Journal SEAMUS

Spring/Fall 2021

The Society for Electro-Acoustic
Music in the United States

Volume 32, Number 1–2

ISSN 0897-6473

Contact Information
Drake Andersen
Department of Music
Central Connecticut State University
1615 Stanley Street
New Britain, CT 06050

journal@seamusonline.org

Journal SEAMUS
Volume 32, Number 1–2
Spring/Fall 2021

ISSN 0897-6473

Please visit us at www.seamusonline.org to find out more about SEAMUS

Information for Contributors

Interested contributors should submit manuscripts electronically. Microsoft Word is the preferred format. If another word processor is used, files should be saved in rich text format (RTF) with an accompanying PDF version. Main articles are generally 2,000 to 6,000 words.

Editing Guidelines Please use Times New Roman fonts with font size 12. Manuscripts should be formatted and prepared using *The Chicago Manual of Style,* 17th edition (2017) as a guide. References should follow Author-Date format. Specific citations should be provided in text in parentheses. Footnotes should be used sparingly and reserved for explanation beyond the text of the article. All references should be listed after the text of the article in a section labeled “References.” Any computer code should be placed in fixed-width format to facilitate readability. Images, figures, musical examples, and other graphics should be sent as separate attachments for ease of layout. The approximate location of each graphic should be indicated in the text by a (sequentially numbered) label and a brief caption.

Graphics Any artwork, graphics, photos, and flowcharts should be sent as separate individual files. We recommend uncompressed graphic files such as TIFF at 300 dpi.

Submission All submissions, including articles, reviews, review proposals, and items for *Tips and Tricks* should be emailed to the Editor-in-Chief, Drake Andersen: journal@seamusonline.org.

About SEAMUS

Founded in 1984, the Society for Electro-Acoustic Music in the United States (SEAMUS) is a national non-profit organization of composers, performers, and teachers of electro-acoustic music representing every part of the country and virtually every musical style. Electro-Acoustic music is a term used to describe those musics that are dependent on electronic technology for their creation and/or performance. Many members of SEAMUS—like Jon Appleton, the guiding light in the conception of the Synclavier—are recognized world leaders in their fields. All are dedicated to the use of the most advanced technology as the tools of their trade.

SEAMUS seeks to provide a broad forum for those involved or interested in electronic music. Through its journal, newsletter, national meetings, and its national archive at the University of Texas, SEAMUS seeks to increase communication among the diverse constituency of the relatively new music medium.

The Society's objectives include:

To encourage the composition and performance of electro-acoustic music
To develop a network for technical information and support
To promote concerts and radio broadcasts of electro-acoustic music both in the US and abroad
To create an exchange of information through newsletters and other means of communication
To establish and maintain a national archive and information center for electro-acoustic music
To attract a wide diversity of members and supporters
To advocate licensing and copyright concerns

SEAMUS strives to address not only relevant technology but also the non-technical issues pertinent to the electro-acoustic music community. In a field usually dominated by technical concerns, it is refreshing to hear paper sessions devoted to aesthetics, collaboration, education, and the ethical and social issues facing electro-acoustic musicians. The provocative sessions provide fuel for lively discussions during the national meetings.

SEAMUS Board of Directors

Contents

Journal SEAMUS

Volume 32, Number 1–2, 2021

From the Editor

It is with great pleasure that I announce the publication of Journal SEAMUS Volume 32. In this issue, you will find a wide variety of scholarship that touches on a range of technical, aesthetic and humanistic concerns. Kicking things off is Mei-ling Lee's study of data-driven instruments, which explores design and mapping considerations across a range of software and hardware platforms, including Kyma, Max, the Gametrak controller, the Leap Motion, and the Wacom tablet. George Edmondson's article examines the intersection of aesthetics and ethnography, focusing on site-specific work in Birmingham, UK. Finally, a multi-authored contribution by Ted Moore, James Bradbury, Pierre Alexandre Tremblay, and Owen Green reflects on teaching using the Fluid Corpus Manipulation toolkit, or FluCoMa, a powerful tool for creativity and pedagogy that may already be familiar to many in the SEAMUS community.

The cover art for Volume 32 is an illustration created by composer, filmmaker, and multi-instrumentalist Eliza Gelinas.

As we begin the new year, the Journal continues to seek new ways of serving and connecting with the SEAMUS community, including through new calls and initiatives. If you have an idea for something you'd like to see in the journal, please don't hesitate to drop us a line at journal@seamusonline.org!

We are always on the lookout for new articles and content—especially submissions that take advantage of our new digital format. (For example, you'll notice that the article on FluCoMa is full of helpful links to online resources.) For more information about submitting to Journal SEAMUS—as well as the latest additions to our ever-growing digital archive—check out this page: https://seamusonline.org/journal-seamus/.

As always, I will reserve the last word to express my appreciation to the Journal SEAMUS staff, SEAMUS leadership, and all of the SEAMUS members who continue to support the Journal's mission. I look forward to hearing from you!

Drake Andersen, Editor-in-Chief

Exploring Data-Driven Instruments in Contemporary Music Composition

Mei-ling Lee

Department of Music
Haverford College
mlee4@haverford.edu

Abstract

How can data-driven instruments, through data mapping strategies and the use of performance interfaces, transform compositional practices and expand the creative possibilities of contemporary music production?

This paper illustrates the concept of data-driven instruments and their connection to conventional musical instruments. It examines three original compositions written by the author to demonstrate the utilization of data-driven instruments in musical contexts, focusing on the application of sonic materials and data mapping strategies. Developed using Max, a programming language by Cycling '74, and Kyma, a sound creation environment by Symbolic Sound, these compositions utilize diverse performance interfaces to explore interactive, real-time possibilities. By exploring these instruments' innovative potential, this paper demonstrates their transformative role in contemporary music creation, offering new potentials and expanding the horizons of musical artistry.

Introduction

Throughout the chronicle of human history, music has consistently served as a crucial vehicle for comprehending and distributing humans' life experiences: from religious ceremonies in ancient Greece (Grout 2001, 3), to the Italian madrigals at a variety of aristocratic social gatherings of the 16th century (Grout 2001, 184), to more recently, the electronic dance music that served as a significant social-defining element for the millennial generation (Matos 2015). In contemporary times, the infusion of modern technologies into music creation has yielded transformative consequences.

The evolution of music-making tools, such as simple MIDI keyboards and alternative electronic music controllers, has initiated the start of an era where important uses of data-driven instruments can be realized.

In the context of our increasingly data-driven world, where electronic devices are omnipresent, the relevance of traditional musical instruments persists, yet there arises a compelling need for fresh perspectives and innovative approaches to musical creation. Given the ubiquity of data in our daily lives, utilizing its power to articulate our musical ideas seems entirely appropriate. Data-driven instruments emerge as a pertinent choice to fulfill this aspiration.

This paper aims to demonstrate the concept of data-driven instruments by presenting three original compositions as case studies. These compositions are interactive, real-time, and multichannel works, developed using Max[1], a programming language by Cycling '74, and Kyma[2], a primary sound creation environment by

[1] Created by Miller Puckette at IRCAM, Max is an object-oriented programming language for the production of music and multimedia. Additional information can be found about Max at https://cycling74.com/products/max

[2] Kyma is a domain-specific programming language that includes a software component created by Carla Scaletti and a hardware component created by Kurt Hebel. The complete Kyma system is dedicated to synthesizing, modifying, exploring, and constructing sound, and to the creation and performance of musical compositions.

Symbolic Sound. The three compositions – *Giant Dipper*, *Farewell*, and *Summoner* – utilize a range of performance interfaces and sound synthesis engines, thereby exemplifying the creative possibilities engendered by data-driven instruments in contemporary music composition.

