Towards a Taxonomy of Musical Metacreation: Reflections on the First Musical Metacreation Weekend

Arne Eigenfeldt
School for the Contemporary Arts,
Simon Fraser University,
Vancouver, Canada
arne_e@sfu.ca

Oliver Bown
Design Lab,
University of Sydney,
Sydney, Australia
oliver.bown@sydney.edu.au

Philippe Pasquier
Interactive Arts & Technology,
Simon Fraser University,
Surrey, Canada
pasquier@sfu.ca

Aengus Martin
Faculty of Engineering,
University of Sydney,
Sydney, Australia
aengus.martin@sydney.edu.au

Abstract

The Musical Metacreation Weekend (MuMeWe), a series of four concerts presenting works of metacreation, was held in June 2013. These concerts offered an opportunity to review how different composers and system designers are approaching generative practices. We propose a taxonomy of musical metacreation, with specific reference to works and systems presented at MuMeWe, in an effort to begin a discussion of methods of measuring metacreative musical works. The seven stages of this taxonomy are loosely based upon the autonomy of the system, specifically in its ability to make high-level musical decisions during the course of the composition and/or performance. We conclude with a discussion of how these levels could be interpreted, and the potential difficulties involved in their specification and terminology.

Introduction

The Musical Metacreation Weekend (MuMeWe)[1] was held in Sydney, Australia, on June 15th and 16th 2013. A series of four concerts were presented that featured peer-reviewed works: these concerts presented a range of music generated by systems with varying degree of autonomy. We feel the weekend festival provides a useful measure on the current state of musical metacreation systems[2]; furthermore, these events will allow us to provide a brief overview and discussion of several of the systems, and conclude with a proposed taxonomy of metacreative systems. Such a taxonomy will allow the comparison of systems without regard to perceived musicality, complexity, or traditional autonomy, concepts we feel are related to, but separate from, describing the level of metacreativity of a system.

Musical metacreation is concerned with autonomy and agency in composition and performance: simply put, having the computer as an active participant in the creative process. Several of the systems presented at MuMeWe were within the paradigm of improvisation; interestingly, improvisation provides a useful parallel to that of musical metacreation, in that there is a varying relationship between intent (composition) and practice (performance). In the case of improvisation, one needs to consider what the relationship is between the underlying composition, and its realisation – the relationship between different performances of the same work – and thus the influence of the performer/improviser. Similarly, in metacreation, we need to consider the relationship between the composer/system designer's direct influence on the system, and the potential differences between performances of the same work, in order to determine the system's autonomy.

Previous Work

Benson provides his taxonomy of improvisation (2003), in which he specifies sixteen different "shades" of improvisation, each of which present a relationship between a given original composition and its realisation. These begin with a "filling-in" of details that are not notated in the score, such as tempi, timbre, and dynamics, to which he notes "no performance is possible without some form and degree of improvisation". The next eleven progress through various degrees of freedom in relationship between performer and original composition, with the thirteenth being one in which the relationship between the performance and the original composition is evident only to the performer and composer of the piece, a model that one tends to find most often in contemporary non-jazz improvisation. Three additional forms of improvisation are provided that Benson considers "compositional", such as Mozart "improvising" upon opera buffa in *Così Fan Tutti*.

Boden and Edmonds (2009) provide a taxonomy of generative art, separating various forms into categories in an

[1] http://www.metacreation.net/mumewe2013/index.php
[2] See Bown et al. (2013) for a detailed discussion of the event.

effort to "help toward a critical framework" of new forms in art technology. The authors provide an interesting discussion in regards to their definition of computer-generated art (CG-art), in which they first define as artwork resulting from autonomous systems that allow for "zero interference from the human artist"; however, they immediately acknowledge that such a definition is extremely strict, which ignores a prominent subclass of computer-generative art that includes interaction. Their more general classification of generative art (G-art) – artwork produced that "is not under the artist's direct control" - raises the question of what exactly "direct control" implies. It is this level of control that we will attempt to categorise in this paper: while some metacreative systems, as demonstrated by those presented at MuMeWe, may strive towards LeWitt's concept of a "machine that makes the art", where "all of the planning and decisions are made beforehand and the execution is a perfunctory affair" (LeWitt in Boden 2007), it turns out that having some direct control over the system is still an important element of metacreation. This reflects Boden and Edmonds statement that such interactivity is the result of artists who "rely on their intuitive judgment to make selections during an artwork's evolution."

