

The Memory Cloud: Personal media libraries as affordance and constraint

Yann Seznec

yannse@kth.se

KTH Royal Institute of Technology
Stockholm, Sweden



Figure 1: Playing a Memory

Abstract

The Memory Cloud is a musical instrument that uses a player's own library of personal recordings as sonic material. This paper presents the design of the instrument, situating it within sustainability HCI studies and constraints-based design, before describing the instrument being used by two musicians in a professional context. Over 2000 sounds from the musician's personal cloud library, dating back over 10 years, were placed in the instrument as the only sonic material available for exploring. I argue that a radically small scale and personal approach could be one strategy for addressing the issues of longevity in NIME, and I suggest that using personal media libraries presents a potential affordance and constraint for musical instrument design.

Keywords

instrument, nime, cloud, sustainability, performance, design

1 Introduction

My cloud library is full. 34,638 photographs and 1,308 videos are taking up my 2TB allocation. As a designer working with sound it is not difficult to see the videos in particular as a form of sonic

content - it is raw, recorded mainly on my phone and full of wind noise, clipping, muffled by a poorly positioned hands. And yet it is unmistakably mine. No one else even has access to this library, built up almost accidentally over many years, representing virtually all of the major events in my life. Chances are, you have one of these libraries too (though you may manage yours better than mine). What would it mean to *use* these sounds? What kind of musical relationship can we build through encounters with this media?

This paper presents an instrument that takes a user's entire cloud library and repurposes it for musical use: The Memory Cloud. Pressing a button calls up a memory from the past and allows the player to break it apart or reconfigure it in any number of ways. In doing so the player builds a relationship with their own sonic history, unfolding new affordances for musical expression.

2 Background

2.1 Cornucopian design

Our overly saturated cloud media libraries, and indeed the entire existence of cloud storage, can be situated within what has been called the *cornucopian paradigm* of computing [16]. This is an approach to designing computer technologies with a fundamental assumption that all options and data should be made instantly available to the user at all times. Cloud libraries provide the potential for users to store all of their media, and the priority from a design standpoint is for the user to be able to seamlessly access



This work is licensed under a Creative Commons Attribution 4.0 International License.

NIME '25, June 24–27, 2025, Canberra, Australia

© 2025 Copyright held by the owner/author(s).

it anytime and anywhere. This is clearly unsustainable from a number of perspectives [9], with a growing body of sustainable HCI (SHCI) research calling for post-growth [20], degrowth [3], computing within limits [2, 15, 22] and deeper inner transitions [4] when considering the design of computing technologies.

2.2 Affordances and constraints

The work described here engages both with constraints-based design from a sustainability standpoint as well as in terms of the creative opportunities it presents. As discussed by Magnusson [11], among others, constraints and affordances are "two sides of the same coin", at least for musical instruments (pg 71). Design constraints, such as the number of strings on a guitar or the number of polyphonic voices on a synthesizer, are things that a musician discovers and engages with through exploration and play, potentially creating a relationship. That embodied, situated, and evolving dynamic can be seen as the affordances of an instrument. As Holzer argues through Barad's notion of agential cuts [8], the affordances of an instrument are thus comprised of the specific set of possibilities that are built between an instrument and player.

2.3 Longevity and NIME

The instrument described in this paper is by definition very personal, and thus it is unlikely to be developed in any large-scale way. There will be no "future developments" section in this paper where I suggest that a commercial version could be released; indeed, the paper will argue rather that the radically small-scale nature of this work is precisely what generates the musical affordances in the instrument. Nevertheless there are aspects of this work that link it to some of the recent discussions around longevity in NIME [5, 12]. In particular this work engages with concepts of *ownership*, which has been suggested as a potential mechanism for encouraging long term use of an instrument [14, 23]. Additionally, by foregrounding the conceptual positioning of the instrument within sustainable HCI we align with recent calls for more "outward-looking" work [13] to encourage wider discussions of the socio-political ramifications of designing musical technologies.

