Musical Meta-Creation: The Discourse of Process

Andie Sigler
School of Computer Science,
McGill University
Montreal, Canada

Abstract

Musical meta-creation is an artistic and technical field in which the results (i.e. music) often cannot be empirically tested. The processes by which music is generated, however, afford a rational level of discourse. This paper suggests ways in which we might orient our discussion and our system-building in order to explore musical possibility and use computational modeling to grow into a new kind of creative musical practice.

Evaluating Musical Meta-Creation: Process and Discourse

The development of musically creative systems is both an artistic and a technical endeavor. Designers of computational music generating systems may be artists using computation as a private means of creating musical works. Other artists may consider the meta-artwork to be the piece they're presenting, itself deserving of aesthetic appraisal. Musicologists, cognitive scientists, and mathematicians may develop music generating systems in order to test hypotheses. Artists, scientists, or artist-scientists may want to model a corpus, model a composer's musical mind, develop a new style, explore the limits of musical listening, or map an open space of musical possibility.

Musical meta-creation is a technical, mathematical field, often practiced by scientists, and yet it is invested with an irreducibly informal, subjective criterion at the ouput level. *Music* can't be formally, scientifically evaluated.[1] But the professional requirement of scientists, as well as the influence of science, currently culturally ascendant, on humanists and artists, brings the desire for empirical validation methods for meta-creative systems.

The hypothesis and method of testing depends on the goal of the experiment – and the general problem of *create (good) (great) (groovy) (new) music* may have no available formal test at all. For the problem of modeling a corpus, however, the "indistinguishability test" is common (Pearce and Wiggins 2001), in which listeners are enlisted to decide whether output is indistinguishable from a target corpus.[2]

Supposing a system passes the indistinguishability test, however, it's not clear what this (by itself) reveals, since we don't know just how and at what level the test subjects listen. For instance, (Tillmann and Bigand 1996) found that chopping up works of Bach or Mozart into short segments and reversing the order of the segments has little effect on expressiveness and coherence for non-musicians. And given that many generative systems rely on corpus recombination, it's possible that a positive indistinguishability test could be interpreted as a confirmation (or elaboration) of these findings – but presented as a generative model of music.

And it's not just recombinative systems that are subject to this danger. For instance, it's easy to use genetic methods to relatively quickly "evolve" a sentence of Shakespeare – given that the sentence is available as a model! (Dawkins 1996)

This is not to say that corpus-bound, evolutionary or recombinative musical systems are no achievement – they may in fact hold extraordinary lessons for us. But those lessons may be subtle. If we want to know just what has been modeled, a careful look at the generative *process* is necessary. As in Borges' "Pierre Menard", the attempt is made to recreate existing work in a new way, at a new historical moment. And as with the few paragraphs of "the Quixote" that Menard managed to recreate, the interest in musical re-creation lies precisely in the *means* by which it may occur.

And it remains that many of us engaged in exploratory artistic and mathematical work don't have goals that can be verified through indistinguishability. It doesn't serve us to have an empirical standard in "testing" output.

Music, generally speaking, is "evaluated" by critics and audiences, within cultural, social, and political environments, and often with results betrayed by history.[3] These

[1] Measuring *effects* of music with respect to human test subjects is one workaround which may be satisfactory for some purposes, but nonetheless is not equivalent to a direct evaluation of music.

[2] It's also been noted that despite a superficial similarity, this is not a Turing test, e.g. because it doesn't deal with semantics, general knowledge of the world, etc. (Ariza 2009)

[3] (Slonimsky 2000) is an entertaining catalogue of critical assaults on great composers including Beethoven, Chopin, Stravinsky, Tchaikovsky, and Wagner. One might ask "where does music function if not within cultural, social, and political environments" (as one reviewer did). Aside from the fact that cultural and social environments change and therefore evaulation within one or

evaluations may be helpful to artistic development, but they are probably more often irrelevant and occasionally disasterous.[4]

In some sense, what our computational systems generate is *our* music, and as artists we're entitled and obliged to judge it by ourselves. But in another sense, we find the music is not ours, it's the result of an experiment. Perhaps we find it musically lacking in some way, but it represents an intermediate result. Perhaps we find it musically interesting and engaging, but still embryonic.

Papers in which authors conclude that they liked the music generated by their system drive scientists crazy. But I think we have to come to terms with the fact that, in a fundamental way, music meta-creation is (mostly) not a scientific field. A conclusion of "I like it" is vacuous in itself, but it doesn't follow that a paper concluding thus is vacuous. This area of research may be invested with irreducible subjectivity, but it's still a *technical* field, in which mathematical, algorithmic, conceptual, and formal design advances and innovations may take place. And therefore it is this level, our *process* level, at which our discourse is valuable.