Bown and Martin (2012) provide a discussion on the nature of autonomy in software systems, and suggest that building systems that attain a high level of autonomy may be trivial in comparison with building systems that produce output that is contextually relevant. Referencing Newell et al. (1963), Bown and Martin suggest that the requirement for usefulness should never be abandoned in favour of autonomy. Furthermore, they suggest that a measure of a system's agency, complexity, and intelligence are an equally important measure of the system's autonomy.

Metacreative systems, as proposed by Eigenfeldt et al. (2012), and as demonstrated at the MuMeWe, tend to fall into two distinct genres: improvisational systems (i.e. online), and composition systems (i.e. offline). In the case of the former, it makes sense to judge such systems in relation to how a system interacts with a live performer: in Bown and Martin's terms, the complexity, intelligence, agency, and autonomy of the system in response to the live input. In the case of offline systems, it makes sense to judge such systems in relation to an underlying compositional idea: to what extent does the system produce its own musical structure and details, and how much does it depend upon the user to move it forward?

Determining the autonomy of either system is problematic, as it would require multiple exposures to discover the true range of a system's output. In fact, Bown and Martin conclude that "a singular measure of autonomy itself may not actually be as interesting as the various measures that can be considered to constitute it", and suggest the potential for a second software system to probe the creative system in order to get "a better sense of the system's autonomy in a range of contexts".

Lastly, the amount of control an artist has over the system may be difficult to ascertain in non-performative situations. As Collins (2008) points out, the established notion of algorithmic composition, as discussed in Pearce, Meredith and Wiggins (2002), place the computer in the role of composer's assistant in an offline system. Collins, in suggesting methods of analysis for generative music, clearly outlines what he considers to be "generative music" – in Boden and Edmonds terms, CG-music – and excludes from his discussion art that is algorithmic, such as live-coding, that allows for human intervention during execution. Our notion of musical metacreative systems, and the taxonomy suggested here, is less rigid, and acknowledges the continuum between strictly reactive or inert music software and fully generative ones. This continuum goes through the various stages of computer-assisted creativity, as detailed below.

Descriptions of selected pieces

In curating the concerts, it was extremely difficult to compare various systems and representative performances, without regard to the artistic output of the systems. While some submissions were artistically strong, they were only vaguely metacreative; others were, perhaps, less successful in their artistic results, but much more autonomous. As such, we are presenting and describing a few of the systems/compositions performed at MuMeWe as examples.

Improvising Algorithms

The first concert presented systems in reaction to a live improvising performer. Five of the systems interacted with an acoustic instrument (Ollie Bown: piano, Oliver Hancock: piano, Bill Hsu: clarinet, Ben Carey: saxophone, Alesandro di Scipio: trumpet), while Andrew Brown/Toby Gifford's system interacted with a MIDI synthesizer. As a result, the latter piece was the most pitch-based in its interactions, most likely due to the availability of exact pitch information from its MIDI input.

Bown and Hsu were present, and operated their own software; in both cases, they interacted with their software during performance. Carey and Brown performed as musicians with their own systems. Hancock and di Scipio sent their software, and required someone else to rehearse with the performers, and operate the software during performance: in both cases, it was not clear to whether any interaction between operator and system occurred during performance.

While each system was unique, and each successfully exhibited "musical" interactions – which, in itself, were an indication of the current state of maturity in improvisational metacreative systems – we will choose Bown's system as an example to describe in more detail. One characteristic that separated Bown's work from the other improvisational systems was his avoidance of the record–playback paradigm demonstrated by the other systems. The fact that the other systems were not playing a distinct timbre or instrument, and instead played back recordings made of the live instrument, made the perception of their identity and autonomy more difficult to determine: playing a processed version of the input blends it in a way that is reminiscent of effect processors, and of the turn-taking mechanisms of reflective interaction, as established and studied in the context of the Continuator (Pachet 2003) and MIROR (Addessi 2013) platforms.

