are looking at problems for which there is no notion of optimality. A direct consequence of this is that most measures used to evaluate the quality of the production of a system in computer science and AI do not readily apply. *A priori*, a melody generation system cannot be assessed through typical methods found in rational problem solving, which measure error and distance to an optimal solution, because there isn't such a thing as "an optimal melody".

Second, when we resort to human subjects in order to evaluate a given system for a given task, subjective and cultural impressions and judgments are involved. Furthermore, the presentation of musical artifacts is multidimensional, and *framing* can play a role [Charnley et al. 2012]. People's judgments are influenced by who produced the work, and how was it produced. For example Salganik et al. [2006] provided experimental evidence of winner-takes-all effects in musical evaluation. They compared a blind evaluation task with one in which people could see the overall scores of the pieces of music, as rated by others, the latter having more extreme winners and losers. One study in computational creativity has suggested that judges may be biased against musical metacreation [Moffat and Kelly 2006].

In reality, though, it may be naive and ineffective to treat evaluation as a simple measurement, like taking a temperature. A person experiencing a piece of art or music does not perform a single judgement, but builds a complex and multifaceted perception of it, that may be manifest in different ways over time and in context. There are serious questions about whether reducing this interaction to a numeric rating is at all meaningful and the debate around the appropriateness of the various research instruments is very much alive [Yannakakis and Martinez 2015].

4.2. Practical Aspects of Evaluation

Finally, there is a plethora of practical reasons why the evaluation of MuMe systems is a challenging topic. Consider the following list of practical matters that are intrinsic to generative systems or affect the performance of the system in very complex ways. In doing so they add to the difficulty in evaluation:

- Choice of the corpus: for corpus-based systems, the choice of music from which to learn will influence the system's output. This is the "inspiring set", as defined by Ritchie [2007]. How big should it be? How diverse, or compact, should it be? If it does well on one set, and less so on another, does this mean the results were "programmed in" – a concern within computational creativity – or that, like humans, it may just be better at doing certain things than others?
- Choice of the parameters: most MuMe systems are heavily parameterized, with different choices of parameters leading to different behaviors that can significantly affect the performance of the system across various dimensions. While this is a problem common to most complex computer systems – in particular the ones using machine learning – the problem is worsened by the non-objective nature of the tasks at hand.
- The system needs to be evaluated on a sample output: generative systems can create *ad infinitum*, so does this mean that any or all output is valid for testing? Or can the designer select from these outputs?[3] Ritchie [2007] discusses this, and the above two issues, as problems in evaluating *typicality, novelty,* and *quality* of creative artefacts.
- The various uses of the system need to be taken into account: MuMe systems do not exist in a vacuum, and a system might perform well for a particular use and not for others. For example, an online rhythm generation system might be a useful tool as an inspiration machine in the context of a composition program, but a terrible musical agent in the context of live improvisation.
- The isolation of the task: most musical tasks interact with their context. For example, composition and interpretation are contingent upon one another. A well-crafted composition, human- or computer-generated, is likely to be perceived unfavourably by listeners if performed without expressivity. As most MuMe systems also involve performance, can a listener separate the composition from the performance?

Due to these difficulties, evaluation methodologies for MuMe systems remain an open research area. Evaluation strategies need to be tailored to specific research goals in ways that are relevant and have integrity, and may range from highly positivist strategies to approaches such as ethnography, design research methods or informal discussion and reflection. The research community must explore existing and novel research instruments that can support qualitative and quantitative comparative studies.

5. ADDITIONAL POINTS OF DISCUSSION

5.1. Generation vs Creation: The Question of the Authorship

A clear goal for practitioners within MuMe is the generation of musical material that would be considered equally creative as output produced by humans. The notion of systems – computational or otherwise – generating artistic material has been ongoing for centuries: as Galanter [2003] suggests, systematic art generation is as old as art

[3]While this may be considered "cherry-picking" of results, and thus negatively viewed from a scientific standpoint, it would be difficult to find a human artist that publically presents every instance of work created without any form of selection.

itself. Although generative art does not require the use of a computer, the prevalence of digital tools in the last fifty years has allowed for a dramatic increase in art and music using generative techniques. With the concurrent increase in power of computational systems, coupled with the application of machine learning and AI more generally, generative music has reached the point where complete works can be generated by machine.

At what point does the output of the system become the machine's, and not the designer's? In asking this, we must remain aware of the collaborative, social, networked factors involved in any act of creativity, and reject the automatic presumption that attributions of creativity must lie with any one easily defined agent [Saunders and Bown 2015].