Conceptual Background and Theoretical Frameworks

Many concepts and techniques are involved in creating electronic music compositions for data-driven instruments. In this chapter, I provide an overview of the most essential concepts and techniques important to the three compositions presented in this paper.

The Structure of Data-Driven Instruments

Data-driven instruments comprise three main components: data acquisition through an interface, data mapping and routing, and sound-producing algorithms (Stolet 2013). When playing traditional instruments, such as piano or violin, we exert force into their physical systems to create a sound. However, when working with data-driven instruments, a performer first operates an interface to generate the data. The data functions as a replacement for energy within the traditional performance model. This data is then processed through software layers, analyzed, recalculated, and mapped to specific ranges. The resulting data stream is directed to the sound-producing algorithm to generate sound. The architectural framework of data-driven instruments is depicted in Figure 1.

One common misinterpretation of data-driven instruments lies in the concept of tracking and the notion of gesture control. In performing with data-driven instruments, the interface itself is not tracking the performer's movements, but is instead reporting the data stream related to the current status of the performance interface. Physical "gestures" do not create or control sounds; however, during a data-driven musical performance, *performance actions* produce data that then can be used to control the sound-producing mechanism of the data-driven instrument.

Instrumental Modularity

Modularity is the foundation for constructing various entities (Stolet 2021, 25-27). Instrumental modularity is also a key feature of the data-driven instrument. The concept of modularity centers on the idea of combining and connecting smaller and simpler things to construct larger and complex things.

Most traditional instruments are contained within a single physical structure. For instance, a piano encompasses keys, pedals, hammers, and a soundboard, all assembled into one body. In contrast, the components of a data-driven instrument can be considered as independent entities or as part of the instrument.

In this approach, elements like the interface, the mapping mechanism, and the sound-producing algorithm operate as distinct modules.

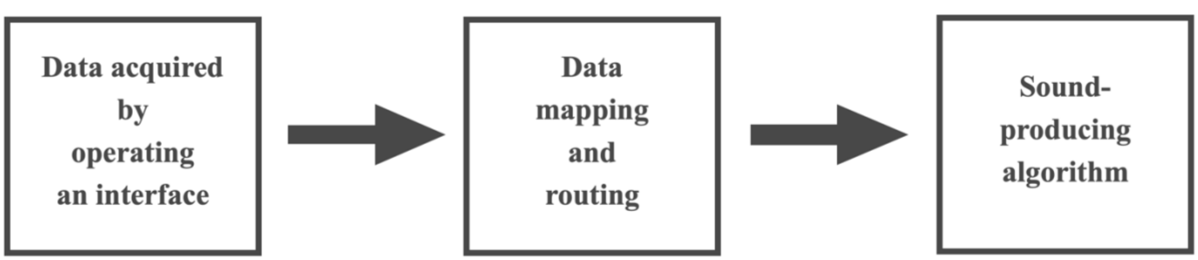

Figure 1. The structure of data-driven instruments.

For example, we can use a Gametrak[3] as the interface for a data-driven instrument, then use Max as the data mapping area of our instrument. Max then sends the newly mapped data to Kyma, the sound-production mechanism. On another setup, we can take the same Gametrak as the performance interface to create another composition and route the data for sound production to Logic Pro instead of Kyma as shown in Figure 2.

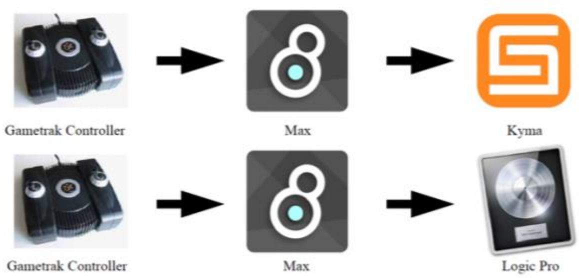

Figure 2. Modularity features on data-driven instruments.

Instrumental Mutability

One of the main features of a data-driven instrument is its mutability (Stolet 2021, 81), allowing significant sound changes during a performance. Traditional instruments have a limited capacity to change their sound during the performance. Mutes on trumpets and stringed instruments, damper pedals on pianos, and different mallets for percussion instruments are examples of mechanisms for mutability of traditional instruments.

For data-driven instruments, however, the ability to mutate is expanded dramatically. Unlike traditional instruments in which mutability happens in the physical hardware of the instrument, mutability of data-driven instruments predominantly occurs at the software layers through data mapping, routing, and sound production.

Data Mapping

Data mapping is a vital process in data-driven instruments, converting interface-generated data values into new output values. It plays a critical role in effectively linking various interfaces to sound-producing algorithms, allowing control over musical parameters.

When performing with a data-driven instrument, the raw data acquired from operating the interface might not be directly suitable for sound production. For example, we may receive a data stream from a MIDI controller with a data range from 0 to 127. We want to use this data stream to generate some MIDI notes. However, if we use this raw data directly, we are likely to produce notes that are pitched too low and produce undesirable musical consequences. To resolve this issue, a data mapping process is employed to convert the data from its original range to a more suitable one that achieves the desired musical outcome. This transformation occurs in software layers, adapting the data to a range better suited for the sound-producing algorithm.

Synthesis Techniques

There are a variety of sound synthesis techniques applied in the three example compositions. These techniques include sampling, additive, subtractive, and granular synthesis, as well as analysis and re-synthesis, and Kyma's Time Alignment Utility (TAU) algorithm.

Sampling synthesis includes playback of the original audio sound source into a digital medium and applies various audio editing techniques to transform the recording such as looping and pitch-shifting. Additive synthesis involves combining simple waveforms at various frequencies, amplitudes, and phases to create more complex waveforms. While additive synthesis involves combining frequencies, subtractive synthesis shapes the sound by removing or attenuating frequencies of a waveform by applying various types of filters.

[3] The Gametrak is originally designed for a virtual golf video game, based on a 3-dimensional position tracking game control system.

Granular synthesis, a pivotal technique in my compositions, involves breaking down samples into tiny sonic particles known as "grains," typically lasting 5-100 milliseconds. This method enables the manipulation of sound to create evolving soundscapes by controlling parameters like grain duration, texture density, and spatialization. Analysis and resynthesis involve studying a recorded sound and utilizing that data to govern the synthesis process (Roads 1996).

In Kyma, the TAU algorithm is crucial for modifying and morphing sounds by blending amplitudes, frequencies, formats, and bandwidths. For example, one can combine the amplitude characteristics of a sound with the frequency characteristics of another sound, then mix in the formant characteristics of the third sound, to create a unique sound result (Stolet 2012, 171).

Musical Works

Composition: Giant Dipper

Giant Dipper is an example of transforming sound materials taken from human life experiences by repurposing the Gametrak spatial-position controller. This composition aims to convey two distinct journeys: one simulating the experience of a roller coaster ride, while the other encapsulates the thrill and speed of a car race seen through the perspective of a seven-year-old girl's mind.

The Gametrak controller serves as a performance interface, transmitting data to Max for mapping and routing to Kyma to control real-time sound algorithms. Initially designed as a floor-based unit for golf simulation on PlayStation, the Gametrak base unit, featuring retractable nylon cables and internal gears, allows for three-dimensional articulation of X, Y, and Z axes. Because of its ease of use, expressive potential, and affordability, the Gametrak spatial position controller has become one of the most popular and reliable performance interfaces for data-driven instruments.

Giant Dipper's sound materials derive from two primary sources: a home recording of a seven-year-old girl engaging in imaginative play, along with field recordings of the roller coaster at the California Santa Cruz Beach Boardwalk amusement park. The girl's singing was captured using a SONY Linear PCM-D50 recorder, while the roller coaster sounds were recorded via the Voice Memo app on an iPhone. To avoid potential mishaps, I held the iPhone to record ambient park sounds but secured it in my purse during the roller coaster ride. Embracing the incidental sounds of the environment, I accepted the resultant sonic artifacts within the recordings, later using an equalizer during editing to minimize wind and unwanted noise for improved clarity.

