The Music of Machines

Investigating Culture and Technology in Musical Creativity

Florian Grote

Native Instruments GmbH
D-10997 Berlin
florian.grote@native-instruments.de

1 Abstract

Machines are an integral part of musical creativity. This is not only true for the obvious cases of electronic music, but also for classical music recorded digitally, or rock music utilizing software for sound manipulation. In all these areas, artists pose queries to machines, in the hope of finding sounds or parts of arrangements that fit to what they are trying to achieve. Current machines provide static answers to such queries, but recent developments highlight the potential for machines that are aware of the situational context of the query, and provide results that fit better to what the artist searched for. This paper investigates the potentials and hurdles for situation-aware machines, and speculates on the aesthetic effects such technologies might have in the field of music.

2 Introduction

Musicians have always played in groups, whether singing and beating drums at a campfire or in full-scale professional orchestras. The phenomenon of group creativity is in no way limited to the field of music, but rather can be found in all fields where humans are creative [Sawr07; Sawr12; Gabr13; Kyrk13]. In general, groups of collaborating humans seem to be more effective at being creative and coming up with better creative solutions than the proverbial "lone genius". This notion raises questions when collaborators come into play that are not human. In 2004, the musician Squarepusher published an article titled "Collaborating with Machines" [Jenk04]. He describes how working with "machines" has changed how he works, how he sees himself as an artist, and the aesthetics of his results, i.e. the music he writes. The "machines" he refers to are all computation devices running a variety of software. The hardware design of the

devices has in some cases been optimized for use in music applications, while in other cases he is referring to generic consumer-grade computers that run dedicated music software.
Squarepusher is not an exception in the cultural realm of music. When using machines in their music, artists effectively collaborate with the objectified knowledge of the organization that manufactured the machine. Being embedded in a technological device, this information cannot be compared with the knowledge of human collaboration between e.g. artist groups and experts from the same organization. The main difference is that in human collaboration, responses to any questions that arise will be situation-specific, as the participants will apply knowledge which they deem fitting in this particular case. For example, a saxophone player in a jazz ensemble will adjust his melodic and timbral expressions to the situation in which he plays, and this will lead to different results in the intimacy of a small club versus a big stage in front of thousands of listeners.

3 Situation-aware Music Machines

Music machines, as Squarepusher described them, so far do not offer responses specific to the situations in which thy are used. Of course, this is largely due to a lack of situation-awareness on the part of the machines, which currently do not have sophisticated means to detect their surroundings. They have no concept of the differences between the intimate jazz club and the big stage. From this, it could be concluded that the goal of technological development should be reaching a similar level of situation-awareness as we can see in the saxophone player. However, research conducted in the GiantSteps[1] project suggests that human collaborators are not necessarily keen on their music machines reacting in ways that are similar to how other humans might react in collaborations [Andr15].

"Yeah, yeah, well I like to be completely in charge myself, but I like to... I don't like other humans sitting the chair, but I would like the machine to sit in the chair, as long as I get to decide when it gets out." [Interviewee RBMA014]

[1] GiantSteps is a research project which has received funding from the European Community's Seventh Framework Programme, and in which the author is active as a consortium researcher.

"For my taste, it needs to be more gentle and really, like, if you suggest something, you got to know what you are doing there. It's just, if you, for me, if this engine is constantly putting stuff I don't want, I will turn it off and I will not turn it on again. It's just, you lose this kind of trust connection, and then it's over." [Interviewee NIBLN004]

"I mean, when the track is finished, and I load it up and people like it, then I'm like, okay, who had this idea? Was it me having this idea? Was I doing the track, or has it been a lucky day and the algorithm did well... (laughing) You know, there's this little bit of doubt in the end, that can hurt your artistic self-confidence." [Interviewee NIBLN006]

The fact that algorithms devised to react with situation-awareness have so far mainly been designed to mimic human behaviour as closely as possible may be their main shortcoming. This is true for most recommender and expert systems, which see the human recommender and expert as the ultimate goal to match or even surpass in modelling. The field of artificial intelligence is a similar case: in many of its projects, the "intelligence" meant to be recreated using computing machines is the human intelligence. The mathematician and philosopher Gotthard Günther has already expressed his doubts about these goals in the middle of the 20th century. In his book "Das Bewußtsein der Maschinen" [Günt63], he argues that for machines to appear intelligent to their human operators, mimicking human traits or processes in human intelligence is not necessary, and it may even prove hindering or, in any case, not effective. Instead, his argument goes, machines will appear intelligent when they can react with applications of knowledge that are obviously specific to the current situation, and that fit the purpose sought by human operators or collaborators.