To ensure solid contributions, we will have to be careful about describing and evaluating processes for musical meta-creation. The question still seems to remain: how do we know that a given process is "useful" or "good" if we can't formally *evaluate* its output? The answer is this: if a process is described completely, so that it's possible to understand what it does and how it does it, which theories and data are implied and which excluded, which parameters do what, and all assumptions made explicit, then we obtain a *rational* description of a process by which a describable *space* of results may be generated. Whether the results are musically good or lousy, indistinguishable or novel, we've (hopefully) made a technical contribution to the rational and algorithmic exploration of musical space.

Autonomy and Generality

Our task now is to develop an intellectual culture in which description of process is primary, and in which we are able to talk with one another in technical (musical and mathematical) terms without getting bogged down with implementation details. We need to develop a set of categories and questions to ask of any generative system in order to help us characterize and assess it.

As a first broad category, I propose *autonomy*. There are two related ways to talk about autonomy in music generative systems. The first is the sense in which a computational system itself may or may not *behave* autonomously. And second is the autonomy or independence of a system or theory from different *influences*.

Computer programs are not currently capable of truly autonomous behaviour.[5] All decisions are made deterministically or randomly (within constraints).[6] So let us ask of a system, "what are all of the *ways* in which it is *not* autonomous?" To what extent does it depend on explicit templates or rules? To what extent does it rely on recombination or statistical analysis of a database? What assumptions are made about music? What constraints are placed on the output? How deterministic is the program and in what ways? Does it in some sense follow a composing script? At what points does it make decisions and how are these decisions made? Does it develop mathematical structures, and what are the properties and constraints on these structures? What's the role of randomness? In order to understand a system, we need to understand what is being *modelled* computationally, and what is fixed by other means.

This is another way of asking the question "how does it work?" But by casting the question negatively and in terms of autonomy, it's possible that more exhaustive answers will be encouraged. The goal is not primarily to explain how it did whatever it did that was particularly interesting, or to point out any novel ideas, but to *delimit* the space of possibile output. In spite of our desire for computer programs to be autonomous, intelligent, and surprising, if we're going to develop them, we have to keep one step ahead of them! We have to *prove*, in fact, that they're not autonomous, and try as best we can to *predict* their outcome. On the other hand we know that musical meta-creation often *does* produce fantastic surprises – it may be that synergies of sufficiently open structure and randomness can convene on previously unthought designs. This is part of the charm and excitement of working in this area – but nonetheless, as investigators bound by the need for a rational discourse, we must always grow with our systems by working to understand these synergies.

The other side of autonomy is the independence of a system from the influence of a set of ideas (such as conventional musical categories, access to data, principles of information theory, etc.). Naturally, any system will be based upon some set of ideas, excluding others. This is significant because it allows us to take stock of the generative effectiveness of different sets of ideas, including possible overlap where vastly different methods can produce similar results. The *inability* of a given set of ideas to articulate some aspect of musical space will predict that a more general system would require the addition of new concepts. Starting from a *limited* set of

a few such environments is insufficient to really evaluate music (as Slonimsky shows), this author does not believe that music begins and ends with human sociality. A few brief points: 1. music can function for an artist (composer) and for individual listeners even when basically "non-functional" on a larger cultural scale. 2. musical structure may have mathematical or other design properties, some of which may be in some sense universal or near-universal, as is the visual perception of straight lines. 3. as cathedrals have carvings visible only to angels, so might symphonies have structures available only to devine, demonic, or artificial intelligences.

[4] One tragic example: the mental breakdown of talented young symphonist Hans Rott seems to have been related to Brahms' condemnations of his work. Rott's acute crisis involved a the delusion that Brahms was planning to dynamite the train he was travelling on.

[5] (Brown and Martin 2012) discuss *measuring* autonomous behavior in computer programs; in contrast, we admit of no autonomous behavior in computer programs.

[6] At a basic computational level, the only options at each point in time and are to run a deterministic function (given inputs and program state), or invoke the pseudo-random number generator.

ideas, on the other hand, affords an easier and more thorough analysis of those ideas. And we may find that some musical categories, statistical phenomena, or subjective effects may be produced indirectly by simplified and generalized processes, instead of by means of more direct modelling.

