# Machine Learning, Music and Creativity: An Interview with Rebecca Fiebrink

**Simon Holland[1] and Rebecca Fiebrink[2]**

[1] Music Computing Lab, The Open University, Milton Keynes, England.
[2] Department of Computing, Goldsmiths, University of London, England.

**Abstract** The following interview was conducted in London and Milton Keynes by Skype. Rebecca Fiebrink is a Senior Lecturer at Goldsmiths, University of London, where she designs new ways for humans to interact with computers in creative practice. As a computer scientist and musician, much of her work focuses on applications of machine learning to music, addressing research questions such as: 'How can machine learning algorithms help people to create new musical instruments and interactions?' and 'How does machine learning change the type of musical systems that can be created; the creative relationships between people and technology; and the set of people who can create new technologies?' Much of Fiebrink's work is also driven by a belief in the importance of inclusion, participation, and accessibility. She frequently uses participatory design processes, and she is currently involved in creating new accessible technologies with people with disabilities, designing inclusive machine learning curricula and tools, and applying participatory design methodologies in the digital humanities. Fiebrink is the developer of the Wekinator: open-source software for real-time interactive machine learning, whose current version has been downloaded over 10,000 times. She is the creator of a MOOC titled "Machine Learning for Artists and Musicians." She was previously an Assistant Professor at Princeton University, where she co-directed the Princeton Laptop Orchestra. She has worked with companies including Microsoft Research, Sun Microsystems Research Labs, Im-agine Research, and Smule. She has performed with a variety of musical ensembles playing flute, keyboard, and laptop. She holds a PhD in Computer Science from Princeton University.

**Holland:** How did you first become involved in HCI research?

**Fiebrink:** I came into my PhD, back in 2006, with an interest in music information retrieval. At that point in time, even though machine learning had been used in music performance systems much earlier by people like David Wessel, there wasn't a lot of focus in looking at how machine learning techniques could be used in performance, or by creative practitioners, and so I saw an opportunity. I was surrounded by people who were experimental musicians and composers, and I got really interested in this question of what might happen if we took some of these techniques that I saw gaining traction, and showed some interesting proofs of concept to the ISMIR[1] community - what would happen if we gave these kinds of tools to creative practitioners. And so, when I started that work I didn't necessarily approach it from a very formal HCI standpoint, but I was very interested in making tools that were usable by other people who weren't me. Given the goal of understanding what they wanted to do with these tools, and not just having everything be driven by what my own ideas were - I very naturally found that HCI gives a set of methodologies and perspectives and modes of evaluation that support those kinds of values.

**Holland:** Do you make music yourself?

**Fiebrink:** Not as much as I used to! Since coming to Goldsmith's University, I haven't had an ensemble that I'm active with, and the kind of electronic music that I make is very much social. I'm not a solo performer, but even when I was working with the Princeton Laptop Orchestra and associated band - when I was at Princeton before moving to Goldsmiths - I often looked at the creative work that I was doing as having dual functionality. From one perspective, it was simply fun - engaging in creative expressive activities that brought me satisfaction just for their own sake. But at the same time, as a researcher, I was able to justify taking the time to do it from the perspective of 'dog fooding' – in the sense that if I'm going to make a tool that other people are going to use, it's always a good idea for me to make sure that it's at least good enough to support my own

[1] International Society for Music Information Retrieval Conference

practice - and so I approached it from that perspective. Obviously I'm going to listen to other people and work with participatory design processes, but at the same time complement that with my own hands-on work - trying to get to the heart of how do I make this more efficient, are there possibilities here that I'm only really going to find by getting my hands dirty.

**Holland:** Many researchers come to work of this kind from a perspective of being a performer or a composer but it's interesting that you mentioned that your work started in part in the context of music information retrieval - can tell us a bit more about that?

**Fiebrink:** My background is multifaceted. I have an undergraduate degree in computer science and I also have an undergraduate degree in flute. I did a master's degree in music technology at which time I became really interested in music information retrieval but was also doing some side projects that were more related to NIME[2] so my interests were always quite broad. On this basis, I approached the work not with the main goal of making things for me to use as a performer, however I can speak the same language as performers, because I have a lifetime of experience being a performer.