One benefit of generative systems for artists is the ability to explore the variety of output from a designed process; the artistic act can be found in the design of the system, as well as the selection from its output. This is a commonly encountered example of the complex overlap that confounds attributions of creativity. But whilst we may not be able to accurately pinpoint the creative capacity of the system, there is absolutely nothing stopping us forming a rich description of the system in a working context, as both designers and ethnographers are experienced at doing. There is a lot of work to be done in the analysis of the system's output [Collins 2008b] as much as in the observation and evaluation of the process.

5.2. Community and Scope

While MuMe has tended to focus on systems and the music they produce, the field covers the wide spectrum of practices around generative music from purely artistic to purely scientific, including the humanities.

Given that MuMe systems involve knowledge about both computers and music, the field is inherently multidisciplinary, and projects are often interdisciplinary [National Academies 2004]. While the arts-science dichotomy is sometimes an issue for the computational creativity community, there are many ways in which these aspects can be better integrated rather than set in opposition to each other, without compromising the rigour of either approach. Continued development in this interdisciplinary field requires expertise from various domains, including both arts and sciences, in order to develop successful systems. As suggested above, design, characterized by *designerly ways of knowing* [Cross 2006], can act as a successful mediator between these areas of expertise. After all, this is one of the key domains influencing how the work is experienced by people.

Finally, MuMe is not only concerned with the design, development and evaluation of generative musical systems, but also with any discourse about – and study of – these systems, be it musicological, philosophical, psychological, neurological, sociological, anthropological, economical, mathematical, communicational, ecological, biological or medical. The diversity of community and broad scope are inherent interests and strengths of the domain.

5.3. Roadmap and Opportunities

At the time of writing, we identify a number of key developments that would help consolidating a collaborative research community:

—*Benchmark Tasks.* The Music Information Retrieval community[4] has created a competition – MIREX[5] – to test and compare algorithms for a variety of tasks. The MuMe community should add tasks to this collection.

[4]http://www.ismir.net/.
[5]http://www.music-ir.org/mirex/wiki/MIREX_HOME.

—*Shared Datasets.* Similar to how the MIR community has common datasets, the MuMe community should build and share common corpora and representations. Such shared datasets would enable progress on benchmarking tasks.
—*Deployment and Accessibility.* The MuMe community is a diverse collection of composers, sound designers, and scientists; as a result of this diversity, MuMe systems have tended to be idiosyncratic, non-idiomatic, and *ad hoc* [Merz 2014]. Sharing and replication of results has been difficult, and efforts should be made to share existing systems and collate libraries compiling the best algorithms in the field.
—*Standardization and interoperability*. A consequence of the above is that the various systems developed are typically not interoperable. Standardization of MuMe interfaces are desirable; some efforts at collaborative MuMe systems are underway [Bown et al. 2015].
—*Real-World Applications.* There is a plethora of opportunities to place current systems within the real world, commercial applications. Despite the obvious potential, only few applications of MuMe are reaching the video game industry, or electronic dance music production. A clearer and more visible presentation of MuMe would help engage the industry further.
—*Community Structure and International Organization*. We count MuMe practitioners by the hundreds, likely thousands, spanning the artistic, academic, and industrial worlds. While some umbrella forums do exist, the formation of an international organization providing structure and visibility to the field would be beneficial.

Meanwhile, the field is growing and is well placed to benefit from and contribute to a large number of current technological development that are the trends of our time. These include: big data [Mayer-Schönberger and Cukier 2013], since artists want to extract knowledge from cultural repositories; deep learning [Arel et al. 2010] because we need to extract meaningful features and structures that cannot be easily formulated from art manuals; online collaboration [Garrison 2006] in order to better share creative resources efficiently; better coding and creative computing environments – including commercial products such as Max for Live – and ecosystems of practice; human-computer interaction (HCI) applied to MuMe; and new application areas, such as interactive music installations, Live Algorithms in Music [Blackwell 2009], automated remixing [Vihavainen et al. 2011], film composition tools [Sorensen and Brown 2008], music education tools [Nilsson 2003], and copyright free generation [Nierhaus 2009], to name a few.

6. CONCLUSION

With hundreds if not thousands of academic publications and equally numerous real-world applications, MuMe is a thriving field of research that speaks to the trends and challenges of our time. The myriad of approaches mirrors the variety of researchers, both artistic and scientific. Its simplest definition – automation of aspects of the musically creative process – can be as simple as randomisation of selected parameters, to the pursuit of a creative system that exhibits volition.

Without any claim of exhaustivity, this short introduction aims to provide basic definitions and salient distinctions that structure MuMe, a field that spans the diverse dimensions of the computer simulation of creative musical tasks.