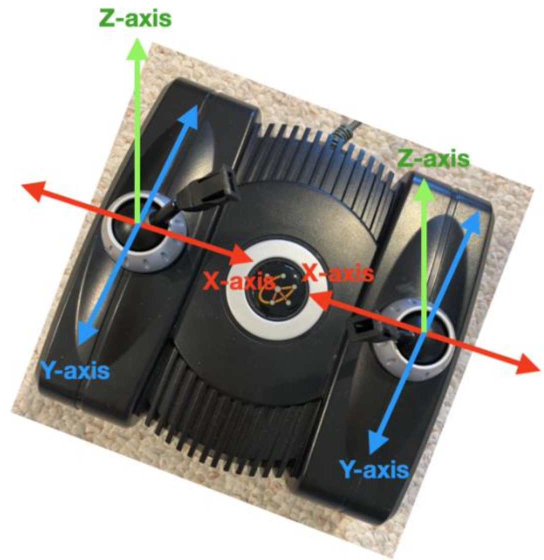

Figure 3. Gametrak X, Y, and Z Cartesian coordinates.

This composition uses various audio playback techniques to present the sound material. For instance, the initial sound is that of a mix of a crowd chanting in cheers at the theme park along with the playback of a resynthesized voice of a girl. In this instance, the left-hand pulling controls the amplitude of the crowd chanting and the right hand controls the amplitude of the resynthesized girl’s voice.

Where, in virtual space, a sound occurs is vital as I simulate the experience of a roller coaster ride. Spatialization of the sounds was controlled by the Y-axis of the right-hand cable. Figure 4 shows the Kyma sound design of the panning.

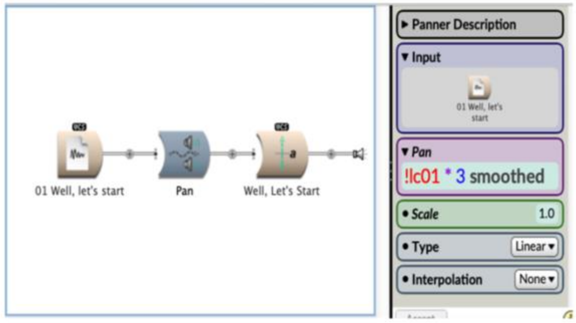

Figure 4. Kyma sound design shows right-hand Y-axis control the panning.

Along with segments of audio that are directly played back, *Giant Dipper* uses granular processing, analysis and resynthesis, and Kyma's TAU algorithm. The movements of both of the Gametrak's cables are designed to control the various parameters of the sound, including amplitudes, frequencies, spatializations, reverbs, and playback speed in real-time by manipulating the combinations of X, Y, and Z coordinates.

One important aspect of this composition is its Graphic Score. In Western societies, the primary method for composers to preserve their musical ideas was to notate pitches and related musical information on staff paper. This notational practice was highly pitch-focused and evolved over many centuries. When composing for the data-driven instrument, however, the traditional notation system proved inadequate because it could not provide important information needed for the performance. To compensate for traditional notation's limitations, I created a musical score that provides performative instructions for 12 separate "scenes" showing not only how to perform the composition in time, but also how to play and control the unique data-driven instrument to control specific musical parameters in real-time. Figure 5 displays the graphic score for the "Dream" scene, illustrating these instructions.

In performing this scene, the left-side string controls the manipulation of the "Low Humm" sound while the right-side string controls the "Da La Singing" audio file manipulation. In Figure 4, the blue horizontal curved arrows denote that pulling the left-side or right-side nylon string towards the left or right will manipulate the panning of the "Chop" sound or "Da La Singing" sound. In contrast, pulling the Gametrak string on the left-side straight forward or backward generates a data stream to send to Kyma, which then changes the frequency of the "Chop" sound so it goes higher or lower in pitch. The red vertical arrows control volume: pulling either string up or down changes it. Releasing both strings triggers Kyma's WaitUntil Sound object[4] at the score's bottom, prompting the start of the next scene.

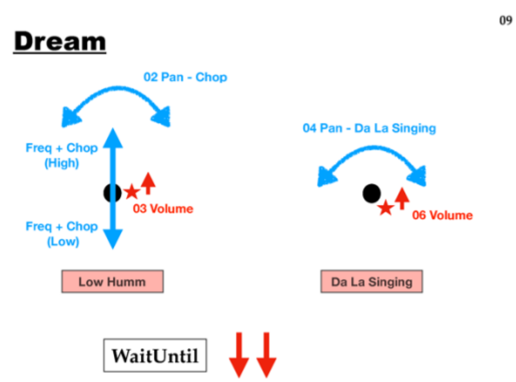

Figure 5. Graphic score for the "Dream" scene.

Performance instructions for each section of the composition were rendered as graphic (PNG) files using Apple's Pages software. Then, using the ImageDisplay Sound object, each section of the score appears at just the right moment in sequence in real-time as the composition progresses.

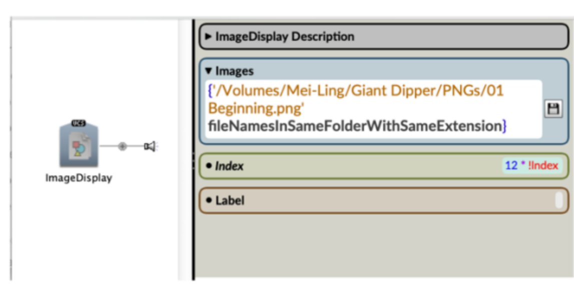

Figure 6. ImageDisplay Sound object in Kyma.

[4] The WaitUntil Sound object is Kyma's version of a fermata.

Composition: Farewell

Farewell is a composition inspired by a story about friendship and built upon audio recordings of my own voice reading the Chinese 賦得古原草送別 "Farewell on Grassland," written by Tang dynasty poet Bai Juyi (白居易).[5] This poem metaphorically relates the departure of a friend to the changing seasons of grasslands in China. I selected this poem due to the enduring cultural resonance of its second stanza, still commonly referenced in daily conversations across Chinese-speaking cultures despite being over 1200 years old. The poem's translation is provided below:

離離原上草，一歲一枯榮。
野火燒不盡，春風吹又生。
遠芳侵古道，晴翠接荒城。
又送王孫去，萋萋滿別情。

How luxuriantly the plains grass grows,
Wilting and rising again once every year.
Wildfires burn, but they are never exhausted.
Spring breezes blow, and up they spring again.
Ahead, wild growths overrun the ancient path,
And surround the old fort under cloudless skies.
Again, I'm sending the royal friend off,
My sorrow at parting is rich as the grass richly grows.

(Vacant Mountain 2023)

Designed for the Wacom tablet and Kyma, the complete structure of the data-driven instrument for *Farewell* is shown in Figure 7.

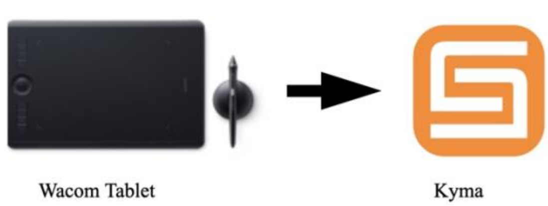

Figure 7. The structure of data-driven instrument for *Farewell.*

The Wacom tablet, initially intended for digital drawing, enables users to generate data by drawing with a stylus pen or fingers on its surface. With high precision, it offers 8192 levels of pen pressure, up to 60 levels of pen tilt recognition, and a spatial resolution of 5080 lines per inch (Wacom n.d.). This abundance of data allows for detailed sound control, which can be routed to Kyma for managing musical parameters.