2.4 Personal sonic content

The specific constraint that is being using in the Memory Cloud is the sound library. Once the instrument has been set up for the player, there is no way of adding new sounds or recording samples into the machine¹. Instead, the only sounds available to be loaded into the granular engine come from the player's own cloud library. Admittedly this can be quite large (and in that sense pushes the boundaries of the term "constraint"), but nevertheless virtually guarantees that the sonic output will be associated with historical and emotional connection to the sonic source. Using what are essentially field recordings as musical content creates obvious parallels with the musique concrète practices of Schaefer [18] and Pade [1], among many others. Through the use of personal sounds in particular, the work described here aligns strongly with Matthew Herbert's concept of the "audible subject" [7] in terms of how the context of the source sound is crucial to the relationships that are built between listener, performer, and

¹As discussed in the "Design" section, this is not a technical so much as a design limitation, as the engine is certainly capable of adding or recording new sounds, and anyone with a screwdriver and Raspberry Pi experience could load new sounds into the instrument. However these features are hidden from the user and thus for the purposes of this work can be considered impossible.




Figure 2: The Memory Cloud

instrument (moving away somewhat from a strict acousmatic approach).

3 The Design

The Memory Cloud is an instrument based on a granular synthesis engine made in Pure Data. This engine has been developed over several years and used for a number of different projects². The precise nature of the granular code is not important for the purposes of this paper, and the specifics of granular synthesis are sufficiently well understood by the NIME community and do not need to be outlined here. For details on how the instrument itself functions and sounds please see the instructions document that is included in the code repository.

The instrument itself is comprised of the granular engine running on a Raspberry Pi, with a Teensy microcontroller taking analog control input from the interface elements and sending them over Serial to the Pure Data patch using the *comport* object³. The interface is made up of 8 knobs, one slider, four switches, and three buttons. The majority of these controls are mapped to granular manipulation parameters such as grain size, grain location, grain direction, pitch and location randomness, etc.

The sounds contained in the instrument are perhaps the most important, or at least the most distinctive aspect of the Memory Cloud. Although technically any folder of sounds could be used, the instrument is intended specifically to contain the sounds from all of the videos of the user's personal media library. In initial tests I used approximately 1,200 audio files extracted from videos in my own iCloud photo library, dating back 12 years. The version that was made for a musician (which will form the bulk of the analysis of this paper) included 2,632 sound files from roughly 10 years worth of cloud storage.

The sound library works as follows: a folder of sounds is exported from all of the videos in a user's cloud library and loaded into the local memory of the Raspberry Pi. When switched on for the first time the Pure Data patch in the device will "shuffle" these sound files and select one to be loaded into the granular synthesis engine for manipulation. A dedicated button on the instrument allows the user to change this sound file - when pressed, the next sound in the shuffled order will be selected and loaded. There

²The code is fully open source and available for download and use at <https://github.com/yannsez nec/ys.granular>, along with a detailed explanation video.

³<https://git.iem.at/pd/comport>

is no way of going back to the previous sound: Each time the players loads one sound they are also leaving the current sound behind. In order to play with a sound file a second time, the user must first load every other sound in the folder.

Preparing, exporting, and transferring these files is something of an undertaking. The most efficient method I devised involved downloading all of the videos from an iCloud library to a local drive, using a command line batch audio processing system (ffmpeg) to extract the audio, and finally batch converting it with SoX to the suitable audio format for the granular engine. The resulting sound files could then be loaded over SSH or USB to the Raspberry Pi's local storage.

This process ensures that each Memory Cloud is absolutely unique. While the engine which creates the sonic output, and the interface for controlling it, can be seen as a relatively straightforward and even fairly standard granular system, the content that forms the basis of this sonic output is tied inextricably to the owner of the device.