As a second broad category under which to discuss a music generating system, I propose *generality*. We can ask: are we dealing with a system that *specializes* in one kind of music? Can it make wildly different pieces? What controls, encourages, or constrains *difference*? In what dimensions can output differ, and what things are always the same? Can (all) parameters be accessed for experimentation? Can underlying assumptions and theories be (easily) changed? What does the system treat as basic or atomic musical structures? Can basic materials be combined flexibly to describe different kinds of larger-scale structure?

Generality, like autonomy, prompts a discussion about the scope and boundaries of the system. It also begins to address the *interactivity* and *extensibility* of the system. A more general system with a broader range of output might be desirable because it would allow us to map a larger, more diverse, yet *connected* subspace of music. If all assumptions and parameters are randomly modulated, this may create the effect of (more) "autonomy," but a controlled investigation is preferable. This means that not only should constraints and decision points be exposed in an *explanation* of the system, but ideally they should be exposed in the system itself, either in a front-end user interface, or, perhaps even better, in the modular design of the code base. Here we can begin to think of our programs not just as parametrized processes, but as extensible, long-term, experimental architectures.

Computational Creativity: Reflexivity and Open AI

Generating musical art, developing a systematic understanding of the structure and design of music, modeling corpora, and testing the limits of human music perception are all reasons to engage in musical meta-creation. Another purpose is to explore the phenomenon of *creativity*.

Human musical creativity is an extremely diverse and hard to theorize phenomenon. Attempts have been made to describe the output characteristics of a musically "creative" (i.e. *seemingly* creative) computer program. (Ritchie 2007) identifies *value*, *novelty* and *typicality* among key concepts; (Jordanous 2012) extracts keywords from written discussions of creativity, and comes up with terms like *communication*, *originality*, *competence*, *development*, and *spontaneity*.

These refinements are helpful if we want to do empirical audience testing with a more specific set of questions, but they do not come toward *rationalizing* the concept of creativity (i.e. making it computable), which in any case might not be possible. We are left with a set of terms that are still subjective, and still *output* based.

I propose that through discussions of generative *process*, through the kinds of questions outlined above, we can start a discussion about what kinds of theories and processes might be, in some sense, creative. What kinds of processes could model creative ways of thinking about musical materials?

Discussing properties of creative process, Harold Cohen finds a capacity for *analysis* to be essential. Cohen, wanting to avoid that problematic term, "creativity", writes about "behavior X," one requirement of which is *reflexivity*, the ability to reflect upon and re-interpret one's own work:

> The individual has to find something in his work that he never consciously put there. Whether these emergent properties result from the juxtaposition of disparate conceptual elements, or whether they result from the technological complexity of the mode of production, the individual must find properties in the work that were neither consciously sought for nor intended, but which may nevertheless lead subsequently to new conceptual structures, new ways of using the technology, or even suggest a modification of existing, consciously-held, goals. (Cohen 1999)

This suggests that a reflexive generative system might be able to analyse its own work-in-progess from premises other than those that generated the structures in question, find emergent structure or pattern not placed there explicitly in previous iterations, and react to its discoveries in a new generation phase by extending or otherwise commenting on them. This is the driving idea behind Eliot Handelman's "Jack and Jill" conception, a theorised (and partially implemented) bicameral composition system in which Jack composes and Jill analyses (Handelman and Sigler 2012).[7]

Non-reflexive music generating systems may do some form of music analysis as well, and non-generating analytic systems may have implications for possible meta-creation systems. Computational music analysis, then, is a subset of the work that musical meta-creatives do, asking: what sort of analytic system might address the kind of musical *listening* we might attribute to a Bach or a Mozart, or to a new composer? And since analysis can be thought of as a dual to generation, every generating system in some sense must address this question.

The creative potentialities of a reflexive system are hard to predict: for example, it's been theorized that reflexivity (feedback looping) is the source of consciousness itself (Hofstadter 2007). Speculation about singularities is beyond the scope of this paper – and anyway, it may be that a totally autonomous, general musical AI, is not the musical future most worth pursuing.

Computation, after all, is not (just) a tool, it's an environment in which we live and in which we *transform ourselves* (McLuhan 1964). Through this transformation, our goals may also change. Tools are useful when goals are explicit; computation, on the other hand, is an open medium for exploration and discovery. And that's why computation is a suitable medium for exploring processes of creativity.