REFERENCES

K. Agres, J. Forth, and G. Wiggins. 2016. Evaluation of musical creativity and musical metacreation systems. *ACM Computers Comput. In Entertainment, Special Issue: Musical Metacreation.* 14, 2.

I. Arel, D. Rose, and T. Karnowski. 2010. Deep machine learning-a new frontier in artificial intelligence research. *IEEE Computational Intelligence Magazine* 5, 4, 13–18.

G. Assayag, C. Rueda, M. Laurson, C. Agon, and O. Delerue. 1999. Computer-assisted composition at IRCAM: From patchwork to openmusic. *Computer Music Journal* 23, 3, 59–72.

P. Beyls. 2007. Interaction and self-organisation in a society of musical agents. In *Proceedings of ECAL 2007 Workshop on Music and Artificial Life (MusicAL 2007)*.

D. Birchfield. 2003. Generative model for the creation of musical emotion, meaning, and form. In *Proceedings of the 2003 ACM SIGMM Workshop on Experiential Telepresence*. 99–104.

T. Blackwell. 2009. Live algorithms. *Dagstuhl Seminar Proceedings*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.

M. A. Boden. 2004. *The Creative Mind: Myths and Mechanisms*. Psychology Press.

S. Bosley, P. Swire, and R. Keller. 2010. Learning to create jazz melodies using deep belief nets. In *Proceedings of the 1st International Conference on Computational Creativity*. Lisbon.

O. Bown. 2012. Generative and adaptive creativity: A unified approach to creativity in nature, humans and machines. *Computers and Creativity*, Springer, Berlin, 361–381.

O. Bown, B. Carey, and A. Eigenfeldt. 2015. Manifesto for a musebot ensemble: A platform for live interactive performance between multiple autonomous musical agents. In *Proceedings of the International Symposium of Electronic Art*.

J. Bradshaw, R. Hoffman, D. Woods, and M. Johnson. 2013. The seven deadly myths of "autonomous systems". *IEEE Intelligent Systems* 3, 54–61.

D. Brown. 2012. Mezzo: An adaptive, real-time composition program for game soundtracks. In *Proceedings of the AIIDE 2012 Workshop on Musical Metacreation*. 68–72.

J. Charnley, A. Pease, and S. Colton. 2012. On the notion of framing in computational creativity. In *Proceedings of the 3rd International Conference on Computational Creativity*. 77–82.

P. Chordia and A. Rae. 2010. Tabla gyan: An artificial tabla Improviser. In *Proceedings. of the 1st International Conference on Computational Creativity*. Lisbon.

K. Collins. 2008. *Game Sound: An Introduction to the History, Theory and Practice of Video Game Music and Sound Design*. MIT Press.

N. Collins. 2008a. Reinforcement learning for live musical agents. In *Proceedings of the International Computer Music Conference*. Belfast.

N. Collins. 2008b. The Analysis of Generative Music Programs. *Organised Sound*, 13.

D. Conklin. 2003. Music generation from statistical models. In *Proceedings of the AISB 2003 Symposium on Artificial Intelligence and Creativity in the Arts and Sciences* (Aberystwyth, Wales), 30–35.

D. Cope. 1996. *Experiments in Musical Intelligence*. Vol. 12. Madison, WI: AR editions.

D. Cope. 1997. The composer underscoring environment: Cue. *Computer Music Journal* 21, 3, 20–37.

N. Cross. 2006. *Designerly Ways of Knowing*. Springer London, 1–13.

P. Dahlstedt and P. McBurney. 2006. Musical agents: Toward computer-aided music composition using autonomous software agents. *Leonardo* 39, 5, 469–470.

D. Damschroder and D. Williams. 1990. *Music Theory from Zarlino to Schenker: A Bibliography and Guide* (No. 4). Pendragon Press.

G. Dickie. 1985. Evaluating art. *The British Journal of Aesthetics* 25, 1, 3–16.

A. Dorin. 2001. Generative processes and the electronic arts. *Organised Sound* 6, 1, 47–53.

S. Dubnov, G. Assayag, O. Lartillot, and G. Bejerano. 2003. Using machine-learning methods for musical style modeling. *Computer* 36, 10, 73–80.

M. Edwards. 2011. Algorithmic composition: computational thinking in music. *Communications of the ACM* 54, 7, 58–67.

A. Eigenfeldt. 2008. Emergent rhythms through multi-agency in max/msp. *Computer Music Modeling and Retrieval - Sense of Sounds*. Springer Berlin, 368–379.

A. Eigenfeldt, J. Bizzocchi, M. Thorogood, and J. Bizzocchi. 2015. Applying valence and arousal values to a unified video, music, and sound generative multimedia work. In *Generative Art Conference*, Venice.