The initial design of the Memory Cloud took the form of a first-person Research through Design project, in which I developed and created the system for my own use using my own cloud library. The design stems from thinking around sustainability issues with regard to cloud storage and cornucopian design paradigms, as described previously, and the resulting instrument acts as something to think with [17], with knowledge emerging [6] from the object itself and its use. Working together with another musician extended this exploration one step further, giving me a glimpse of another perspective on the instrument and allowing me to play the role of mediator, as the process for setting up the instrument required a certain amount of specialist guidance. In this way the device offered a way to expand the inquiry of the design space beyond my own designing. It and the unfolding use became a carrier of the research questioning in which I was engaged around radically small-scale values embedded into the instrument.

4 Using the Instrument

The Memory Cloud was used for several months by musicians Sam Beste and Yoshika Colwell, using Sam's personal sound library. During the initial familiarization period the instrument was in Sam's studio and was used by visitors:

Sam: So it's just been sort of set up in my studio, and just as people have been passing through they've been using it, which has been really nice, and that's the way that I've most enjoyed interacting with it, in a way, is with me playing an instrument and somebody else actually playing [The Memory Cloud] so that there's an interaction.

After some time it became clear that the instrument could work well for a concert at a large gallery in London, particularly as Yoshika built a relationship with it:

Sam: I made this collaborative EP with this singer, Yoshika, which kind of came together through us improvising together in a way, and then building songs out of those improvisations. And so when we were offered this opportunity to do this thing at the Gallery, you know, we only have four songs, and they wanted us to play 40 minutes. And at first I wasn't really sure how we would do that. But then, actually, when [Yoshika] came into the studio, the idea was like, well, why don't we play

the four songs but in between the songs, we sort of just improvise? And having the Memory Cloud, it was actually quite a transformative experience for her, because she's often been in environments where she isn't given control of an instrument, or she's...been stuck in certain kind of ways of perceiving herself and others perceiving her. And actually having this box, she was completely amazing at performing with it as well...

Yoshika: I played it a few times leading up to that gig and then I played it at the gig and it was amazing. I absolutely love it, it's one of my favorite things I've ever used. Because I'm not someone who plays with samplers [or other music technology]. I often feel slightly alienated by them...With the Memory Cloud I immediately took to it and I found it very very accessible, which was so exciting for me.

The inherently personal nature of the device was one of the key elements that activated the instrument for Yoshika in particular:

Yoshika: The beauty of it is the personalized element, so I'm sitting there playing with Sam and I can hear "oh that's your kids" and it's just so beautiful and moving. All of those sounds immediately have a weight to them, even though I'm not in his family. There's something very humanizing about that, I really like that.

Sam: I think part of her joy in using it was knowing that it's completely singular and that it's completely related to my own life, and that the things that she's relating to and is playing are so personal and it's like this sort of process of discovery...That's what makes it magical.

The impermanence of the sound selection seems to have been an important part of what made the instrument work well in this context, with Sam eventually beginning to see future potential for building work around it:

Sam: But she also just really loved the concept of it, and loved how the kind of impermanence of it, you know that you can't get attached to anything, because you'll never get there again. And that really worked for the nature of performance, and I think that kind of impermanent nature of something in your setup just invites you to kind of be present in a way, and not be sort of tied to things that are going to occur every time and you know are going to be there. It kind of protects you from stagnation in some way, which I really love. And I'm not entirely sure how I'm going to live without it, to be honest, because I sort of feel like I keep imagining ways that it's going to be involved in different aspects [of my work].

As the designer of the object I was surprised to hear from Sam that he had so readily given this instrument, the sonic contents of which were so personalized and particular to his own life, over to others to play. However he explained that this was a result of his desire to respond musically (mainly through the piano) to processed sound recordings:

Sam: It has fit so well into a world that I was already exploring, but didn't have a means of doing it in a kind of this generative way, live. And so it

was always going to be somebody else playing it in a way, because I'm always interested in the interaction between the piano and sound...So actually, having somebody else control it sort of just makes a lot of sense.

5 Reflections

The use of the Memory Cloud suggests three specific themes that can be discussed here: the specific potential of radical small scale design, how the design of the instrument can be seen as a constraint that encourages creative output, and how the use personal sounds acts as a musical affordance.