Inspired by the above Chinese poem, I selected a controller that mirrors the act of Chinese calligraphy, traditionally used for writing poetry. The Wacom tablet emulates the brush-on-paper action of calligraphy: the tablet surface represents paper, and the stylus pen symbolizes the brush. This metaphorical linkage translates literary calligraphic actions into a sonic realization. *Farewell* combines poetry and calligraphy – two significant elements in Asian culture – by recording the poem and using the Wacom tablet as a symbol of calligraphy through data-driven instruments.

The Chinese poem's recitation contains distinctive vocal intonations, offering rich sonic material for manipulation. To amplify these unique intonation features, I extensively utilize Kyma's TAU algorithm, a powerful tool for manipulating the human voice and transforming speech into pitched sound material. In addition to the TAU algorithm, the synthesis methods employed in *Farewell* encompass granular processing, analysis and re-synthesis, as well as sampling techniques.

The beginning four semi-pitched notes at the start of Section A are created using Kyma's TAU algorithm. Each note is built using four TAU objects to manipulate pitch contour. Each TAU object is slightly detuned and have random values added to continuously change their pitches. The frequencies of each TAU sound are controlled by the degree of tilt of the stylus. When the pen is held straight, the pitch changes less. When the pen is tilted to the side at steeper or shallower

[5] Bai Juyi (772–846), was a well-known government officer and Chinese poet during the Tang dynasty. He wrote many poems about his experiences in government and his views of daily life while serving as the governor of three different provinces. In addition to his straightforward, clear, and approachable poetic style, Bai Juyi is also known for his social and political criticism.

angles, notes are detuned more or less. The result of nuanced detuning and random pitch shifting in micro scales creates a shimmery sound effect. Figure 8 shows the design of the first semi-pitched note in Kyma Sound as well as the Capytalk for each frequency parameter inside the TAU object.

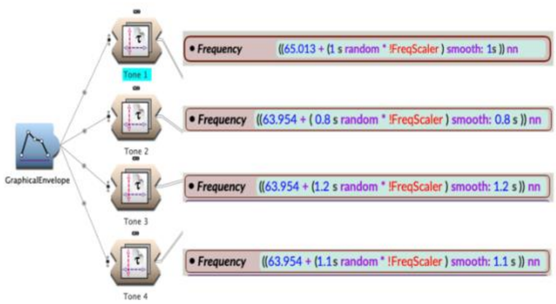

Figure 8. Four TAU objects and frequency parameter design.

Section A' begins at 6:52 and is based on different sound design techniques. Additionally, the four different notes are spatially positioned in a reverse order from their original position in Section A. Instead of using TAU to generate sound, Section A' uses the Kyma Sound object Chopper to extract a very small segment of my reading of the poem. The design of the Kyma Sound is shown in Figure 9.

Beyond pitch manipulation, the spatialization of the four semi-pitched notes in both sections A and A' are carefully designed. When the stylus touches the upper right corner of the Wacom surface, the sound is triggered and is heard coming from the front right speaker. As the stylus pen draws on the surface diagonally to the bottom left corner, the spatial position of the note moves diagonally to the rear left speaker. The three notes that follow are also designed to be spatially positioned in a similar fashion. Eventually, all four notes are spatially positioned in the four separate corners of the listening area.

Composition: Summoner

Summoner weaves a narrative inspired by an imagined story that centers around my mother-in-law who was a great lover of all kinds of animals. The primary sounds contained in this composition are actual recordings of peacock calls, owl hoots, and various other bird vocalizations, directly recorded from her backyard. The composition transforms these original recordings to embark on an exploration of the mysteries, mythologies, and mysticism surrounding these creatures, serving as a summoning not just to bring forth the birds, but also to articulate their compelling narratives.

The data-driven instrument for *Summoner* consists of a Leap Motion, custom software created in Max, and Kyma. Using the "aka.leapmotion" Max object (Masayuki. n.d.),[6] the data from Leap Motion is parsed and scaled into individual streams, then routed to Kyma.

The Leap Motion is an optical hand-tracking device equipped with two 640 x 240-pixel near-infrared cameras and three LEDs. It functions by tracking hand and finger movements within a designated three-dimensional zone, extending up to about 60 cm, with a wide field of view ranging between 120° and 140° (Ultraleap. 2020).

For optimal performance, a dimly lit concert hall with a minimal number of LED lights is preferred. Controlled darkness is necessary due to

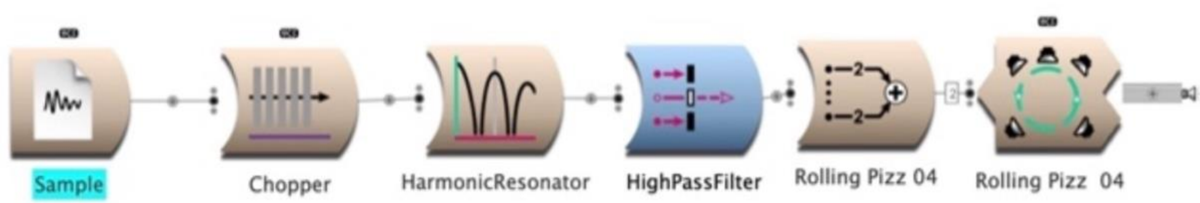

Figure 9. Kyma sound design for Section A'

[6] Masayuki Akamatsu created a number of external Max objects that are widely used. "aka.leapmotion" Max external object package was downloaded from https://github.com/akamatsu/aka.leapmotion.

the infrared sensors' sensitivity. Incandescent or halogen lights, as well as daylight, can disrupt the Leap Motion's accuracy in detecting hand and finger positions.

Figure 10. Leap Motion hand-tracking interaction zones (Reallusion 2024).

All performative movements in this composition are based on the concept of turning an intangible sound into an imaginary physical object. By breaching numerical thresholds, I can trigger or stop the sounds in the composition. I also use hand distance and speed to trigger musical events. The 3-D distance between the two hands as measured by the infrared sensors is calculated by Max and those values are used to trigger musical events. For example, at 3:00 in the video recording, when the distance of both hands reaches a threshold that I predetermine, an explosive sound is triggered. The metallic banging sound in Section B is triggered by how quickly my hand exits the Leap Motion's observable sphere and returns to it. The pitch of each banging sound is indeterminate and is selected through an algorithm resident in Kyma.

Changing my hand shape is a key performance technique within the composition. For instance, during the transition section at 3:04, I use all ten fingers – altering hand shape – to control and shift the frequencies of the peacock sounds. The speed of my fingers' vertical movements directly influences how fast the frequencies change and intensifies the buildup of sound.

Conclusion

Employing data-driven instruments, these compositions demonstrate the great potential of creating music in a non-traditional way. With the current state of technological development, data-driven instruments carry two unique features, modularity and mutability, which allow us to create unlimited instrumental variations with tremendous musical flexibility, as well as the boundless potential and ability to transform and contextualize sound material in much more radical and multidimensional ways than with many traditional instruments.

Creating compositions for data-driven instruments presents a complex challenge. It demands one to simultaneously represent the roles of inventor, composer, and performer. Balancing these distinct yet interrelated roles is a constant struggle throughout the composing process.

Through the endeavors detailed in this paper and through the execution of these compositions, I present the potential of the data-driven instruments. Despite the existing technological, intellectual, and conceptual demands, I believe this approach expands our capacity as composers and performers to communicate human experience and to powerfully reach and impact audiences. The utilization of data-driven instruments might open avenues to novel ways of sharing the human experience, allowing individuals, families, and communities to delve deeper into our collective experiences and relationships.