**Holland:** What influences affect the way you develop your work?

**Fiebrink:** I'm drawing on a lot of different perspectives in my work. I am a computer scientist and a programmer. I also take ideas from what's happening in the machine learning community and what's happening in the music information retrieval community. Certainly, being able to prototype new technology myself is really crucial to the way that I look at the space of possibilities. It allows me to engage in really hands-on participatory experimental processes, when I'm making stuff and people are trying it out - as opposed to being limited to approaches that are more removed – for example simply trying to understand people's existing practices. Also many of the research questions that I'm most interested in are not just technical research questions; there are wider questions about things like: What is machine learning good for? How do we make better tools for creative practitioners? What should creative

[2] New Interfaces for Musical Expression

practitioners learn or know about particular technical practices in order to use tools effectively? How do we educate people or build interfaces that might feel well matched to what people already know? I have a driving interest in advancing, changing and developing the types of creative practice that people can do with technology - I would say that's my main motivation as opposed to 'how I can make a better algorithm?'.

**Holland:** When you're working on research that involves music does it have implications any wider than music?

**Fiebrink:** Definitely, yes. When I did my PhD work I was focusing quite explicitly and narrowly on music and on building tools for electronic musicians and composers - but there are immediate applications to other domains. It's not a big leap from thinking about building a gesturally controlled musical instrument to building a gesturally controlled game or building interactive art installations. A suitable domain might be audio-visual, or might not even have a musical component. My work can be applied in any situation in which people are interacting in a space with sensors, where information about what they are doing gets translated into some aspect of what's happening. I pretty soon I started working with folks who wouldn't necessarily describe themselves as musicians but they were nevertheless people doing creative stuff with sensors in closely related fields. And then there's a piece of my work where I ask questions about what it means to support composers or musicians in their practice. There's a piece of that which applies far more widely and gets fundamentally at what does it mean to support people involved in a creative practice - that's not domain-specific.

I draw a lot on work by Ben Shneiderman, and more recent work from Celine Latulipe, who are looking at things that are common across creative practice. These are not the types of things that we might think about when we're designing user interfaces for word processing or web browsing. Things like making it really easy for people to prototype an idea so that they can get a hands-on feeling for whether their idea is any good - making it easier for people to explore lots of ideas in a given space, rather than forcing them to commit to one initial idea.

Certainly one of the big themes in my work is about making interfaces that allow people to communicate embodied information or embodied practices and ideas to the computer. The body is highly important, and keyboards and mice aren't great for that, Neither is programming ideal for communicating ideas about human movement. Those issues manifest quite clearly in music, but they're shared across a lot of other fields.

**Holland:** For people who don't know, tell us a little about your PhD.

**Fiebrink:** My starting point was asking wahat might be needed to enable musicians and composers to use machine learning in their work without requiring them to get a PhD in computer science first. In order to explore that, I did a lot of building of software prototypes and did that in conjunction with several composers using participatory processes. The software that came out of that is called Wekinator. Wekinator allows people to do machine learning in real time. That's something that I've continued to develop and release and it's been downloaded about 10,000 times now. It's used in a lot of teaching around the world. A more research-oriented outcome is that in order to make machine learning useful and usable to people doing things like making new instruments, it tums out that the conventional approach and conventional assumptions around machine learning aren't necessarily appropriate. For example, this idea that you have a ground truth data set that you want to model as accurately as possible goes out of the window - because what people often care about is solving some bigger creative problem, or building something that functions within a particular context - and the training data set may often be quite malleable. Or you may start with one data set, and you build a model that models that data set quite well, but it doesn't really do what you wanted to do as a person. Often you're able to change that data to give better results. So one of the other main outcomes of this work was identifying human-in-the-loop processes that make sense for applying machine learning to creative problems, and of course Wekinator embodies those ideas.

**Holland:** There are certain criteria for whether research is successful and other criteria for whether musicians or creative people think

you're helping them. Is there a tension between these two things, and if so, how do you navigate it?

**Fiebrink:** There's definitely tension between those things. I'm happiest as a person when I'm making things that are useful to people, but I make my department happiest when I'm publishing highly cited research papers! Sometimes the first thing does lead to the second, but not always. Sometimes, for instance, it's hard to communicate the particular challenges and goals of computer music to a broader set of reviewers, for example HCI reviewers. I don't necessarily think that's a bad thing, but it's a fact of life. It can be hard to try to tell the story of why solving a particular problem can be extrapolated into broader findings without having, for example, a comparison of multiple domains, or showing very clearly that this one thing that experimental artists or musicians care about is actually a manifestation of this deeper set of concerns. Sometimes the storytelling around that is easy to do in a compelling way, but sometimes, especially if people haven't thought much about how to build on a new approach for music or the arts - it's hard!