5.1 Radically microscale design, politics and longevity

The Memory Cloud is an instrument that demands care. From the outset, my own role of designer involves a significant amount of practical work and guidance, with the eventual goal of collecting a user's own personal sounds and creating a customized version of the instrument. On the one hand this is an illustration of the inefficiency and biases of the technological systems that we are entangled with. Each step of the process is frictional and demonstrates how the systems on which we rely for managing our personal media are simply not designed for these purposes. Clicking and dragging 2000 video files from a photo management application immediately creates graphical and functional slowdowns of the operating system even on a relatively new computer. Exporting that many files is simply not expected behaviour - the iCloud browser system, for example, only allows a user to download 1000 videos at a time (and does not function reliably).

Indeed, the process of creating Sam's Memory Cloud from his cloud library was challenging. I needed to guide him through the process of exporting his entire video library from iCloud, installing and using command line batch processing tools for the first time, and sending me the resulting sounds to load onto the instrument. Working through these steps with Sam, whilst living in two different countries, took several days (and many text messages), and gave him some understanding of both the scale and power of his own media:

text message from Sam: in trying to find a way through i accidentally opened all 4.5k videos⁴ with quicktime omg! it almost collapsed my system! haha! But was really emotional to see all the videos popping up one by one...quite amazing that we have this archive of our lives

This process was essentially reactivating a relationship between Sam and his own media library, a relationship in which I as the designer was acting as something of a mediator or guide, carefully helping him collect and prepare his personal memories for use in his own instrument.

While the technical system that underpins the Memory Cloud is a relatively straightforward granular synthesis engine, the use of the personal media library as the only sonic content available in the instrument is a constraint that serves to strengthen its emotional power. The Memory Cloud takes a radically anti-innovation approach that explicitly uses an existing technical engine and instead is about building a connection between a specific user and their own personal memories. Put simply, each

Memory Cloud is unabashedly a singular instrument for a single person. My own Memory Cloud could not possibly have the same power for someone else - and that is precisely the point.

Thus the core elements that come together to form what I call "radically microscale design" are:

- A user willing to engage with their own media library and confront the biases embedded in our technologies
- A thoughtful and caring relationship between designer and user
- An overt avoidance of technological innovation

The way this concept dovetails with sustainability and computing issues as detailed in the background also points towards a form of media reuse, which could form the basis of more work in the NIME community. Where systems such as sample libraries, physical modeling, and latent spaces all push for creating *new* sounds, what about the thousands of sounds that we are all carrying around in our pockets? Perhaps these sounds are worth considering as material for creative expression - particularly as many of them will contain far more emotional weight than anything generated by machine learning or other methods. Therein lies the political stance in the work, which is both embedded into the design of the instrument and is revealed through setup and use.

The way this was manifested in the Memory Cloud, through the use of the vast personal library spanning a number of years, prompted reflections from Yoshika around the nature of these sounds and how they differ from musical material that is often used in musical composition and performance:

Yoshika: It's such a scrapbook of somebody's life, so intimate. That ambient noise, it's not constructed, which does resonate with the philosophy of the [instrument]. It's like...all of the background stuff is all of the important stuff, really. The stuff of everyday life, the noise, the laughing, the rustling, it's not performative in any way. And I think that is so mirrored in the instrument, because it feels like it's a tool to stimulate your thoughts about those sounds.

It has been suggested that a preoccupation with technical innovation and an avoidance of political stances has potentially contributed to the problems with NIME longevity [13]. I would like to position the Memory Cloud within these discussions that have taken place within the NIME community, suggesting that the creative process for this instrument stemmed from a specific political stance based on anti-solutionism, sustainable HCI, care, and radically small scale design. Whether or not this results in an instrument with longevity remains to be seen, but I present this work as a way of attempting to actively engage with the discussions.

5.2 Constraints of personal sonic material

Having briefly touched upon the results of a constraints-led design process, it is worth considering exactly how these are applied in the design of the instrument and how the resulting interaction manifests.