Acknowledgments

This paper is adapted from the author's dissertation, titled *Storytelling: The Human Experience Through Data-driven Instruments*, submitted to the University of Oregon in 2023. Significant modifications have been made to adapt the content for this publication (Lee 2023).

References

Akamatsu, Masayuki. n.d. "aka.leapmotion, akalogue." Accessed November 11, 2024. http://akamatsu.org/aka/max/objects/.

Bachmann, Daniel, Frank Weichert, and Gerhard Rinkenauer. 2014. "Evaluation of the Leap

Motion Controller as a New Contact-Free Pointing Device." *Sensors 15*, no. 1: 214–33. https://doi.org/10.3390/s150100214.

Collins, Nicolas. 2009. *Handmade Electronic Music: The Art of Hardware Hacking.* Abington: Routledge.

Cox, Christoph, and Daniel Waner. 2004. *Audio Culture: Reading in Modern Music.* New York: Bloomsbury Continuum International Publishing Group.

Dodge, Charles, and Thomas A. Jerse. 1997. Computer Music: Synthesis, Composition, and Performance. New York: Schirmer Books.

Freed, Adrian, Devin McCutchen, Andy Schmeder, et al. 2009. "Musical Applications and Design Techniques for the Gametrak Tethered Spatial Position Controller." In Proceedings of SCM 2009 - 6th Sound and Music Computing Conference, 189–94. Porto, Portugal. https://doi.org/10.5281/zenodo.849645.

Grout, Donald Jay, and Claude V. Palisca. 2001. *A History of Western Music*. New York: W. W. Norton.

Hinkle-Turner, Elizabeth. 2006. *Women Composers and Music Technology in the United States.* Burlington: Ashgate.

Holmes, Thom. 2015. *Electronic and Experimental Music: Technology, Music, and Culture*. 4th ed. Abington: Routledge.

Hunt, Andy, Marcelo M. Wanderley, and Matthew Paradis. 2002. "The Importance of Parameter Mapping in Electronic Instrument Design." In Proceedings of the *2002 Conference on New Instruments for Musical Expression* (NIME-02). Dublin, Ireland.

Lee, Mei-ling. 2023. "Storytelling: The Human Experience through Data-Driven Instruments." DMA diss., University of Oregon.

Matos, Michelangelo. 2016. *The Underground Is Massive: How Electronic Dance Music Conquered America*. New York: Dey Street Books.

Reallusion. n.d. "Leap Motion: The Essential Hand Tracking Device." Accessed December 16, 2024. https://mocap.reallusion.com/iclone-motion-live-mocap/leap-motion.html.

Roads, Curtis. 1996. *The Computer Music Tutorial*. Cambridge, MA: MIT Press.

Scaletti, Carla. 2004. *Kyma X Revealed!: Secrets of the Kyma Sound Design Language*. Champaign, IL: Symbolic Sound.

Scaletti, Carla. 2016. "Looking Back, Looking Forward: A Keynote Address for the 2015 International Computer Music Conference." Computer Music Journal 40, no. 1: 10–24.

Stolet, Jeffrey. 2012. *Kyma and the SumOfSines Disco Club*. Morrisville: Lulu Press.

Stolet, Jeffrey. 2013. "Twenty-three and a Half Things about Musical Interfaces." Keynote address, Kyma International Sound Symposium, Brussels, Belgium, September 14.

Stolet, Jeffrey. 2021. *Do: Notes about Action in the Creation of Musical Performance with Data-Driven Instruments*. Lulu Press.

Ultraleap. 2020. "How Hand Tracking Works." Posted September 3. https://www.ultraleap.com/company/news/blog/how-hand-tracking-works/.

Ultraleap. n.d. "Leap Motion API Overview." Ultraleap Developer. Accessed October 25, 2023. https://developer-archive.leapmotion.com/documentation/csharp/devguide/Leap_Overview.html.

Vacant Mountain. n.d. "Translation Thursday: 赋得古原草送别 Bidding Farewell on the Plain, by 白居易 Bai Juyi." Accessed October 28, 2023. http://www.vacantmountain.wordpress.com/2013/05/30/translation-thursday-!%22#%&'-bidding-farewell-on-the-plain-by-()*-bai-juyi).

Wacom. n.d. “Wacom Intuos Pro.” Accessed December 3, 2022. https://www.wacom.com/en-cn/products/ pen-tablets/wacom-intuos-pro.

Winslow Burleson Hansen, Camilla Nørgaard Jensen, and Alex Mesker. 2009. “Musical Applications and Design Techniques for the Gametrak Tethered Spatial Position Controller.” In Proceedings of *the 6th Sound and Music Computing Conference*, 189–94. Porto, Portugal: CMS 2009.

Media References

Video recordings for the discussed compositions can be accessed below:

Giant Dipper: https://youtu.be/jz9tfMLZ40w

Farewell: https://youtu.be/nyFtCIpFcNY

Summoner: https://youtu.be/Og7dSErldOs

George Edmondson

University of Birmingham
gxe370@student.bham.ac.uk

Introduction

What does it mean to hear a place, to immerse in the sounds that illustrate, curate, and elicit the identities and histories of the people it hosts? The interplay between sound art and ethnography offers a unique opportunity to explore this question, reshaping how we document and engage with the world around us. This article coincides with fieldwork conducted in collaboration with community members in Northfield, Birmingham.

By centering listening as both method and medium, this work invites us to reconsider the role of sound—not only as an artifact of culture but as an active agent in shaping how communities perceive themselves, their surroundings, their memories and perceptions. For sound artists, ethnographers, and scholars alike, it offers insight on using sound to bridge the gaps between past and present, individual and collective, local and global.

Northfield's Industrial and Cultural Evolution

Northfield's trajectory from a thriving industrial center to a post-industrial community mirrors the complex transformations of Birmingham itself—an interplay of economic shifts, urban reinvention, and evolving collective identities. Once emblematic of Birmingham's reputation as the "workshop of the world," Northfield stood at the heart of a city defined by its manufacturing prowess. As Kew reminds us, Birmingham was celebrated for producing "finely crafted goods such as buttons, glass and guns," industries that provided employment for thousands and fostered a strong sense of identity rooted in skilled labor and artisanal pride (Kew 2023). Yet, Northfield's industrial significance reaches even further back. As Vinen notes, it was the largest settlement in the Birmingham area listed in the Domesday Book, valued several times higher than Birmingham itself during the Anglo-Saxon period (Vinen 2022)—a testament to its longstanding importance.

The late 20th century, however, brought waves of economic decline that irrevocably altered Birmingham's fortunes. Between 1971 and 1986 (Flynn and Taylor 1986), the city lost nearly 191,000 jobs, a substantial % of its total employment, with the manufacturing sector hit hardest (Vinen 2022). Once aligned with the "affluent South," Birmingham found itself reclassified among the "declining North," a shift symptomatic of a broader "urban crisis" marked by fiscal stress, unemployment, and social unrest. Northfield bore the brunt of these changes, epitomized by the closure of the Longbridge car plant—a symbolic and material loss for the community. At its height in the 1960s, Longbridge employed 25,000 workers, yet the plant's gradual decline culminated in its final closure in 2005 (Vinen 2022). Vinen observes that while some workers transitioned to new employment, many faced reduced wages and precarious conditions, eroding the economic stability that had once defined the area (Vinen 2022).

In response, Birmingham embarked on an ambitious program of urban redevelopment. Flagship projects like the International Convention Centre (ICC) and Symphony Hall sought to reimagine the city as a hub for global business and culture (Chinn 2016). Yet these efforts often deepened existing divisions. Critics argue that the city's redevelopment agenda adhered to a neoliberal framework, prioritizing prestige and profit at the expense of inclusivity. Resources were diverted from essential services such as education and housing, sidelining the economically deprived and ethnically diverse communities that form Birmingham's backbone (Adams 2011; Chan 2007; Gale 2004). Kew highlights the paradox of these policies, noting how Birmingham's motto, "Forward!" often translated into a vision of progress that marginalized its most vulnerable populations (Kew 2023).