Another obvious issue is that some of the evaluation methods that are expected at venues like CHI (the premier international conference on Human-Computer Interaction) are very different from what you would want to do in practice to understand whether you have built something that's useful for musicians or not. Some of the things that are really meaningful for me in understanding if the thing I built works need to take place over really long timescales. Has something been adopted and propagated over a period of years? Or at a very local level - this tool that I built for this music teacher, are they still using it, or are they developing a curriculum around it? So there are all sorts of factors. You generally can't measure whether one musical interface is better than another using established criteria - you're often building technology to enable something totally new and the criteria may change. Developing new technology often entails developing new evaluation methods as well, so there's all sorts of challenges.

**Holland:** Have you developed any strategies for explaining work that straddles these boundaries to HCI referees, or does it have to be ad-hoc?

**Fiebrink:** A bit of both. In terms of general strategies, one approach is to link it to existing threads of research in the HCI community. So, for example, my work with interactive machine learning and music is not just about music, it's also relevant to a larger space filled with people who couldn't care less about music but might be interested in how we can improve machine learning systems by putting humans in the loop. So I feel I have something useful to say about that and I've written some papers where you can talk about creative use cases of using machine learning in the interface as a complementary perspective in other perhaps more traditional or unambiguous use cases.
There is a similar situation with the discourse around what makes a good creative technology such as the work that Celine Latulipe is doing. There's a thread of that woven through the CHI community where I can currently engage and bring a different set of perspectives and show what machine learning can bring to creativity.
So I think its good that you have to contextualize your work against the types of things that a particular community cares about – but I'm not always successful!

**Holland:** Are there areas where music interaction still has lessons for mainstream HCI?

**Fiebrink:** That's a good question. One of the challenges for musicians and people who research in music is that often we're not particularly good at articulating what makes something a positive experience or an engaging interface, or at any rate articulating that in ways that that naturally suggest linkages to HCI. That doesn't mean that they're not there. One thing about music and the arts is that there's a tradition of practice-based research and there's a tradition of self reflection on one's work. You can find this done well in different music conferences, for instance, where somebody writes a paper as a composer and talks about their rationale for doing things the way that they do. Obviously auto ethnography is not a new method, but a lot of those papers are fascinating to read and may perhaps contribute to the dialog around formalizing and refining methods of this kind and importing them into other fields. That's something that I do feel is appropriate and valuable. In my work, when I interview people who use my software, I'm getting information of that kind, and

trying to trying to understand as deeply as possible why they're doing what they're doing and all the different factors that that come into play when they're composing a piece. I think that some methods and findings of this kind can be encapsulated as case studies, but perhaps there is more to be understood and articulated methodologically.

**Holland:** Are there things that mainstream HCI knows about that music HCI is neglecting?

**Fiebrink:** Not that come to mind immediately! I can't think of any HCI papers that I've read recently where I think oh that would be great to apply to music and nobody's done anything like that before. That's not to say they aren't out there but in general I feel like there is a contingent of people within the music community who are pretty on top of what's happening in the HCI community!

**Holland:** What are you researching at the moment?

**Fiebrink:** Let's see if I can give you a lightning overview. I have one set of projects around continuing to look at ways of making end-user machine learning more usable especially in creative contexts. For instance, looking at how to make feature selection by musicians or artists easier, because that's something that Wekinator doesn't do and other tools don't do, and deep learning doesn't solve - even though people think it does. That's one of the big practical barriers to people using machine learning in creative work.

**Holland:** Lets just dig in a little - how might you do that?

**Fiebrink:** We don't know yet, we're experimenting with a few different approaches. My intuition is that probably the way to do this is a mixed initiative approach where you're combining some user driven information that you can encapsulate into some kind of objective function with some automated search or recommendation of features. But there are other approaches. Just providing good visualizations and providing better interfaces that meet users where they are, in terms of their understanding of data and recognising that they may not know signal processing for instance. Taking advantage likewise

not just of users clicking on a GUI, but users demonstrating invariants for instance.

**Holland:** So how might a user demonstrate invariances?

**Fiebrink:** Instead of a user just showing an example of a gesture, they could intentionally show several different gestures meant to communicate the same intent. So for example, if speed is invariant, then with some kind of point and click interface maybe the user could somehow specify that, but far more flexibly, from a bunch of different example gestures, the system might infer that this feature is not a discriminating factor.

**Holland:** What other projects are you working on right now?

**Fiebrink:** I've got another project working with music therapists and kids with disabilities where we are looking at how to build better on-the-fly instrument-making tools. A central idea is that a music therapist could sit down with a kid with potentially quite severe physical disabilities and quickly make a customized interface. That has been the seed of that project, but it's branched out into a few different variations. For example, if we make highly configurable tools in terms of what gestures people can use to control music, what else needs to go into tools for music teachers to effectively build curricula around them? Can we build tools that allow kids with disabilities - and kids without disabilities who are playing acoustic instruments, for example - to participate collaboratively in music classrooms? That's a project with collaborators in Northamptonshire.