It is informative to compare two pieces. Bown's system, *Zamyatin*, is "comprised of an audio analysis layer, an inner control system exhibiting a form of complex dynamical behaviour, and a set of output modules that respond to the patterned output from the dynamical system. The inner systems consists of a Decision Tree that is built to feed back on itself, maintaining both a responsive behaviour to the outside world and a generative behaviour, driven by its own internal activity" [Bown, program notes].

Hancock's system is "a group of computer agents, each one acting as a delay, and basing its playback on the rhythms of the live improviser, and all the other agents in the system. Together they behave like chorusing insects or frogs, with some degree of irregularity and unpredictability. Broadly the system matches the activity of the improviser, but it can blossom and race unexpectedly, carry on alone in a lull, or leave the improviser to play in relative isolation." [Hancock, program notes].

Bown's approach consciously avoids designing behaviour into the core of the system in a preconceived manner. A general purpose system is instead created and given basic functionality through an evolutionary goal. However, the musical work is then additionally developed around this core system. In contrast, Hancock follows the paradigm of resampling the performer's input and applying different composed responses to the playback of these samples. In both cases, the system in no way directly mimics or tracks the performer, and is understood as an independent agent that stimulates the performer through interaction. In Bown's case this is extreme, and the performer can be understood as perturbing the output of an otherwise reasonably stoic system, whereas in Hancock's case, musical materials (timbral, melodic, rhythmic) from the performer become part of the system's response and the interaction has the form of a rich loop or echo system.

Formally, in both cases, the performer does not directly control the system – thus, a lack of heteronomy – the system is neither fixed nor highly random, and is in part influenced by the performer. In both cases the system has low predictability in its short-term decisions but high-predictability in its long-term style.

Other performances included Oliver Hancock, echo-system; Bill Hsu, *Figment*; Benjamin Carey, *_derivations*; Andrew Brown and Toby Gifford, *Unity in Diversity*; Agostino di Scipio, *Modes of Interference / 1*.

Algorave

A second event was held in a club venue, focusing upon generative dance music. The term Algorave specifically relates to the notion of algorithmically generated dance music, an event begun by Alex MacLean and Nick Collins circa 2012. The event began with a 30 minute set of generated EDM tracks in the Breakbeat style by Chris Anderson's *GEDMAS* system (Anderson et al. 2013). As the tracks were pre-generated, there was no performance aspect, and thus no interaction. Anderson's set was generated offline, using a corpus of 24 transcribed Breakbeat tracks. The system was autonomous in that it allowed for no user interaction during generation; however, timbres were hand-chosen once the tracks were generated.

Following the presentation of pre-generated audio, two live-coding performances were presented. Bell's performance of *Live with Conductive* demonstrated "post-dubstep bass music", which was followed by the live-coding duet of Sorensen-Swift, who provided a mellower and completely tonal set of electro and minimal house. Although the live-coders used different environments – Bell used *Haskell*[3], while Sorenson/Swift used *Extempore*[4] – both triggered gestures and events that were typed by hand. In the case of the latter, the laptops shared a common sync, which allowed them to maintain a common pulse.

Other works presented included Sick Lincoln, *Algoravethm for 31770*; Arne Eigenfeldt, *Breakbeats by GESMI.*

Composed Process / Studies in Autonomy

Two additional concerts were held in more traditional concert venues, and provided a somewhat more formal atmosphere. Roger Dean's *Serial Identities* began with Dean improvising atonally at the piano; he paused to initiate his software that began generating similar atonal/serial music that triggered a sampled piano module. The system was reactive, rather than interactive, in that Dean reacted to the generated material, but the system was not "listening" to his performance.

Kenneth Newby's *The Flicker Generative Orchestra* was a fixed media example of "the application of a variety of generative processes that result in music that resembles that of large symphony orchestras playing in a contempo-

[3] http://www.haskell.org/haskellwiki/Haskell
[4] http://benswift.me/extempore-docs/

rary idiom" (Newby, program notes). The system generated individual musical lines that were convincing and believable in their specific representations. For this particular rendition, Newby selected a specific pitch and rhythmic set for the system to explore.