It is immediately noticeable that the instrument itself has no screen. Whilst some discussions around musical instrument design has suggested that moving away from screens might be a small step towards sustainability [19], in practice there are also cost and technical reasons for this approach. The instrument runs on a Raspberry Pi, which could certainly have a screen attached

⁴the final number was smaller due to file compatibility, doubles, and naming issues, but was deemed sufficient in any case

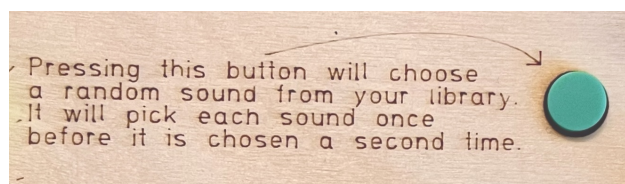


Figure 3: Choose a new sound, leave the old sound behind

through HDMI, but integrating this into the user interface would likely have created more problems than it solved. Nevertheless an instrument that requires a sound file to be loaded in order to start making sound (and contains potentially thousands of files to load) will certainly require some method of choosing a sound file to load.

This constraint is what led to the button for loading the sound into the granular engine. This button is prominently featured in the middle of the device, accompanied by an explanation (Figure 3). As mentioned earlier, the entire sound library is shuffled, so the order of loading the sounds is predetermined but unknown to the player.⁵ The next sound file may be from a video recorded on the same day as the previous, or could be years apart. It could be a short clip of a rehearsal or it could be the sound of a dog running on a beach. The only thing that is certain is that pressing the button will leave the current sound behind, and you will not hear it again until you have heard every other sound in the library - with over 2000 sounds in Sam's library it is unlikely that he will be hearing any given sound again any time soon.

There is a peculiar balance to this, where there is a combination of unknown (which sound will come next) and known (all of the sounds are from a single and personal source). It is likely that the player will recognize any given sound, though the order in which they are presented will be surprising and potentially jarring. This mirrors the dual nature of the design constraint - on the one hand this is an instrument filled with a staggering amount of sonic content. In Sam's case the total size of the library came to 46 gigabytes, or roughly 80 hours of audio. This could, admittedly, be seen as adopting the very cornucopian design paradigm that I was trying to avoid, in the sense of providing the user not just with quite a lot of sounds, but with *every single sound they have ever recorded*. However, by only providing a single button as the interface for selecting a sound (and by overtly avoiding a screen to enable selection or give information on upcoming files), this glut of sound is transformed instead into a constraint where each button press must be carefully considered. Is it time for a new sound? Is it worth sticking with this one a bit longer? When will I hear it again?

For Yoshika in particular this constraint was a liberating experience from her very first time playing the Memory Cloud:

Yoshika: I was playing one sample and [Sam] was like "the fun thing about it is you go forward and you can never get back [to the previous sound]" ...So you can never get attached to it, which I absolutely love. And immediately you make your peace with that, just psychologically for me, you can't get stuck. I think that's because you immediately accept that you can't be attached to a sound...I think in terms of not being attached, and how generative that is

creatively, I think that's amazing and very freeing because you feel like you are improvising with it, it feels like a two way thing with the machine...you're not attached to the outcome being exactly the same. You want to capture the feeling and the essence rather than keep something in a very constricted space.

5.3 Personal media as affordance

Sam's relation with his Memory Cloud was intensely personal, describing it as "singular and...completely related to my own life". This is perhaps not surprising as it contained only his own sounds from his own media library, though it also seems likely that the process of exporting the videos and preparing the sounds contributed also to his connection to the finished instrument. Playing the instrument did not exist in a vacuum - Sam began building a relationship with the library of sounds well before he received the Memory Box in the post. By the time he received it he was already aware of the personal nature of the sounds that it would create, as well as having a handle on the sheer quantity of files and the associated technical challenges of using them.

These interwoven relationships between myself (the designer), Sam and Yoshika (the performers), the instrument, and the personal media content can be seen as reflecting the complex and entangled nature of musicking with technology [10, 21]. The Memory Cloud creates affordances not primarily through interface or sound design, but instead through the use of personal media. It leverages sonic content which by its very existence will already be filled with meaning, particularly for the one individual who created that content. The affordances that make the Memory Cloud function are built upon and with the user's past - potentially years of memories that are re-configured and presented back to the creator, generating encounters with their own history and giving the opportunity to dissect a singular moment from their own experience.

