AIAS
AARHUS INSTITUTE
OF ADVANCED STUDIES

AARHUS
UNIVERSITY

AIAS Conference:
Music and Artificial Intelligence: Pasts and Futures, Opportunities and Risks
Date: 28 May

Abstracts

Session 1: Histories

Music AI: Copyright, Compensation, Commons

Eric Drott

Since 2015 a number of startups have emerged seeking to commercialize music AI. Two trends stand out: one involves firms marketing services directly to consumers, in the form of adaptive music that responds to contextual and/or activity-related cues; the other involves companies marketing AI-generated music to other cultural producers, in the form of algorithmically-generated, royalty-free production music. Initiatives like these have generated debate among legal scholars about notions of copyright and authorship. But until recently discussion has focused on who (or what) should be awarded rights over the products of so-called "expressive AI": Its programmers? Its users? Or the AI itself? Largely overlooked in such debates is the status of another repertoire: not the music *put out* by an AI, but that which is *put into* it, the music that constitutes the training set necessary for machine learners to learn.

The dependence of machine-learning systems upon the human labor they threaten to displace raises important questions of distributive justice. To address them, this paper examines precedents for music AI that also relied on the "dead labor" of musicians (specifically, the "Music Composing Machine" developed by RCA in the 1950s), as well as proposals that have been advanced in response to more recent developments. I argue in particular that existing copyright regimes are ill-equipped to remedy the potential economic harms of music AI. Commercial practices premised on the extraction of value from a special kind of common-pool resource—the shared knowledge of a given music community—demand remedies grounded not in the possessive individualism of copyright law, but commons-based responses instead. Here, too, history provides precedents. Specifically, the Performance Trust Fund instituted by the American Federation of Musicians offers one model of how to respond to technology-induced unemployment, one that resulted from an earlier innovation in music technology: sound recording.

Biography: Eric Drott is Associate Professor of Music Theory at the University of Texas at Austin. He is author of *Music and the Elusive Revolution*(California, 2011). Current projects include *The Oxford Handbook of Protest Music*, co-edited with Noriko Manabe, and a book on the political economy of streaming platforms.

Cybernetics and AI in electronic music 1968-2018: Continuities, divergences, tensions

Christopher Haworth

The field of electronic music has been enduringly marked by the influence of cybernetics and artificial intelligence. The most tangible evidence of this is in the field's discourse, where vaguely defined 'cyber-' entities are frequently invoked in order to evoke an aesthetics of post-human machine agency, and 'artificial intelligence' has almost become a byword for a kind of computer programmed funk--think of Warp Records' 'Artificial Intelligence' series from the early 1990s, or Doris Norton's Apple-endorsed album of the same name from 1985. Cybernetics and AI had more than just a rhetorical influence on electronic music, though, informing its theory, practice and pedagogy; moreover, the two are not isomorphic, but are continuous and divergent. Disentangling them and drawing them into critical focus will be the aim of this paper.

I will approach the problem by analysing three case studies that cover a period from the heyday of cybernetics, to the beginning of the commercialisation of AI, to the present. The first is the influential Cybernetic Serendipity show from 1968 at the ICA. Alongside a public exhibition of computational and cybernetic art, Cybernetic Serendipity featured an exhibition catalogue and vinyl record that reframed the computer-assisted works of Cage, Xenakis, Peter Zinovieff, and Herbert Brün in terms of cybernetic open systems. The second case study I will examine is David Tudor's Neural Synthesis works from 1992. Where previously Tudor's works had incorporated audio and signal feedback systems that were analogous to the famous thermostat system of first-wave cybernetics, Neural Synthesis saw him work with an analog neural-network microchip that electronically emulates the brain's neuron cells, creating a system that would co-compose alongside Tudor's electronics. The third case study will be Young Paint, a project developed by Darren Cunningham who records under the moniker 'Actress'. Described as a 'learning program', Young Paint uses the recorded output of Actress as a training set, co-producing sonically-complimentary outputs alongside Cunningham himself.

Although part of my aim in this talk is to isolate the technical and aesthetic differences that characterise these systems, and thereby to distinguish the 'transition' from cybernetics to AI in musical terms, I do not aim to reduce the cybernetic and AI influence to only these domains. Since the differences in emphasis that take shape as these technoscientific concepts make contact with artworlds often open onto different articulations of race, ethnicity, gender and ability, the aim is also to bring the technical specification into productive tension with the hopes and desires that have been and are ascribed to cybernetics and AI by electronic musicians.