A precursor to the tension— the physical transformation of Birmingham's urban landscape in the 1960s. The city's modernist planners, led by Herbert Manzoni, envisioned a futuristic cityscape, replacing Victorian housing with urban motorways and concrete underpasses (Adams 2011). Such changes are considered to have partially erased the city's industrial heritage, leaving behind what Kew terms a "concrete palimpsest." (Kew 2023). For Northfield, the nostalgia associated with sites like Longbridge was replaced by an uncertain future tied to retail parks and technology corridors, further alienating residents.

Northfield, like the rest of Birmingham, also grappled with the profound demographic and cultural shifts brought about by post-war migration. Immigration from the Caribbean, South Asia, and other Commonwealth regions reshaped the city's identity, further positioning Birmingham as a "community of communities" (Myers and Grosvenor 2011). These migrations were underpinned by policies such as the British Nationality Act of 1948, which expanded citizenship to Commonwealth migrants and played a central role in Birmingham's establishment of the "multicultural midlands" (Kew 2023). Yet, as Kew observes, urban planners often viewed these migrants as a disruption to visions of "low population density" and "de-congested suburban spaces," leading to segregation and the concentration of minority communities in specific areas (Kew 2023).

This history of migration cannot be disentangled from Birmingham's imperial legacy. As Henry et al. argue, the city was deeply embedded in British colonial networks, producing "guns and other apparatus of colonialism" while benefiting from the labor and resources of colonized peoples (Henry, McEwan, Pollard 2002). The post-colonial era brought these imperial connections into Birmingham's urban fabric, with migrant communities contributing to the city's cultural and economic resilience. Industries like the British Bhangra music scene (Henry, McEwan, Pollard 2002) and the establishment of the Chinese Quarter (Wun Fung Chan 2007) are emblematic of how these communities shaped Birmingham's identity. Yet these contributions have often been overshadowed by narratives of economic decline and cultural tension.

The rhetoric of Brexit brought these tensions into sharper relief. In areas like Northfield, where the collapse of industrial hubs like Longbridge had fostered resentment towards globalization and centralized policies, the "politics of nostalgia" gained traction. Vinen identifies this sentiment as a longing for a past defined by industrial stability and local pride, a past seemingly undermined by the forces of globalization and the European Union. In 2016, all wards in Northfield voted to leave the EU, reflecting these deep-seated anxieties (Vinen 2022).

Yet, Northfield's history is not solely one of loss and division. It is also a story of resilience and reinvention. Initiatives like the Birmingham People's History Archive (Binnie 2023) have worked to preserve the voices and experiences of the city's working-class communities, offering a counter-narrative to the "condescension of posterity" (Thompson 1980). By capturing oral testimonies and archival materials, projects like the BPHA ensure that the histories of communities like Northfield are not forgotten but celebrated. Northfield's journey from industrial heartland to post-industrial community encapsulates the broader dynamics of Birmingham's transformation. Its story is one of economic upheaval, urban reinvention, and cultural flux, but also of resilience and hope.

Engaged Sound Art: Methodologies and Critical Intersections

Imagine a quiet gallery, its walls lined not with paintings but with speakers, each emitting fragments of voices, environmental sounds, and melodies that seem to interact with the very architecture. The visitors move through the space, not merely observing but responding— sometimes actively contributing to the composition through digital interfaces or their own movements. This is the landscape of contemporary engaged sound art: a space where technology, community, and creativity converge to redefine the boundaries of artistic practice and social engagement.

The field is as varied as it is vibrant, and the methodologies it encompasses are anything but uniform. This section overviews recent approaches to engaged sound art, from large-scale digital collaborations to intimate, site-specific installations, to explore how sound art today grapples with pressing questions: How can it remain inclusive without sacrificing complexity? What roles do artists and participants play in the co-creation of meaning? And how do these works navigate the line between documentation and abstraction, between the ephemeral and the enduring?

Tsuruoka, Ellis, and Chang's Ear Talk (2022) exemplifies how digital platforms can revolutionize collaborative music-making. Using YouTube live-streaming and web-based systems, the project invites participants from around the world to contribute real-time sound inputs. These inputs—ranging from vocalizations to instrumentals—are layered into a constantly evolving composition.

What sets Ear Talk apart is its ability to create a sense of shared authorship across vast distances. The platform itself becomes a collaborator, its algorithms influencing the prominence and layering of contributions. Yet, this reliance on technology is also where challenges emerge. Who gets to participate meaningfully, and whose contributions risk being obscured by the platform's mechanics? The algorithms, while designed to facilitate inclusivity, inadvertently impose a form of editorial control that participants cannot easily contest. Ear Talk raises vital questions about the agency of both artists and audiences in the digital age, making its strengths and constraints equally integral to its impact.

Vadim Keylin's (2022) exploration of participatory sound art includes Katrine Faber's Let Us Sing Your Place, a project rooted in site-specific performance. Faber works with communities to craft compositions that draw from local histories and landscapes, embedding collective memory into the act of singing. These performances transform the relationship between participants and their environments, offering a deeply personal engagement with place. Similarly, Benoît Maubrey's Speaker Sculptures invite interaction in a physical space. These installations—constructed from repurposed loudspeakers—emit sounds triggered by participants' movements or recorded inputs. The sculptures become alive with layers of sound: voices, ambient noises, and musical fragments coalesce to create a communal auditory experience.

Both works challenge the notion of the artist as sole creator, foregrounding participant contributions. However, this decentralization comes with its own complexities. In Faber's case, the act of collective singing risks flattening individual experiences into a homogenized narrative of place. For Maubrey, the spontaneity of audience interaction can lead to moments of cacophony that obscure the coherence of the work. These tensions, far from diminishing the projects, highlight the delicate balance required to merge individual agency with a cohesive artistic vision.

Theorizations like Barrett's concept of music "after sound," as employed by Moorehouse, Matthews, and Maia (2022), push the boundaries of what constitutes music. This framework situates sound as a secondary consideration, placing greater emphasis on the social and cultural contexts of music-making. For Barrett, music becomes a vehicle for sociopolitical engagement, a tool to interrogate power dynamics and cultural narratives.

What makes this approach so provocative is its departure from sound as a purely aesthetic phenomenon. Instead, it treats sound as a lens through which broader social issues can be explored. Yet this abstraction also risks alienating audiences who find resonance in sound's visceral qualities. By focusing so intently on context, the framework occasionally leaves unanswered questions about the sensory and emotional dimensions that are central to musical experience. Nonetheless, Barrett's work invites us to rethink music as an active participant in shaping cultural and political realities.

Arian Bagheri Pour Fallah's reflections on récit music (Bagheri Pour Fallah 2022) embraces fragmentation and modularity to reimagine musical collaboration in the networked age. Rooted in the concepts of reciprocity and dispersion, récit music invites participants to contribute, transform, and exchange sonic elements, creating compositions that evolve through collective iteration rather than static finality. This approach reflects Manuel Castells's notion of networks as boundaryless structures (Castells 2008) and Marcel Mauss's view of reciprocity as central to societal cohesion (Mauss 2002).

The framework critiques the insularity of fixed media and traditional concert-hall practices, which Bagheri Pour Fallah argues have fostered divisions between progressive and regressive musical poles (Bagheri Pour Fallah 2022). Instead, récit music encourages dynamic exchanges across temporal and spatial divides. For instance, a string quartet might gift an improvisation to an electroacoustic composer, who transforms and passes it on, continuing the reciprocal cycle until no further changes are needed (Bagheri Pour Fallah 2022).