A. Eigenfeldt, O. Bown, P. Pasquier, and A. Martin. 2013. Towards a taxonomy of musical metacreation: Reflections on the first musical metacreation weekend. In *Proceedings of the Artificial Intelligence and Interactive Digital Entertainment (AIIDE'13) Conference* (Boston, MA).

A. Eigenfeldt and P. Pasquier. 2010. Realtime generation of harmonic progressions using controlled Markov selection. In *Proceedings of ICCC-X-Computational Creativity Conference* (Lisbon, Portugal). 16–25.

A. Eigenfeldt and P. Pasquier. 2011. Negotiated content: Generative soundscape composition by autonomous musical agents in coming together: Freesound. In *Proceedings of the 2nd International Conference on Computational Creativity* (Mexico City), 27–32.

A. Eigenfeldt, M. Thorogood, J. Bizzocchi, and P. Pasquier. 2014. MediaScape: Towards a video, music, and sound metacreation. *CITAR Journal* 6, 1.

G. Englert. 1989. Automated composition and composer automation. *The Music Machine,* C. Roads (Ed.). MIT Press, Cambridge, MA, 131–142.

P. Galanter. 2003. What is generative art? Complexity theory as a context for art theory. In *Proceedings of the 6th Generative Art Conference*. Milan.

R. A. Garcia. 2001. *Automatic Generation of Sound Synthesis Techniques*. Ph.D dissertation. MIT.

D. Garrison. 2006. Online collaboration principles. *Journal of Asynchronous Learning Networks* 10, 1, 25–34.

R. Groves. 2013. Automatic harmonization using a hidden semi-Markov model. In *Proceedings of the Artificial Intelligence and Interactive Digital Entertainment Conference* (Boston, MA).

E. Handelman, A. Sigler, and D. Donna. 2012. Automatic orchestration for automatic composition. In *Proceedings of the AIIDE 2012 Workshop on Musical Metacreation*. 43–48.

S. Hedges. 1978. Dice music in the eighteenth century. *Music & Letters*, 180–187.

A. Kirke and E. Miranda. 2009. A survey of computer systems for expressive music performance. *ACM Computing Surveys (CSUR)*, 42, 1, 3.

G. Koenig. 1983. Aesthetic integration of computer-composed scores. *Computer Music Journal* 7, 4, 27–32.

G. Lewis. 2000. Too many notes: Computers, complexity and culture in voyager. *Leonardo Music Journal* 10, 33–39.

M. Macret and P. Pasquier. 2013. Automatic tuning of the OP-1 synthesizer using a multi-objective genetic algorithm. In *Proceedings of the 10th Sound and Music Computing Conference (SMC-2013)* (Stockholm, Sweden). Logos Verlag, Berlin, 614–621.

M. Macret and P. Pasquier. 2014. Automatic design of sound synthesizers as pure data patches using coevolutionary mixed-typed Cartesian genetic programming. In *Proceedings of the Conference in Genetic and Evolutionary Computation (GECCO-2014)*. Vancouver, Canada, 309–316, ACM Press, 2014.

B. Manaris, D. Johnson, and Y. Vassilandonakis. 2013. Harmonic navigator: A gesture-driven, corpus-based approach to music analysis, composition, and performance. In *Procedings of the 9th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* (Boston, MA), 67–74.

A. Martin, C. Jin, and O. Bown. 2011. A toolkit for designing interactive musical agents. In *Proceedings of the 23rd Australian Computer-Human Interaction Conference*. 194–197.

J. B. Maxwell, A. Eigenfeldt, and P. Pasquier. 2012. MANUSCORE: Music notation-based computer-assisted composition. In *Proceedings of the International Computer Music Conference*. 357–364.

V. Mayer-Schönberger and K. Cukier. 2013. *Big Data: A Revolution that Will Transform How We Live, Work, and Think*. Houghton Mifflin Harcourt.

E. Merz. 2014. Implications of ad hoc artificial intelligence in music. In *Proceedings of the Artificial Intelligence and Interactive Digital Entertainment Conference* (Raleigh, NC).

D. Moffat and M. Kelly. 2006. An investigation into people's bias against computational creativity in music composition. *Assessment* 13, 11.

D. Murray-Rust, A. Smaill, and M. Edwards. 2006. MAMA: An architecture for interactive musical agents. *Frontiers in Artificial Intelligence and Applications* 141, 36.

National Academies. 2004. Facilitating interdisciplinary research. *The National Academies Press*, Washington DC.

G. Nierhaus. 2009. *Algorithmic Composition: Paradigms of Automated Music Generation.* Springer Science & Business Media.