David Kim-Boyle presented *Line Studies no. 1*, a real-time generated graphic score which was immediately interpreted by a software improviser, performed on a sampled piano. The score itself was visually beautiful – a series of "concentric rings segmented into arcs of varying lengths and thicknesses". During the performance, the arcs are continually "drawn, extended, overwritten, rotated, and erased. Each arc denotes a pitch that is performed by the pianist, with arc lengths and arc thicknesses corresponding to pitch durations and dynamics respectively" (Kim-Boyle, program notes). Although the piano was interpreting the score, the unusual presentation suggested the reverse (to the first author): a visualisation of an improvised musical performance.

Nomos ex Machina was an example of Shlomo Dubnov's *Audio Oracle* in action, as it reacted to a live performer and responded with audio recordings made during the performance. Unfortunately, the system's ability to interact with the performer through its complex processes of time mapping wasn't clear, as the performer choose to improvise in a consistent "full-on" manner that ignored the system's responses.

Other works presented included Brian Kelly, *Phi*; Micquel Parera Jaques, *neix_2013a*; Ivan Zavada, *Algosonic*; Anisha Thomas, *Kundaliktra*; Kerry Hagan, *Morphons and Bions*; Benjamin O'Brien, *Densité*.

Installations

Eigenfeldt's *Roboterstück* was one of three installations that ran continuously during the weekend. Eigenfeldt's work played every 15 minutes, when agents would "wake-up" and begin interacting. The agents sent MIDI information to the *NotomotoN* mechanical musical instrument, which struck a variety of objects, including clay flower-pots, metal bars, plastic blocks, and a drum. The agents negotiated a texture – based upon opposing metrics of slow-fast, loud-soft, dense-sparse, rhythmic-arrhythmic – then improvised within that texture until a majority of the agents "got bored", at which point a new texture would be negotiated. Once the same texture was performed three times, the performance ended. Each performance lasted between three to six minutes.

Other installations included Miles Thorogood and Philippe Pasquier, *Audio Metaphor*; Fredrik Olofsson, *low-life*.

On Reflection

The weekend provided the opportunity to witness state-of-the-art metacreative systems, and demonstrated a wide variety of approaches to the notion of autonomous creative systems. Without commenting on the artistic merit of any of the compositions and performances, the experience did suggest a basic distinction between the approaches.

Musical metacreation, borrowing from the more general definition of computational creativity, is the idea of endowing machines with creative behavior. Musical metacreation can be divided into a number of canonical problems:

1. **Composition** – being the process of creating a series of performance instructions for musical performers, (i.e. a score);
2. **Interpretation** – being the process of performing a musical composition and producing an audio rendering;
3. **Improvisation** – which combines (1) and (2) in real-time performance;
4. **Accompaniment** – being the process of following a live performer in an accompanying role, possibly performing pre-composed music;
5. **Continuation** – being the process of continuing a given musical input in the same style.

The MuMeWe provided two examples of #1 in the Algorave (albeit, the directions for performance were Ableton Live scores). While #2 was not presented on its own, it was encompassed within several instances of #3 through system performance. #4 was tangentially demonstrated within the improvisation systems; in fact, some of the performances were more suggestive of accompaniment, rather than equal interaction, due mainly to the performer's choice to remain as the foreground instrument. #5 was mostly demonstrated by the Brown/Gifford piece, itself inspired by the reflective interaction paradigm of Pachet (2003).

A Taxonomy of Musical Metacreation

The taxonomy presented is meant as a classification system for metacreative systems in terms of their relationship to the human designer/composer's control over the final musical result, which can also be interpreted as the level of musical autonomy. Just as there are different levels of improvisation (and different relationships between a performer's effect upon the final result in comparison to the original "composition"), there are different degrees that a system has control over the final musical artifact, whether it is generated live (online) or fixed (offline). This can be reduced to how much creative decision-making was left to the system. Another way of looking at it is how much influence is required from a human in order for a metacreative system to perform musically[5].