While this decentralization can risk fragmentation, Bagheri Pour Fallah sees the resulting tension as integral to the form's ethos. Unlike spectromorphology, which prioritizes the composer and listener, récit music includes performers as active contributors, fostering a more collaborative musical community (Bagheri Pour Fallah 2022). However, the reliance on digital networks introduces challenges, potentially reducing reciprocity to commodified exchanges.

Developed by Simurra, Messina, Aliel, and Keller (Simura et al. 2022), Radical Creative Semantic Anchoring (Radical ASC) reimagines sound art as a democratic process. By integrating linguistic and ecological strategies, the framework prioritizes accessibility, enabling non-specialists to engage with music-making. Radical ASC emphasizes process over product, encouraging participants to focus on the act of creation rather than its outcome.

This inclusivity disrupts traditional hierarchies of expertise, opening sound art to a broader audience. However, the framework's radical departure from established practices raises questions about its integration with existing musical traditions. Can such an approach maintain its accessibility without becoming detached from the historical and cultural weight of sound as a medium? Radical ASC's emphasis on inclusivity offers a compelling vision for the future of sound art, even as it grapples with the complexities of its ambitions.

Simon Fox's New Amateurs and Tricksters (Fox 2022) champions the Trickster as a transformative figure in music-making, embodying disruption, play, and the democratization of creativity. The manifesto introduces the New Amateur—creators without formal training—who expand collaboration beyond traditional boundaries to include materials, ideas, and non-human entities. This approach echoes Maria Lind's extended models of collaboration and critiques conventional notions of expertise.

The Trickster plays a vital role, equipping New Amateurs with the courage to challenge professional norms and hierarchies. Fox describes the Trickster as a culture hero and provocateur, adept at navigating contested spaces and unmasking exclusivity. This archetype enables amateurs to leverage their lack of formal training as a strength, fostering experimental collaborations unburdened by tradition (Hyde 2017; Fox 2022).

Critiquing traditional music education for its rigidity, Fox contrasts it with the exploratory frameworks of visual art, which prioritize individual creativity. His vision aligns with Adam Harper's "n-dimensional modernism," advocating for music-making without restrictions or predefined assumptions (Harper 2011). By embedding personal creativity within collective, socially engaged practices, the New Amateur disrupts entrenched norms.

Finally, Erik DeLuca and Elana Hausknecht's Ears to the Ground (2023) integrates sound art with dialogic pedagogy. Drawing on Pauline Oliveros' practice of deep listening and Paulo Freire's critical pedagogy, the project uses sound to explore contested spaces and challenge dominant narratives. Through collaborative workshops and site-specific installations, Ears to the Ground transforms listening into an act of resistance.

By engaging participants in both creation and critique, the project offers a model for how sound can amplify marginalized voices and uncover silenced histories. Yet its structured frameworks, while effective in fostering dialogue, raise questions about how open-ended such participatory practices can truly be. Ears to the

Ground exemplifies the potential of sound art to bridge the personal and the political, creating spaces for reflection and action.

These projects demonstrate the expansive potential of contemporary sound art, where methodologies range from the deeply personal to the globally networked. Each approach grapples with tensions—between structure and freedom, inclusivity and expertise, abstraction, and embodiment. Rather than weaknesses, these tensions are sites of productive inquiry, driving the field forward. These works invite us to consider sound art not as a fixed discipline but as a dynamic, evolving practice. They challenge us to listen differently—to the sounds themselves, to the spaces they inhabit, and to the communities they engage. In doing so, they remind us that the true power of sound lies not in providing answers but in framing new ways of asking questions.

The Convergence of Sound Art and Ethnography

The relationship between sound art and ethnography has emerged as a fertile ground for exploring cultural, social, and spatial dynamics. This convergence reflects a shared commitment to engaging with lived experience, memory, and place, often through participatory and sensory methodologies. By integrating artistic and ethnographic practices, sound art expands the scope of ethnography beyond textual or visual documentation, creating new possibilities for understanding and representing cultural realities.

Sound art serves as both a creative and analytical tool within ethnography, enabling the documentation of cultural practices while fostering collective expression. This dual role is evident in Amelides' (2016) work on acousmatic storytelling, which integrates interviews, archival recordings, soundscapes, and music to craft hybrid narratives. This method not only preserves cultural memories but also invites listeners to co-create the final narrative by interpreting and reimagining the sonic material. As Amelides notes, "The listener participates in the creation of the 'final version' of the story, thereby becoming, in a sense, its co-creator" (Amelides, 2016, 220).

Boersen (2022) builds on this idea by highlighting the interplay between soundscapes and personal imagination. His research emphasizes the role of enactive listening, where ambiguous sounds allow participants to construct individual narratives. This approach underscores the potential of sound to act as a dynamic medium for cultural interpretation, bridging artistic abstraction with ethnographic depth.

The integration of sound art with sensory ethnographic methods further exemplifies this convergence. Walking, for example, is a methodological tool that enables researchers to engage with the spatial and sensory aspects of everyday life. Lawhon and Pierce (2015) argue that walking facilitates an embodied understanding of place, allowing ethnographers to participate in the spatiality of lived experiences. Similarly, Aduonum (2021) highlights walking's potential for serendipitous discovery, describing it as a means of uncovering hidden narratives and insights within a locale.

Sound art enhances this process by embedding auditory elements into the act of walking, transforming it into a multisensory exploration. The compositions studied by Boersen (2022) illustrate how soundscapes can deepen spatial engagement, encouraging listeners to interact with their surroundings through auditory cues. This synthesis of walking and listening creates a richer ethnographic experience, one that captures the fluidity and complexity of cultural landscapes.

The convergence of sound art and ethnography also brings critical attention to memory and representation. De Leon and Cohen (2005) employ object and walking probes to evoke deep memories through physical artifacts and locations. These methods foster trust and equality between researchers and participants, aligning with the ethical imperatives of ethnography. By giving participants control over the narrative, sound art and ethnographic practices converge to create inclusive and participatory frameworks.

Claire Bishop (2012) highlights the ethical complexities of participatory art, particularly in its drive to democratize creation by transforming audiences into co-creators. While these practices aim to counter societal alienation and forge collective authorship, they often navigate power imbalances, with participants' labor sometimes commodified as unpaid contribution. Bishop emphasizes that the artist-participant relationship is a dynamic interplay of dependency and tension, rather than inherently egalitarian. Participatory art, she argues, is neither a privileged political medium nor an automatic solution to societal challenges, but a precarious practice requiring ethical scrutiny and contextual negotiation to balance empowerment and exploitation.

But sound art can enrich ethnography by extending its sensory and temporal dimensions. The new mobilities paradigm proposed by Sheller and Urry (2016) situates movement as a central element of ethnographic practice, highlighting how "activities occur while on the move" and how mobility itself becomes a site for "occasioned activities" (Sheller and Urry 2016, 213). Within this framework, sound art provides a unique way to document and interpret these interactions and activities in motion, emphasizing the embodied and relational aspects of ethnographic research.

The role of silence, as explored by DeVito (2021), expands the sensory and emotional dimensions of ethnography, offering a space for reflection, imagination, and connection. Silence allows what DeVito describes as the "anticipation and imagination to grow," fostering both intrapersonal and shared dialogues that amplify the emotional resonance of ethnographic narratives. The strategic use of silence, paired with sound memories or subtle triggers, creates what DeVito terms an "emotional Easter Egg," a moment that draws participants into a deeper connection with the work. This aligns with Voegelin's (2006) assertion that sound's "quasi-virtual, immaterial blindness" renders it uniquely immersive, while the application of silence heightens empathy and engagement. Additionally, Norman (1996) underscores that the reflective and referential listening facilitated by such techniques enables participants to deconstruct and reinterpret sound in ways that connect past and present experiences. Together, these elements form a dynamic auditory space where cultural meanings can be negotiated, reimagined, and shared.