**Holland:** Does this line of work carry a responsibility to ensure it’s sustainable?

**Fiebrink:** Yes. Making things sustainable is always tough without sustainable funding. But part of my approach is always at a minimum to make it open source and free to download, and to provide as much documentation as we can in different formats, and strive to develop a community of users.

We just wrapped up a project called Rapid Mix, which is a Horizon 20203 project. Our part of the project was largely focused on making better tools for machine learning for creative developers. So for example, to serve the needs of hackers[4], makers, developers in gaming and audio, and creative coders, we crafted a programmer-level API for machine learning that you can interact with without needing to be a machine learning expert. This work was tailored for people who want to interact with machine learning where their training set may be a moving target - where conventional evaluation metrics might not be as useful as just building something that you can play with in real time and trying it out to improve it iteratively.

**Holland:** Can you give an example of what that might have made possible that perhaps wouldn't have existed before?

**Fiebrink:** For example we produced an open source set of JavaScript examples making it easy to use machine learning to create flexible interactions using sensors such as those in phones. When we started this project, JavaScript developers making mobile or desktop apps faced steep barriers in doing this kind of thing. They had to deal with large quantities of boilerplate code and make many low-level decisions about what algorithms to use, as well as dealing with libraries that assume that your training data must be stored in a file or a database. What you can do instead now is go to the Rapid Mix website. There's JavaScript and web-based tools that let you see the source code, see the executed code right next to it, edit it live in real time, and there is a suite of demos showing how to carry out the entire learning process, from collecting data, to training a model, to testing it out, to changing the model in real time, and showing how do that with sensors, with audio and with video. As well as the value of these tools for developers, we are also now finding them useful for teaching. The web makes these resources highly accessible, and of course JavaScript is a standard introductory language for many people.

---

3 A major research programme funded by the European Union

4 'hackers' in the original benign sense for this word.

**Holland:** What are the interesting problems in Music and HCI that people should be working on?

**Fiebrink:** My answer to this hasn't changed dramatically from what it was several years ago. Even though research keeps advancing, what tends to happen is that the answers become more nuanced over time rather than the question changing. For me, the big questions are around what are the uses of machine learning in creative practice.

There is a lot of focus both among machine learning researchers and the general public about using machine learning and AI to replace people and duplicate human creative processes and I find this such a limiting viewpoint. From an artistic and humanist perspective, people derive a lot of value from making creative work, so we need to make tools that people actually use, where we are adding value to people's lives, rather than replacing people.

There are also commercial and practical reasons to broaden the question beyond replacing people. Some of the algorithms proposed are effectively big red buttons that you push and the music happens. That's not really that useful if you want to generate music for a particular application. Often you are going to do a better job if you have nuanced ways of taking into account information about the context, the users goals, the characteristics that they would like to see in that finished product. There's such a gap between these autonomous algorithms that just do something that's fairly fixed for a given training set, as compared with something that flexibly integrates into a real-world media content production context. So I think there's so much to explore in looking at how machine learning algorithms can be can support the work that humans are doing - how they can be collaborators, and how they can be embedded into tools. There's all these different roles that algorithms might take rather than replacing a human. You can think about a learning algorithm as something that challenges or pushes back against you. You can think about it as something that's analogous to a paintbrush. There's all sorts of other metaphors that we could draw: collaborator, opponent, mentor, student, instrument, and many others.

**Holland:** How can machine learning take into account context better in supporting creative work?

**Fiebrink:** Many different kinds of context can be relevant. If you think about systems like Jukedeck, these are applications that in some sense are straightforward AI generation systems, but the most compelling use cases are for casual YouTube creators who don't have the budget to hire a human composer and don't want to just get a stock piece of music - they want something that matches their video a bit better. In that kind of system, a lot of the ways that context could in theory be imposed would involve people communicating higher level goals of what they want the system to do. Systems like Jukedeck do that by asking questions such as: 'what genre do you want?' - which presumably influences which model from sets of models trained on slightly different training sets get chosen – 'how long do you want the piece to be?' and 'where's the peak moment?'. Those are really rough measures, but they are things that are easy to provide and certainly make a system much more useful than if you didn't have those features. Now there are probably lots of ways of extending sensitivity to context in worthwhile ways, but it's hard to be prescriptive in the abstract, because the types of information that people can give, or the aspects of context that matter, are just so dependent on the domain and the goal, so the starting point is simply understanding what things are important to people when they're doing creative processes.