Thus, intelligence cannot be said to just "be", in a sense of true existence. Rather, it is ascribed to a thing or a being in the situation, by the respective observer describing what the situation is. Similarly, there is no predefined group of participants, but participation is rather established through discourse each time it is expected to happen. This view is informed by constructivism in sociology, especially the sociology of the arts and of inclusion by Niklas Luhmann [Luhm00; Luhm12]. According to this view, a group of musicians like a band is held together not by a generally known fact (they belong together!), but through shared expectations among them as well as among members of the public who follow the band. Each time the band is mentioned

without indicating a change in composition, the expectation of continuity is fulfilled. However, this is nothing a band member can fully rely on, as being kicked out of a band always remains a possibility. This unreliability also has a big advantage: who and what is participating in a creative collaboration is in principle subject to change with each new instance of a situation in which it can happen. Human members of a band may remain rather static (albeit unreliably so), but machines can be exchanged more frequently. This is not to say they do not have an identity which can contribute to the overall identity of a music group or its human members. There are numerous examples where machines have played a major role in the public perception of artists. But their identity is of a different kind, which has to do with their thing-ness [Heid71] and their technological qualities, rather than the originality of the music they produce themselves – if they themselves are capable of producing music at all. Their utility is defined by this identity matching that of the artist or band, and the availability of knowledge it brings into the situation where it is supposed to be used. This makes it easy to add them to a situation, to leave them out, or indeed to remove them, which would not be as easy with a human collaborator.

So what can we expect from machines that react with situation-awareness? What would change in the settings of creative collaboration involving humans and machines? Research conducted in the GiantSteps project suggests that humans expect the behaviour of machines to change mainly in situations where the human operator is searching either for something specific or, more general, for inspiration in the objectified knowledge of the machine. The take-away message here is that musicians do not usually expect a machine to come up with its own recommendations, but rather see the usefulness of better results for queries they gave to the machine, i.e. they would like a more proactive machine primarily in cases where they have already asked it for help anyway. This is mainly the case when they are searching for sounds to be used in their music, or when having to organize an archive of older music projects. Notably, this does not go in the direction of handing off the compositional work to a machine [as is illustrated in Cope00 or Coll06; Coll14], but rather does it express the hope that machines can handle menial tasks on their behalf.

"So it would be really useful to for example have some kind of sorting system for drums, for example, where I could for example choose: 'bass drum', and here it is: 'bass' and 'bright', and I would like it to have maybe bass drum 'round' and 'dry', and you can choose both, the more I choose, of course, the less results I will have, and I would have one left at the end or

maybe five. So it is filtering it down, that is really helpful, if it works well of course." [Interviewee RBMA002]

The designers of such machines are therefore well-advised to focus on aspects that help users in situations where they have already actively requested help. The main difference will then be how this help is provided. With the traditional file-based or database systems available in current machines, a given query will always lead to the same results. The usefulness of these results directly depends on how well their organization scheme matches the current needs of the user. Situation-aware machines can, at least in theory, adapt to the context in which the query was generated, and present results that may be more useful right away. The user may have to go through and dismiss fewer results until making a positive choice, i.e. effectively reducing the number of negative choices needed to get to a positive choice. This could, for example, speed up the search for a sound that matches the mood of a music piece.

3.1 Four Dimensions of Situation-Awareness

Situation-awareness in music machines can be based on four dimensions: cultural context or genre, relative origin of the query, personal preferences of the operator posing the query, and success of previous responses generated by the machine. Each of these dimensions can be analysed separately, yet they clearly interact in many significant ways.

Let us start with cultural context or genre. Knowing the genre into which results of a query are supposed to fit could allow a machine to limit its results to exactly those it associates with this genre. While theoretically elegant, this approach is met with several obstacles in real-world applications. First of all, detecting a genre is a non-trivial, highly observer-dependent task. In music culture, there exists no single entity which could determine into which genre a given cultural artefact falls. The problem becomes obvious for music genres when looking at services which try to become such entities for recorded music, such as Discogs.com or AllMusic.com. While Discogs.com relies on the input from its users, AllMusic.com employs a service provider, which categorizes music recordings into genres and their sub-classes. While both Discogs.com and AllMusic.com try to enforce a strict ontology of genre relations, it remains hard to define borders between them. Therefore, it seems unlikely that any piece of music can be easily categorized as belonging to a very specific genre without also looking at the cultural context in which it was created and is being received. The classification then has to check whether a detected genre actually fits the cultural context in which the classification was triggered, ideally

also taking into account the self-descriptions of the artists and any classifications they might have used for themselves.