6 Conclusion

The Memory Cloud is an instrument that uses personal media to generate new affordances for musical interaction. It leverages our vast libraries of videos towards a radically small scale and personal experience, embracing constraints as a design strategy and a political stance.

7 Ethical Standards

This research took place at KTH Royal Institute of Technology in the context of my doctoral studies. The research is exempt from formal ethical approval under the interpretation of Swedish law⁶. Data collection does not directly relate to sensitive personal data. The musicians involved in testing the instrument (Sam Beste and Yoshika Colwell) were informed from the beginning of the research context of the work and Sam fully consented to using his own sound recordings in the instrument, aware that this formed an integral part of the design of the instrument and the resulting research. Both Sam and Yoshika consented to our conversations being recorded and transcribed for the paper. They were not compensated financially for their participation. The research also necessarily involved the use of my own personal media library in the development of the instrument. While my supervisory team

⁵When the device is switched on it will automatically reload the same sound that was in the buffer when the instrument was switched off. The shuffled order is retained.

⁶https://www.riksdagen.se/sv/dokument-och-lagar/dokument/svensk-forfattningssamling/lag-2003460-om-etikprovning-av-forskning-som_sfs-2003-460/

was fully supportive of this work, the use of my own sounds was entirely of my own volition and was not the result of any coercion from supervisors or anyone else in a position of power.

Acknowledgments

This work is supported by the Energimyndigheten (Swedish Energy Agency) project *Interaktionsdesign för ett Soldrivet Internet* (Interaction Design for a Solar Powered Internet), 51885-1.

I would like to thank my supervisor Rob Comber for advice and support throughout this project.

Special thanks also to Sam and Yoshika for their open minded but critical embrace of the concepts, constraints, and outcomes of using the instrument.