Importantly, a human designer/composer's sense of control over the musical result of an action depends strongly on his/her understanding of the mapping between the input and output of the system. A mapping that is completely deterministic but complex might deny a sense of control and therefore exhibit a significant level of metacreation. Thus, the level of metacreation is at least in part something that is attributed to a system by a user or observer, and not something that might be directly measurable. This is consistent with Paine's view that interactive MuMe systems require "a level of cognition" (Paine, 2002), i.e. something clearly attributed to it by a human observer. It is also consistent with Jòrda (2007), who links interactivity in MuMe systems to the potential of a system for producing unexpected results, whether through `non-linearity, randomness or ungraspable complexity'.

Our taxonomy consists of seven levels, and is ordered from the least autonomous to the most. Later levels do not need to exhibit properties of earlier levels (although they could): instead, it should be understood that earlier levels are incapable of exhibiting defining properties of later levels.

An eighth level (level 0) can be inferred as musical compositions that contain no metacreativity; however, note that such a level does not imply fixed-media works, as autonomy can exist in the creation, separate from the performance/execution.

1. Independence: the use of any process on a musical gesture that is beyond the control of the composer.

The most basic creative influence a system can have on a musical work is the application of a process upon a gesture[6], in which the process itself can vary with each application; in other words, a process that involves some randomisation, for example. In a studio environment, such processes can be auditioned and selected, while in performance, a certain amount of trust will be placed upon the system to apply the process in a predictable way: in other words, delegating some creative responsibility to the system.

Examples of processes that fall into this category would include:

-complex signal processing that involve random parameters, in which the final result is heard as a single gesture (i.e. reverberation);

-a random selection of playback speed (tempo);

-using an algorithmic process to alter volume and/or onset data in a MIDI sequencer. (Note that the "humanisation" of pre-generated sequences through such variation can be considered within the *interpretation task* of a metacreative system.)

Such variation can be random (discrete distributions, or continuous), constrained random (within a preset range), constrained random whose range is varied by a performer, random variation of a preset value, or a process – such as an LFO based upon a complex mathematical function – whose output may not be entirely predictable. It can also be stochastic, or statistically learned from a corpus of human interpretations.

2. Compositionality: the use of any process to determine the relationships between pre-defined musical gestures.

The previous level applied to a single gesture; this level is concerned with the relationship between two gestures and/or processes that are in themselves fixed. This may involve initiating multiple layers of pre-generated material, triggering pre-recorded material, or even initiating complex signal processing in relation to a live performance in which the process is heard as a separate gesture to the original (i.e. complex delays). Initiating events and/or processing in relation to a fixed score (i.e score-following) would be another example: from the perspective of predictability and causality, whether events are initiated by the press of a button, or waiting to be triggered by an event in a score, the results are the same.

Similarly, aspects of live-coding (Bell's *Live with Conductive*, Sorenson/Swift's *Palle in Aria*) in which direct sequences are initiated, are likewise catagorised here.

In all these cases, the exact relationship between individual musical lines, while deterministic (due to their fixed material), possibly remain unknown when initiated, and are thus considered metacreative.

3. Generativity: the generation of musical gestures.

Similar to #2, albeit with the addition of an algorithm that generates and/or substantially varies pre-generated sequences; any reactive system that requires input in order to function.

An example would be triggering processes that contain pitch/rhythm generation algorithms; another example would be to trigger gestures that are undefined (i.e. generative) in their pitch/rhythm/volume in response to a performer's action. Similar to #2, live systems that are reactive, in that they use live input to generate material, are categorised here for the same reason score-followers are placed in #2.

[5] By "musically", we mean as done by a human, either improvisational or compositional, in a way that would be considered "musical".

[6] By "gesture", we may any musical idea that has a beginning and an end. This includes not only traditional "phrases", but also "performer-less" notions, such as LFO drones.

Feedback systems are contained here, in that the gesture is indeterminate in its details, yet require the proper environment to be instantiated.

Examples from the MuMeWe described include Kim-Boyle's *Line Studies no.1*, and Dean's *Serial Identities.* Additionally, some aspects of live-coding (Bell's *Live with Conductive*, Sorenson/Swift's *Palle in Aria*) are included here, in which sequences were initiated that included random, or stochastic, selection from a constrained set of pitches.