The next dimension to be mastered for situation-aware machines is the structure of relations in which a query originates. The machine should have a concept of the role its operator plays in relation to other musicians or the audience, both in an overarching cultural sense and in the actual situation at hand. It could then offer results which directly support this role, or which may help a user to break out of it. Again, this goal is very hard to achieve in reality. How should a music machine become aware of other musicians, let alone the audience? One approach could be to connect all musicians and perhaps the audience via a technical protocol, as it has already been done with MIDI to some extent. This standard from the early days of personal computing does, of course, not contain a concept of roles played by musicians and machines which it connects. The challenge becomes even harder if the audience and its sentiment are to be taken into account. To achieve a notion of the overall situation in terms of social and cultural relations, real-time sensors such as cameras or laser scanners might be required, alongside access to the participants' online social network relations, to gain a consistent picture of who initiates a query and together with who and in which technological context the results are to be used.

The matter of personal preferences in human operators becomes even more important when the machine starts to appear more independent. When creating the order of results to present to the user, the machine would ideally have a concept of how new this presentation will be for the user, and how comfortable or uncomfortable she will be with the results. This would theoretically be possible, since the machine has many different sessions with the same user, but in practice a meaningful assessment of what to show and what to hide based on personal preferences of the user will depend on a correct assessment of the relational social and technological structure in which the query originated. The difficulties associated with this approach have been lined out above. Nevertheless, it may be easier to lay some groundwork in this dimension, as the machine will be able to collect data on previous usage and use this accumulated knowledge to derive the importance of certain specific results for the user. For example, favourite sounds which have been used over and over again should probably not be hidden from the results because they get a lower score in other dimensions, unless the user is trying to break away from old habits, which will again be hard to detect – just as it is hard to detect for a human collaborator.

This leads to the fourth dimension of situation-awareness in machines, the notion of success. As has been stated above, current databases used in music machines commonly present static results in response to queries by their operators. If a hypothetical situation-aware music machine could gain an understanding in all three dimensions mentioned above, then it would, under the same circumstances, also always render the same results, only more tailored toward the situation at hand. Nevertheless, the machine would be entirely determined by external states. This would change if the machine had some concept of success, i.e. would be able to know whether one of the results it presented was chosen, and what happened with the piece in which the result was used.[2] Was it a success in relation to similar pieces? How was it received in the public, and which genre attributes were associated with it? Did it change the identity and role of the artist who made the choice, both in their self- and in their public image? Obviously, the data needed to derive such knowledge is hard to come by for a pre-programmed machine, but it exists and may be integrated at some point in the future. After all, the knowledge of how successful a music piece is has been made available for algorithmic exploration by all major music labels, and efforts have already been undertaken to automatically generate pieces that are successfully received by human audiences [Stein12]. For music machines, we see more interesting applications for algorithmic representations of individual success in terms of the machine knowing whether the human operator was happy with the choice she made.

3.2 Outlook: Aesthetics and Culture

Advancements in any of the fields mentioned above will change how humans and computers interact. Depending on how the changes are implemented, humans may no longer be able to see themselves as operators, but rather accept the machines as collaborators, as Squarepusher has already stated. Notably, with the possible steps lined out above, this much greater degree of machine integration could be achieved without any specific approaches of human-level artificial intelligence, and particularly without the aim of creating self-composing machines.

Which implications might this have on the aesthetic results of the creative work? Without any situation-aware music machines available, it is hard to predict what might happen with the creative output of musicians. We can make two assumptions, though.

First, the way musicians view cultural genres and their sub-genres will change considerably, which may lead either to more cultural cross-pollination, or indeed to the opposite, a well-

[2] A similar line of research in robotics aims to create "socially intelligent robots" by incorporating feedback from their human operators into future actions [Brzl05; Knox13].