Session 2: Ethnographies

Machine Listening and Aesthetics: The Case of AI-Assisted Mastering

Jonathan Sterne

This talk will give a brief overview of research I have conducted with Elena Razlogova (Concordia University) on LANDR, a company that uses machine learning (branded as "Artificial Intelligence") to automate music mastering and create a platform strategy around it (supplemented with work on other AI-based audio applications from companies like Izotope). Some of our findings are that 1) the success of AI is defined in part by limiting the problems it is trying to solve while finding ways to market new uses for its products, thus offering a new chapter in the history of mastering, a set of audio techniques that has evolved alongside the history of audio formats; 2) LandR therefore offers an early test case for AI's effects on other media industries, showing that industries that are highly concentrated and whose work is mystified for everyday users are especially susceptible to automation; 3) there is a heavy "bullshit

quotient" in corporate discourse around AI: machine learning is often a single aspect or ingredient of a complex and multivalent set of social and technical processes, and may or may not be the most important one; and, finally, 4) LandR also offers an early test case for arguments about AI and labour, showing its effects on the labor force can be uneven and contradictory, shaped by the specific contours and limits of the industry rather than the "impact" of AI itself.

Biography: Jonathan Sterne is James McGill Professor of Culture and Technology at McGill University. He is author of *MP3: The Meaning of a Format* (Duke 2012), *The Audible Past: Cultural Origins of Sound Reproduction* (Duke, 2003); and numerous articles on media, technologies and the politics of culture. He is also editor of *The Sound Studies Reader* (Routledge, 2012) and co-editor of *The Participatory Condition in the Digital Age* (Minnesota, 2016). Visit his website at http://sterneworks.org.

Preferential Technics and the Cosmology of Overload

Nick Seaver

Why do music recommender systems exist? Ask someone who makes them, and you will likely hear this answer: listeners have so much music available to them that they are overloaded by choice. Within the relatively short history of digital music, many configurations of distribution have stood for this overwhelming archive, from CDs to MP3 downloads to on-demand streaming. Drawing on fieldwork with makers of music recommender systems and the algorithmic recommendation research literature, this paper describes the cosmology within which contemporary claims of overload operate. Understanding this cosmology, which has roots in mid-century cybernetics, is key to understanding how music recommendation makes sense to its makers: it optimizes information flows through channels both human and machine, recasting both taste and algorithmic personalization as techniques for managing the potentially unmanageable plenitude of the world. After Jonathan Sterne, I call this constellation of practices and theories "preferential technics." Preferential technics explains why recommender systems' integration of taste and technology, which many critics find paradoxical, seems reasonable to people who work on them. It also provides a starting point for analyzing new configurations of taste within emerging musical media assemblages (in Georgina Born's terms).

Biography: Nick Seaver is Assistant Professor of Anthropology and core faculty in the Program on Science, Technology, and Society at Tufts University. His research interests include ethnographic methodologies for studying algorithmic systems, the history of anthropological formalisms and their rejections, and the vernacular cultural theories of technologists. He is currently finishing an ethnography of the developers of music recommender systems in the US entitled *Computing Taste*. In a new project, he is studying the technocultural life of attention in machine learning.

Session 3: Accountability, Transparency and Ethics

The Social, Legal, and Algorithmic Study of Artificial Intelligence

Fernando Diaz

The past several years have seen the increasing development and deployment of machine learning and artificial intelligence systems in real world contexts. Notwithstanding the major advances in the theory and algorithms underlying modern artificial intelligence, as these methodologies have moved from offline laboratory experiments – either with fixed datasets or in simulation – to production systems, the complexities of the real world have surfaced for AI researchers, often in catastrophic ways. While our systems are capable of optimizing specific metrics, they often did so while violating social, ethical, and

legal norms. In response to these issues, Microsoft Research started the Fairness, Accountability, Transparency, and Ethics (FATE) research group focused on the societal implications of artificial intelligence. Because social questions require an understanding of algorithms, systems, and society, this group is composed of researchers with backgrounds in science and technology studies, machine learning theory, computational social science, and law. I will cover several of the major results from this group and focus on our research into the broader implications of using artificial intelligence for music information access.

Biography: Fernando Diaz is a principal research manager at Microsoft Research Montreal leading a multidisciplinary group studying Fairness, Accountability, Transparency, and Ethics (FATE). Previously, he was a director of research at Spotify. His primary research interest is information retrieval, including core ranking algorithms as well as evaluation. His work has won the best paper awards at multiple conferences. Fernando has served as the general chair for WSDM 2013, program chair for FAT* 2019, and will be a general chair for SIGIR 2021.