4. Proactivity: system/agents that are able to initiate their own musical gestures.

As opposed to #3, there no is direct trigger that an agent awaits before beginning. Thus, this is the first notion of autonomous agent-based behaviour. Agents are not merely reactive, requiring input in order to function, but can generate material of their own accord; this includes Blackwell's autonomous musical agent, which "goes beyond mere automation in determining responses to musical input, by deriving novel but relevant responses" (Blackwell, in Bown). In the agent literature (Wooldridge 2009), this is know as proactivity, and is a defining characteristic of autonomy.

Examples of proactive systems would be interactive systems that have an independent response to a performer (i.e. Bown's *Zamyatin*, as well as George Lewis' *Voyager*); another example would be a multi-agent system (i.e. Eigenfeldt's *Roberterstück*).

5. Adaptability: a) Agents behave in different ways over time due to their own internal evolution; b) agents interact and influence one another.

This level is a subtle shift in autonomy in comparison to #4 – the primary difference is the amount of change the agents are capable of achieving proactively, in a restricted period of time (i.e. the composition). In some agent-based works, an agent's behaviour is governed by an internal state that can be altered by the composer/performer so as to achieve different reactions in different situations; as such, human interaction is still required. In systems within this fifth level, agents determine *when* to alter their behaviour, and *how*, proactively. This is separate from defined parametric change, in which gestures may alter over time based upon a predefined applied function.

5b includes agent interaction, thus completing the multi-agent model, in which agents are not only autonomous (#4 - proactivity), but they are also social, and their interactions with other agents determine how they operate in the future. However, social agents do not, on their own, guarantee that they will evolve their actions over time *in a significant way*. As such, agent interaction is separate from agent autonomy, and does not warrant its own level.

An example of 5a would be a generative system that generates its own musical structure (i.e. Newby's *The Flicker Generative Orchestra*). Offline examples would include Anderson's *GEDMAS*, which produces complete electronic dance music tracks based upon a restricted corpus.

Examples of 5b would be evolutionary systems not heard at MuMeWe, such as McCormack's *Eden,* and Bown's *Sonic Eco-System*, in which the system produces fundamentally different output over time. Another simpler example would be a physically modeled system that undergoes a single-phase transition during its lifetime.

6. Versatility: agents determine their own content without predefined stylistic limits.

As opposed to #5, which essentially generates a variation of the same composition each time, with variation determined by user settings (i.e. a user-set pitch structure, a fixed relation between agents, a given corpus, or any kind of formal template), here the agents would generate truly different compositions with each generation. The use of (and need for) templates for formal construction have been eliminated. In terms of Boden's exploration vs. transformation, this system would offer the potential for transformation of the creative space. No such system was presented at the MuMeWe.

7. Volition: agents exhibit volition, deciding when, what, and how to compose/perform.

A truly free-standing creative system, independently capable of deciding whether it wants to create, and eventually why it would want to do so (which is the problem of intrinsic motivation); a system capable of deriving its own conceptual spaces (Gärdenfors 2004, in Bown); a system capable of autonomous critical evaluation (Galanter 2012).

Such a system, as of yet, does not exist. Human composers, on the other hand, presumably exhibit such actions whenever they create: at least truly creative artists would, one hopes. Such a system would require long-term learning, sophisticated feedback mechanisms that include peers and community, the ability to form aesthetic judgment, and the ability to derive its own motivations.

Discussion

The reader may notice that the above taxonomy does not take into account any level of complexity of a system. An extremely complex system that requires the designer to "nudge it along" during performance will remain at level 4 (proactivity), while a system that generates simple random melodies, but changes *how* those melodies are produced proactively during the composition/performance – even if only by using a randomly generated form – will be placed at level 5 (adaptability). While this may appear counterintuitive, the autonomy of the two systems is different: the latter system can produce a complete (albeit musically limited) work, while the former requires a response to be triggered and/or altered. This division is a deliberate nod to LeWitt, and his suggestion of a "machine that makes art": without the ability to generate its own forms, the machine will only be able to assist a human in making the artwork, as the human is still required to make fundamental decisions regarding the evolution of the artwork in time. While Colton et al. (2012) have created a fully autonomous poetry generation system which requires the use of form templates, we contend that such use of templates are a substitute for direct interaction: the higher levels in our taxonomy must generate their own forms autonomously.