We can, as McLuhan suggests, *extend* ourselves through the medium of computation, and of computational modeling of musical intelligence. The object of this game is not

[7] In fact an extensive "Jill" and a preliminary "Jack" have been realised at the time of writing, but their interaction is not yet automatic.

to create a mysterious black-box system that astonishes and entertains us. It's to grow our understanding of musical possibility and thereby extend our own musical potentialities.

The culture of deception encouraged, for example, by the Loebner prize, doesn't serve this exploration. But in the same paper where the Turing test was first described, we come upon another idea that may serve us better: an interrogation of an AI poet in order to understand its decision-making process.

> Interrogator: In the first line of your sonnet which reads "Shall I compare thee to a summer's day," would not "a spring day" do as well or better?
>
> Witness: It wouldn't scan.
>
> Interrogator: How about "a winter's day," That would scan all right.
>
> Witness: Yes, but nobody wants to be compared to a winter's day.
>
> Interrogator: Would you say Mr. Pickwick reminded you of Christmas?
>
> Witness: In a way.
>
> Interrogator: Yet Christmas is a winter's day, and I do not think Mr. Pickwick would mind the comparison.
>
> Witness: I don't think you're serious. By a winter's day one means a typical winter's day, rather than a special one like Christmas. (Turing 1950)

This example is whimsical, and probably beyond the current possiblities because of semantic and real-world knowledge and associations. But something very close to this is already attainable in the musical domain. Along with a musical piece, it might be possible to somehow display, query, or manipulate the structures and decisions that generated it.

If we want to pursue artificial intelligence as a medium for discovery and extension, our programs must be just the opposite of autonomous black-boxes, they must be radically *open* to interaction. *"Open AI"* is a way of thinking about computational modeling in which understanding of process and rational mapping of results are the primary discourse. It's a way of thinking about system-building in which flexibility, extensibility, interaction, and communication are key concepts, in order to facilitate systematic experimentation.[8]

Open AI is a commitment to keeping one step ahead of our "AI" systems – no matter how smart they become – so that we bootstrap ourselves into higher understanding of the phenomena we study. It's a way of growing human musical creativity through the medium of computational investigation. We could be at the beginning of a new kind of musical practice, where describing musical thought and musical design becomes an open art project that we all work at together.

[8] "Interaction" here is meant very differently than the way it is often used in discussions of musical meta-creation. Runtime human intervention (e.g. assisted composition, improvisation) probably would *not* facilitate systematic experimentation with computational models of musical thinking unless the human side was *very* well controlled – and in this case, it might be easier to automate. What's meant is that *formal testing* is facilitated by easy interaction with program settings.

Acknowledgements

The author is grateful for ongoing collaboration and discussion with Eliot Handelman, in which all of these ideas were developed, and in which Eliot coined the term Open AI. The author is also grateful for the support of Doina Precup, Prakash Panangaden, and the School of Computer Science at McGill University, as well as the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

Ariza, C. 2009. The interrogator as critic: The Turing test and the evaluation of generative music systems. *Computer Music Journal* 33(2):48–70.

Brown, O., and Martin, A. 2012. Autonomy in music-generating systems. In *Proceedings of the International Workshop on Musical Metacreation at AIIDE*, 8–13.

Cohen, H. 1999. Colouring without seeing: a problem in machine creativity. *Department of Visual Arts, University of California at San Diego*.

Dawkins, R. 1996. *The Blind Watchmaker: Why the Evidence of Evolution Reveals a Universe without Design*. W.W. Norton & Company, New York.

Handelman, E., and Sigler, A. 2012. Simulating the listening of composing. In *Proceedings of the Music, Mind and Invention Symposium*.

Hofstadter, D. R. 2007. *I am a Strange Loop*. Basic Books, New York.

Jordanous, A. 2012. A standardised procedure for evaluating creative systems: Computational creativity evaluation based on what it is to be creative. *Cognitive Computation* 4:246–279.

McLuhan, M. 1964. *Understanding Media: The Extensions of Man*. McGraw-Hill, New York.

Pearce, M., and Wiggins, G. 2001. Towards a framework for the evaluation of machine compositions. In *In Proceedings of the AISB01 Symposium on AI and Creativity in Arts and Science. AISB*, 22–32.

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

Slonimsky, N. 2000. *Lexicon of Musical Invective: Critical Assaults on Composers Since Beethoven's Time*. W.W. Norton & Company, New York.

Tillmann, B., and Bigand, E. 1996. Does formal musical structure affect perception of musical expressiveness? *Psychology of Music* 24:3–17.

Turing, A. M. 1950. Computing machinery and intelligence. *Mind* 59:433–460.