B. Nilsson. 2003. "I can always make another one!"–Young musicians creating music with digital tools. *Musicianship in the 21st Century: Issues, Trends and Possibilities*. Sydney, Australian Music Centre.

F. Pachet. 2003. The continuator: Musical interaction with style. *Journal of New Music Research* 32, 3, 333–341.

F. Pachet and P. Roy. 2001. Musical harmonization with constraints: A survey. *Constraints* 6, 1, 7–19.

F. Pachet and P. Roy. 2014. Non-conformant harmonization: The real book in the style of take 6. In *Proceedings of the International Conference on Computational Creativity* (Ljubljana).

M. Peter. 2009. Milieus of creativity: The role of places, environments, and spatial contexts. In *Milieus of Creativity*. Springer Netherlands, 97–153.

J. Reiss. 2011. Intelligent systems for mixing multichannel audio. In *Proceedings of the 17th International Conference on Digital Signal Processing (DSP)*. 1–6.

J. Reiss and E. Perez Gonzalez. 2008. An automatic maximum gain normalization technique with applications to audio mixing. In *Proceedings of the Audio Engineering Society Convention* 124. Audio Engineering Society.

G. Ritchie. 2007. Some empirical criteria for attributing creativity to a computer program. *Minds and Machines* 17, 1, 67–99.

S. Russell and P. Norvig. 2009. *Artificial Intelligence: A Modern Approach*, third edition, Prentice-Hall.

M. Salganik, P. Dodds, and D. Watts. 2006. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311, 854–856.

A. Sarwate and R. Fiebrink. 2013. Expansion on description-based design of melodies. In *Proceedings of the 9th Artificial Intelligence and Interactive Digital Entertainment Conference*. Boston.

R. Saunders and O. Bown O. 2015. Computational social creativity. *Artificial Life* 21, 3, 366–378.

M. Schedel and A. Rootberg. 2009. Generative techniques in hypermedia performance. *Contemporary Music Review* 28, 1, 57–73

H. Simon. 1960. *The New Science of Decision Making*. Harper and Row, New York.

I. Simon, D. Morris, and S. Basu. 2008. MySong: Automatic accompaniment generation for vocal melodies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 725–734.

J. Sloboda. 1988. *Generative Processes in Music: The Psychology of Performance, Improvisation, and Composition*. Clarendon Press.

A. Sorensen and A. Brown. 2008. A computational model for the generation of orchestral music in the Germanic symphonic tradition: A progress report. In *Proceedings of Sound: Space - The Australasian Computer Music Conference*. Sydney. ACMA, 78–84.

M. Thorogood, P. Pasquier, and A. Eigenfeldt. 2012. Audio metaphor: Audio information retrieval for soundscape composition. *Sound and Music Computing Conference*, Copenhagen.

B. Truax. 1980. The inverse relation between generality and strength in computer music programs. *Journal of New Music Research* 9, 1, 49–57.

B. Truax. 1985. The PODX system: Interactive compositional software for the DMX-1000. *Computer Music Journal* 29–38, 1985.

S. Vihavainen, S. Mate, L. Seppälä, F. Cricri, and I. Curcio. 2011. We want more: Human-computer collaboration in mobile social video remixing of music concerts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 287–296.

I. Wallis, T. Ingalls, E. Campana, and J. Goodman. 2011. A rule-based generative music system controlled by desired valence and arousal. In *Proceedings of 8th International Sound and Music Computing Conference* (Padova).

Ian Whalley. 2009. Software agents in music and sound research/creative work. *Organised Sound* 14, 156–167.

R. Whorley, G. Wiggins, C. Rhodes, and M. Pearce. 2010. Development of techniques for the computational modelling of harmony. In *Proceedings of the 1st International Conference on Computational Creativity*. Lisbon.

M. Wooldridge. 2009. *An Introduction to Multiagent Systems*. Wiley.

R. Wulfhorst, L. Flores, L. Nakayama, C. Flores, L. Alvares, and R. Viccari. 2001. An open architecture for a musical multi-agent system. In *Proceedings of Brazilian Symposium on Computer Music*.

G. Yannakakis and H. Martínez. 2015. *Ratings are Overrated! Frontiers in ICT*, 2.

M. J. Yee-King. 2011. Automatic sound synthesizer programming: Techniques and applications (Doctoral dissertation, University of Sussex).

M. Young. 2008. NN music: Improvising with a ‘living’ computer. *Computer Music Modeling and Retrieval – Sense of Sounds*. Springer Berlin, 337–350.

Received September 2015; revised April 2016; accepted July 2016