On Proactivity

A defining aspect of many of the proposed levels is related to the system's ability to make higher-level musical decisions proactively: or "on its own". Clearly, proactivity, like autonomy, is extremely difficult to define, particularly from an external viewpoint. How can listeners determine whether a musical event was the result of serendipity in performance, or the emergence of relationships within the system? Conversely, how can a profound change in the system be recognised by the audience if such changes do not result in substantial changes within the music? We cannot assume that designers of metacreative systems can have a clear operational grasp of what it true proactivity means, and how it is applied in a range of human and machine scenarios, nor do we have clear enough methods to determine how to unambiguously make something that does something "on its own".

While we wish to avoid including the notion of complexity within our level criteria, it may need to be included in differentiating level 3 – triggered generativity – from level 4 – proactivity. As previously suggested, if a system waits for a pre-determined input, it would not be considered proactive due to its predictability and causality; but by combining a series of such triggers (i.e. a series of when() methods) with some randomness, the predictability and causality could be eliminated, and the system would seem to operate with greater complexity, and *perceived* proactivity.

In summary, an understanding of a system's ability to act "on its own" may be our deepest analytical and philosophical challenge.

On Thinking and Listening

We can wander into another conceptual minefield by suggesting that we are attempting to distinguish a system's ability to "think" (its autonomy) from its ability to "listen" to any input. Metacreative systems do not require input, and thus do not require any ability to listen; they do, however, require an ability to "think". As such, a system may display an ability to interpret live input exceedingly well, but if it cannot rise above echoing such input, will have difficulty rising above level 3 (generativity) or 4 (proactivity) on our scale. On the other hand, a system may completely ignore a live performer (and thus listen very poorly) yet function as a highly metacreative system if it can determine, proactively, when to make important musical decisions. Of course, a system can also do both; as such, it might be justified to use a cognitive model of listening, or more specifically, *musical thinking*, as proposed by Maxwell et al. (2012).

We acknowledge that it is difficult to ascribe a clear definition of "thinking", which is as hard to define as autonomy and proactivity; however, we believe a colloquial understanding of the term is useful here.

On Usefulness

Roger Dean has suggested that non-interactive systems in improvisational contexts are musically useful: "A generative algorithm might also be conceived as poorly interactive, but providing a sound environment in which other sound is improvised. Unpredictability in the algorithmic output is also a potential positive contributor to its improvisatory utility." (Dean 2003). The musical "usefulness" of any system is not accounted for in our taxonomy, whether it is Dean's application of the word, or in terms of creativity research (Newell et al. 1963). Dean's example system may provide generated material, and whether it is created in real-time or not, places it at level #3 (generativity): the requirement for the next level (proactivity) is that it decides *when* to initiate. We consider the discussion of the usefulness of a system to be separate.

On Autonomous Fitness Functions

Boden and Edmonds discuss autonomy in relation to rule-based systems versus evolutionary systems: since the latter are making decisions proactively as to how to evolve (assuming a non-interactive fitness function), such systems will be placed at level 5 (adaptability) in our taxonomy. While the creation of a fitness function is, in itself, a human decision, if the fitness function operates on its own during the composition/performance, its level of autonomy is greater than a rule-based system that requires interaction. However, the human-created fitness function will also limit it to this location: in order for such a system to be placed at level 6 (versatility), it would require the ability to develop its own non-human fitness function. Similarly, any corpus-based system that involves a fixed corpus provided by a user will limit its level in the same way.

On Large-scale Structure

Lastly, it became apparent at MuMeWe that successful large-scale musical structure is something that is the most difficult for an artist to delegate to a system (voluntarily or otherwise). This makes complete musical sense, as this is one of the most difficult aspects of music for young composers to learn[7]. Just as young (human) composers can generate short forms, so can many metacreative systems; selecting "what to do next" is an incredibly difficult musical decision to make, one that takes into account the current musical context, the past musical context, and the potential future contexts. Asking – or expecting – a computer to successfully navigate this terrain is within the nascent field of computational aesthetics: see Galanter (2012) for an overview of this field and its potential for musical metacreation.