informed, intentional form of cultural segregation. This is because culture is very much based on comparisons how knowledge is applied in specific ways in different areas. Dirk Baecker [Baec12] has developed a theoretical framework around this function of culture, where, in essence, knowledge is about knowing how something is done, and culture is about knowing how something is done in one's own cultural area versus how it is done elsewhere. Such cultural areas may be synonymous with music genres, for example. Whereas so far in history, artists had to attain this knowledge about how something is done in one cultural area versus others by either experiencing it or actively learning about it, this information will increasingly be called up in machines. It may influence how a machine structures its results to a query without the user having any idea why the results are structured in this particular way. Even though this type of knowledge can already be gained with the help of technological means, such as libraries, encyclopaedias, and online services, the situation-aware music machines described in this article will be much more upfront, as they will be able to take the perspective of their operator, and in the best case provide them with an overview of how some things are done differently in other cultural areas it has information on. This may help the operating artist broaden their perspective, and it will also help them discover cultural areas that fit very well to what they themselves are aiming at, potentially leading to more and broader collaborations and cultural cross-pollination in the form of imports of knowledge on the specific ways in which things are done elsewhere, where the musician operating the machine would not have looked without it.

The other aspect has to do with the artists' identities. Music machines that are increasingly aware of the situation in which they are used will seem more and more intelligent. I am carefully choosing the word "seem" here, because as stated above, this is not necessarily connected with the field of artificial intelligence. However, a machine that seemingly knows what its user needs in a specific situation will be perceived as acting more independently than present-day music machines. This could lead to artists questioning their own role as authors of their works versus the role the machine plays in creating them, and potentially to total rejection of situation-aware machines, if users are not willing to accept giving up parts of their independence as authors. On the other hand, it could also lead to an increase in music collaborations and result in a new era of sampling each other's work. Indeed, if the very strong notion of the individual human author owning the intellectual property of a work of art is weakened a bit, this may bring much more freedom of choice to both humans and machines.

4 Acknowledgments

The author would like to acknowledge support from the GiantSteps project, which has received funding from the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement no. 610591.

Literature

[Andr15] Andersen, Kristina, Florian Grote. GiantSteps: Semi-Structured Conversations with Musicians. In: CHI'15 Extended Abstracts. Seoul, 2015, forthcoming.

[Baec12] Baecker, Dirk. Wozu Kultur?. Berlin: Kulturverlag Kadmos, 2012.

[Brzl05] Breazeal, Cynthia. Socially Intelligent Robots. interactions 12, Nr. 2 (March 2005): 19-22.

[Coll06] Collins, Nick. Towards Autonomous Agents for Live Computer Music: Realtime Machine Listening and Interactive Music Systems. Cambridge: University of Cambridge, 2006.

[Coll14] Collins, Nick. Virtual Musicians and Machine Learning: Towards Musical Familiars and Touring Als. In: Karen Collins, Bill Kapralos, and Holly Tessler (Eds.). The Oxford Handbook of Interactive Audio. Oxford: Oxford University Press, 2014: 350-363.

[Gabr13] Gabora, Liane, Steve DiPaola. How Did Humans Become So Creative? A Computational Approach. arXiv:1308.5032, August 22, 2013. http://arxiv.org/abs/1308.5032.

[Günt63] Günther, Gotthard. Das Bewußtsein der Maschinen: Eine Metaphysik der Kybernetik. Krefeld, Baden-Baden: Agis-Verlag, 1963.

[Heid71] Heidegger, Martin. The Thing. In: Poetry, Language, Thought, 163–80. New York: HarperCollins, 1971.

[Jenk04] Jenkinson, Tom (Squarepusher). Collaborating with Machines. Flux Magazine, 2004. http://neugierig.org/pics/livejournal/2004/03/19-jenkinson.pdf.

[Knox13] Knox, W. Bradley, Peter Stone, and Cynthia Breazeal. Training a Robot via Human Feedback: A Case Study. In: Proceedings of the 5th International Conference on Social Robotics - Volume 8239, 460–70. ICSR 2013. New York: Springer-Verlag New York, 2013.

[Kyrk13] Kyriakou, Harris, Jeffrey V. Nickerson. Idea Inheritance, Originality, and Collective Innovation. arXiv:1311.0527, November 3, 2013. http://arxiv.org/abs/1311.0527.

[Luhm00] Luhmann, Niklas. Art as a Social System. Stanford: Stanford University Press, 2000.

[Luhm12] Luhmann, Niklas. Theory of Society. Stanford: Stanford University Press, 2012.

[Sawr07] Sawyer, Keith. Group Genius: The Creative Power of Collaboration. New York: Basic Books, 2007.

[Sawr12] Sawyer, Keith. Explaining Creativity: The Science of Human Innovation. Oxford: Oxford University Press, 2012.

[Stein00] John-Steiner, Vera. Creative Collaboration. New York: Oxford University Press, USA, 2000.