References

- [1] Andrea Bak. 2009. *Else Marie Pade: en biografi*. Gyldendal, København. OCLC: 667231490.
- [2] Jay Chen. 2016. A strategy for limits-aware computing. In *Proceedings of the Second Workshop on Computing within Limits*. ACM, Irvine California, 1–6. <https://doi.org/10.1145/2926676.2926692>
- [3] Giacomo D'Alisa, Federico Demaria, and Giorgos Kallis (Eds.). 2015. *Degrowth: a vocabulary for a new era*. Routledge, Taylor & Francis Group, New York ; London.
- [4] Elina Eriksson and Daniel Pargman. 2023. Exploring Inner Transition: Expanding Computing for Sustainability. In *Ninth Computing within Limits 2023*. LIMITS, Virtual. <https://doi.org/10.21428/bf6fb269.22dea4ad>
- [5] Alessandro Fiordelmondo, Giada Zuccolo, Sergio Canazza, and Raul Masu. 2024. Longevity in NIME research: a case study using time-based media art preservation models. (Oct. 2024). <https://doi.org/10.5281/ZENODO.13904858> Publisher: Zenodo.
- [6] William Gaver, Peter Gall Krogh, Andy Boucher, and David Chatting. 2022. Emergence as a Feature of Practice-based Design Research. In *Designing Interactive Systems Conference*. ACM, Virtual Event Australia, 517–526. <https://doi.org/10.1145/3532106.3533524>
- [7] Matthew Herbert. 2020. *Listen while you work: negotiating power and meaning in post-concrete music*. PhD Thesis. Canterbury Christ Church University. <https://repository.canterbury.ac.uk/item/8x378/listen-while-you-work-negotiating-power-and-meaning-in-post-concrete-music>
- [8] Derek Holzer. 2024. A Piano for Visuals: Affordances of Scan Processing Instruments. *Leonardo* 57, 6 (Dec. 2024), 613–621. https://doi.org/10.1162/leon_a_02584
- [9] Esther Jang, Matthew Johnson, Edward Burnell, and Kurtis Heimerl. 2017. Unplanned Obsolescence: Hardware and Software After Collapse. In *Proceedings of the 2017 Workshop on Computing Within Limits*. ACM, Santa Barbara California USA, 93–101. <https://doi.org/10.1145/3080556.3080566>
- [10] Alexander Refsum Jensenius. 2022. *Sound Actions: Conceptualizing Musical Instruments*. The MIT Press, Cambridge.
- [11] Thor Magnusson. 2010. Designing Constraints: Composing and Performing with Digital Musical Systems. *Computer Music Journal* 34, 4 (Dec. 2010), 62–73. https://doi.org/10.1162/COMJ_a_00026
- [12] Raul Masu, Fabio Morreale, and Alexander Refsum Jensenius. 2023. The O in NIME: Reflecting on the Importance of Reusing and Repurposing Old Musical Instruments. (May 2023). <https://doi.org/10.5281/ZENODO.11189120> Publisher: Zenodo.
- [13] Fabio Morreale, S. M. Astrid Bin, Andrew McPherson, Paul Stapleton, and Marcelo Wanderley. 2020. A NIME Of The Times: Developing an Outward-Looking Political Agenda For This Community. (June 2020). <https://doi.org/10.5281/ZENODO.4813294> Publisher: Zenodo.
- [14] Fabio Morreale and Andrew McPherson. 2017. Design For Longevity: Ongoing Use Of Instruments From Nime 2010-14. (June 2017). <https://doi.org/10.5281/ZENODO.1176218> Publisher: Zenodo.
- [15] Bonnie Nardi, Bill Tomlinson, Donald J. Patterson, Jay Chen, Daniel Pargman, Barath Raghavan, and Birgit Penzenstadler. 2018. Computing within limits. *Commun. ACM* 61, 10 (Sept. 2018), 86–93. <https://doi.org/10.1145/3183582>
- [16] Chris Preist, Daniel Schien, and Eli Blevis. 2016. Understanding and Mitigating the Effects of Device and Cloud Service Design Decisions on the Environmental Footprint of Digital Infrastructure. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, San Jose California USA, 1324–1337. <https://doi.org/10.1145/2858036.2858378>
- [17] Johan Redström. 2017. *Making Design Theory*. The MIT Press. <https://doi.org/10.7551/mitpress/11160.001.0001>
- [18] Pierre Schaeffer. 1998. *À la recherche d'une musique concrète* ([nachdr. der ausg.] paris, seuil, 1952 ed.). Seuil, Paris.
- [19] Yann Seznec. 2022. Music Within Limits. In *2022 International Conference on ICT for Sustainability (ICT4S)*. IEEE, Plovdiv, Bulgaria, 172–178. <https://doi.org/10.1109/ICT4S55073.2022.00028>
- [20] Vishal Sharma, Neha Kumar, and Bonnie Nardi. 2024. Post-growth Human–Computer Interaction. *ACM Transactions on Computer-Human Interaction* 31, 1 (Feb. 2024), 1–37. <https://doi.org/10.1145/3624981>
- [21] Christopher Small. 1999. Musicking — the meanings of performing and listening. A lecture. *Music Education Research* 1, 1 (March 1999), 9–22. <https://doi.org/10.1080/1461380990010102>
- [22] Martin Stojanov, Daniel Pargman, Mike Hazas, Rob Comber, and Jorge Luis Zapico. 2023. How do we arrive at constraints? Articulating limits for computing. In *Ninth Computing within Limits 2023*. LIMITS, Virtual. <https://doi.org/10.21428/bf6fb269.a317d18f>
- [23] Isaac Wallis, Todd Ingalls, Ellen Campana, and Catherine Vuong. 2013. Amateur Musicians, Long-Term Engagement, and HCI. In *Music and Human-Computer Interaction*, Simon Holland, Katie Wilkie, Paul Mulholland, and Allan Seago (Eds.). Springer London, London, 49–66. https://doi.org/10.1007/978-1-4471-2990-5_3 Series Title: Springer Series on Cultural Computing.