Conclusion

Once a system is clearly placed within this taxonomy, a discussion can occur as to the complexity of its output, the level of variation in its output, and its "musicality". Clearly, some very complex and musical systems exist at the lower levels of this taxonomy, as evidenced by MuMeWe: what we propose is a method to compare the metacreative level of systems independent of their "musical maturity".

Perhaps the distinction in all of these examples is to think of each level as a principle, and consider each system in terms of whether (a) it aspires to that principle and (b) it masters that principle. In that way, if certain systems could be said to aspire to principle #6 (versatility), then we can critically examine how these systems fall short.

Lastly, it should be pointed out that there were many musical works at MuMeWe that were of a high musical quality. Artists within this nascent field will always want to create works of which they feel are artistically strong, and may feel the need to retain interactive control of their systems; scientists within the field will be drawn to creating systems that are more autonomous. The dilemma for all practitioners in the field is to balance these two goals, which should not be mutually exclusive.

References

Addessi, A. 2013. Child/machine interaction in reflexive environment: The MIROR platform. In *Proceedings of the Sound and Music Computing (SMC 2013) Conference.*

Anderson, C., Eigenfeldt, A., Pasquier, P. The Generative Electronic Dance Music Algorithmic Systems (GEDMAS). In *Proceedings of the Artificial Intelligence and Interactive Digital Entertainment (AIIDE'13) Conference*.

Benson, B. E. 2003. The improvisation of musical dialogue: a phenomenology of music. Cambridge University Press.

Boden, M., and Edmonds, E. 2009. What is generative art? *Digital Creativity* 20.1-2: 21–46.

Bown, O., and Martin, A. 2012. Autonomy in Music-Generating Systems. In *Eighth Artificial Intelligence and Interactive Digital Entertainment Conference.*

Bown, O., Eigenfeldt, A., Pasquier, P., Martin, A., and Carey, B., 2013. The Musical Metacreation Weekend: Challenges Arising from the Live Presentation of Musically Metacreative Systems. In *Proceedings of the Artificial Intelligence and Interactive Digital Entertainment (AIIDE'13) Conference*.

Collins, N. 2008. The analysis of generative music programs. *Organised Sound* 13.3: 237-248.

Colton, S., Goodwin, J., and Veale, T. 2012. Full face poetry generation. In *Proceedings of the Third International Conference on Computational Creativity,* 95–102.

Dean, R. T. 2003. *Hyperimprovisation: Computer Interactive Sound Improvisations*. AR Editions, Inc.

Eigenfeldt, A., Burnett, A., and Pasquier, P. 2012. Evaluating Musical Metacreation in a Live Performance Context. In *Proceedings of the International Conference on Computational Creativity.*

Galanter, P. 2012. Computational aesthetic evaluation: past and future. *Computers and Creativity,* 255–293. Springer Berlin Heidelberg.

Jòrda, S. 2007. Interactivity and live computer music, N. Collins and J. d'Escrivan, (eds.), *The Cambridge Companion to Electronic Music*, 89–106. Cambridge University Press, Cambridge, UK.

Maxwell, J., Eigenfeldt, A., Pasquier, P., Gonzalez Thomas, N. 2012. Musicog: A Cognitive Architecture for Music Learning and Generation. In *Proceedings of Sound and Music Computing (SMC 2012) Conference.*

Newell, A., Shaw, J. G., and Simon, H. A. 1963. The process of creative thinking, H. E. Gruber, G. Terrell and M. Wertheimer (eds.), *Contemporary Approaches to Creative Thinking,* 63–119. New York: Atherton.

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

Paine, G. 2002. Interactivity, where to from here? *Organised Sound*, 7(3): 295–304.

Pearce, M., Meredith, D., and Wiggins, G. 2002. Motivations and methodologies for automation of the compositional process. *Musicae Scientiae*, 6(2): 119–147.

Wooldridge, M. 2009. *An Introduction to MultiAgent Systems – Second Edition.* John Wiley & Sons

[7] At least of Western art music.