Session 4: Creative Interventions

Beyond Automated Creation: Making Machine Learning Useful, Accessible, and Understandable to Creative Practitioners

Rebecca Fiebrink

In my talk, I will describe my previous work creating machine learning (ML) tools for creative practitioners and studying their use. This work outlines a richer landscape for the usage of ML in creative practice than many people assume—for instance, ML can support the development of embodied interactions between people and technology, due to its ability to build models from examples (e.g., of movement) rather than relying on tacit and embodied knowledge to be expressed in code by a programmer. It can facilitate essential design activities such as rapid prototyping and exploration of a design space. It enables people without technical backgrounds (e.g., students, music therapists working with youth with disabilities) to create new interactive technologies without the need for programming. It also creates opportunities for more complex relationships between human creators and tools, inviting negotiation, surprise, and cooperation—a compelling alternative to the power dynamic embedded in building digital systems with code.

I will also briefly describe my work on teaching ML tools and techniques to creative practitioners and the general public. This work presents some evidence for the feasibility of teaching diverse, non-technical audiences about what machine learning can do, and for providing them the tools and knowledge to support meaningful discussion, ideation, and prototyping around their own visions for ML technology.

Biography: Rebecca Fiebrink is a Senior Lecturer at Goldsmiths, University of London, where much of her research and teaching focus on developing new tools and approaches for people to use machine learning in human creative practice. Fiebrink is the developer of the Wekinator software for real-time interactive machine learning whose current version has been downloaded over 25,000 times and used in hundreds of creative projects in music, art, dance, and games. She is the creator of the world's first online course teaching machine learning to creative practitioners, and she frequently teaches machine learning workshops to children, creative practitioners, and the general public. Current projects include developing machine learning tools for instrument-building for children with disabilities, developing machine learning based systems for aiding people with visual impairments in social interactions, and developing web-based deep learning tools for musicians.

Algorithmic Composition Subverted

Aaron Einbond

Recent narratives of artificial-intelligence for music promise to open, "advance,"[1] or "personalize"[2] creative practice. However, AI can also be used as a justification to narrow and "enclose"[3] existing materials and techniques under the guise of research and experimentation. How can machine learning instead contradict and explode creative expectations and the musical use cases for which it may have been intended? Situating my compositional practice among the fields of music information retreival (MIR) and computer-assisted composition (CAC), I will present three case studies of my own work that suggest possible approaches to machine learning subverted: in my composition *Xylography* from 2015, based on self-made software *CatOracle,* an approach to computer noise improvisation determines both the composed score for live cellist and the computer's unpredictable responses. In *The kind of problem a city is* from 2016, an instrumental score based on field recordings from New York City is transcribed for noise-based playing techniques of the ensemble of two keyboard players and two percussionists. And in upcoming project *Cosmologies,* microscopic sounds produced inside a grand piano will be projected into the concert hall spatially based on machine learning of measured instrumental diffusion patterns: highlighting the electronics' difference by attempting to compete with the richness of the live instrument's spatial presence.

Biography: Aaron Einbond's work explores the intersection of instrumental music, sound installation, field recording, and technology, focusing on audio transcription as the center of a creative process bridging composition, improvisation, interpretation, and questioning the thresholds of perception between instrument, loudspeaker, stage, and place. Recent collaborators include SWR Experimentalstudio, ZKM Karlsruhe, Académie du Festival d'Aix-en-Provence, Opera Lab Berlin, Ensemble Dal Niente, Yarn/Wire, Two New Duo, loadbang, and the Riot Ensemble. He teaches Composition, Sound, and Technology at City, University of London and is Co-Artistic Director of Qubit New Music Initiative in New York.

[1] Dannenberg, Roger. 'Artificial Intelligence, Machine Learning, and Music Understanding,' in *Proceedings of the Brazilian Symposium on Computer Music (SBCM2000)*, Curitiba, Brazil, 2000.
[2] 'AIVA – Music Engine', accessed 19 April 2019: https://aiva.ai/engine.
[3] Kate Crawford and Vladan Joler, 'Anatomy of an AI System: The Amazon Echo As An Anatomical Map of Human Labor, Data and Planetary Resources,' *AI Now Institute and Share Lab*, 2018.