Could the convergence of sound art and ethnography represent a shift toward more collaborative, inclusive, and sensory methodologies? By integrating artistic and ethnographic practices, this approach challenges traditional boundaries between researcher and subject, artist and audience. Ultimately, the interplay between sound art and ethnography offers a framework for capturing the fluidity of cultural experiences while addressing ethical and representational challenges. This convergence not only broadens the methodological toolkit of ethnographers but also redefines the role of sound in documenting and interpreting the complexities of human life.

Listening as Ethnography: Soundscapes and Collaborative Narratives

Soundscapes, voice, and heard action form a vital and ever-changing part of daily life. Whether shaped by human activity or other environmental elements, these soundscapes are not neutral; their meaning shifts depending on who is listening and how. Sound carries layers of memory, identity, and reflection, offering more than mere auditory data—it represents how individuals and communities navigate and interpret their surroundings. This research takes as its premise that sound is not just heard but lived, experienced, and woven into broader narratives of socio-economic change, collective memory, and identity.

The project seeks to merge ethnography and sound art in ways that challenge conventional methodologies. Traditional ethnographic approaches—often centered on text and image—can sometimes fall short in conveying the relational immediacy of human experience. Sound, by contrast, has the potential to offer an immersive and participatory experience, opening pathways to deeper and more affective connections. Drawing on Steven Feld's (1994, 996, 2003, 2015) concept of "sound as a way of knowing," the research positions sound not as a supplement but as a core medium of inquiry and expression. Interlocutors are invited to co-create and disseminate their sonic narratives, revealing how sound shapes and is shaped by their interpretations of place and memory.

This approach is not merely about documenting what is heard; it is about listening through the perspectives of others. Sound here becomes a bridge—linking researcher and participant, personal and collective, present and remembered. It transforms listening into an act of collaboration, aided by sound art in uncovering the hidden or obscured layers of meaning within auditory experiences. By recreating the contexts in which sounds are heard, the research aims to explore the multiplicity of meanings embedded in these soundscapes.

Sound serves as both a method of documentation and a tool for dissemination. Techniques such as sound-walking, collaborative sonic journals, and object probes enable participants to articulate their experiences and connect them to personal or shared memories. These approaches center dialogue and reciprocity, shifting the role of the researcher from an external observer or co-participant to an active facilitator in the co-construction of knowledge. Through experimental sound compositions, these collected materials are transformed into layered, sonic collages, blending ethnographic documentation with co-creativity.

Research Questions

The study is guided by questions that aim to reimagine the role of sound in ethnographic practice:

1. How might sonic field notes, multimodal interaction, and co-authored making provide a model for a more reciprocal, intimate ethnographic approach?
2. Could merging soundscape's "preservation of place" with the acousmatic enhancement of the "sonic reveal" productively further the dissemination of experiences, particularly through narrative engagement?
3. Could a greater reflexive mode of operation—developed through ongoing dialogue around the co-creation and manipulation of sound—challenge canonical compositional assumptions and suggest a more democratic, fluid sound-art form?
4. To what extent could an engaged, cooperative sound practice establish a more reciprocal methodological tool to educate and facilitate sound-based experiences for community-based creative groups?

Methods and Practices

To explore these questions, the project employs sound-based methods that foreground sensory and affective engagement, emphasizing collaboration between researcher and participant:

1. Interviews: These are approached as "inter-views" into the perspectives and experiences of interlocutors. Beyond words, the qualities of voice—intonation, rhythm, and emotion—offer insight into how participants "know" and interpret their auditory worlds.
2. Sound-walking: Participants are guided through their environments, attending closely to the auditory textures of their surroundings. This method encourages reflection on how sound shapes their understanding of space and facilitates shared exploration.
3. Collaborative sonic journals: Participants document the sounds of their daily lives, creating personal archives that trace their auditory landscapes. These journals reveal how individuals navigate and respond to their environments.
4. Object probes: This technique links specific sounds to personal objects or memories, exploring how auditory stimuli evoke the past and anchor it within the present.
5. Sound compositions: Sonic materials collected through these methods are transformed into experimental compositions. These works are both documentation and artistic reimagination, offering new ways to experience the lived realities of Northfield.

By incorporating sound as both method and medium, the project adds to the ways ethnography can be approached as a collaborative process. It emphasizes listening as an active, relational practice that not only documents lived experience but also amplifies voices and redistributes agency. Through the merging of sound art and ethnography, the research fosters a shared exploration of auditory and cultural landscapes, inviting participants and audiences alike to reimagine the relationship between memory, identity, and place.

Micro Vignette: Sound as Ethnographic Method, Documentation, and Art

A conversation begins: one participant recalls childhood memories, their voice rich with emotional resonance tied to specific places. Another reflects on the day their workplace shut its gates, their subdued tone evoking shared loss and transition. These moments—where language and sound converge—highlight the indispensable role of interviews. Beyond the words spoken, tone, cadence, and emotion reveal a deeper layer of meaning, positioning interviews as both a method and metaphor: an inter-view that engages sound to access lived experience.

While the research focuses on environmental soundscapes—what the community hears and how they interpret those sounds—the interview offers a complementary perspective: what the community says and, critically, how it is said. Speech, with its rhythm, inflections, and pauses, becomes a sonic narrative, offering insights that written text cannot fully capture. These auditory nuances convey not just content but also cultural identity and the emotional texture of memory.

Walking through the research site unveils a soundscape alive with layers of significance. Participants pause to reflect on the sounds around them, drawing connections between auditory environments and personal histories. Inspired by scholars like Lawhon and Pierce (2015), Aduonum (2021), Impey (2019), and Jack (2021), sound-walking frames sound as both a reflection of life and an active force shaping it. Subtle rhythms, from birdcalls to distant machinery, act as markers of continuity and change, inviting reflection on the meaning embedded in everyday auditory experiences.

In one session, a participant links a treasured object to a specific sound, illustrating how auditory memory bridges personal history and broader cultural narratives. These connections are deepened by the way sound art, as ethnography, captures and highlights moments of sonic recontextualization. This does not mean altering participants' perceptions but instead acknowledging how sounds can evoke new or layered memories. For instance, hearing a park might evoke childhood memories, but when the park's sounds are accompanied by the distant wail of a siren, the memory might shift, becoming imbued with tension or nostalgia, depending on the listener's associations. These moments reveal how memory and emotion intertwine dynamically with sound, creating a richer understanding of place and identity.

Such interplay is particularly evident in the collaborative nature of this research. Participants actively contribute their interpretations, shaping the representation of their experiences. This shared authorship maintains their agency and ensures authenticity in how their narratives are presented, highlighting the importance of reciprocity in ethnographic practice.

The site's stories emerge not only through its visible landscapes but through its vibrant soundscapes. Everyday auditory elements carry meaning that evolves with memory and context. By attending to these nuances, the research transcends simple documentation, fostering reflection on how sound constructs belonging and identity. Through methods like object probes, sound-walking, and careful engagement with sonic recontextualization, the research redefines "place" as a dynamic narrative—a living archive of shared memory and cultural resonance. In the voices of its people and the rhythms of its spaces, the community comes alive, underscoring the profound connections that sound forges across time and place.

Reflections on Outcomes and Future Directions

Sonic-arts-based ethnography presents a compelling opportunity to document and share community narratives, potentially opening new pathways for exploring identity through auditory experiences. By combining electroacoustic methods with participatory practices, this approach offers a platform for collaborators to engage with and reinterpret their surroundings. In Northfield, as in other post-industrial settings, these methods have the capacity to reframe communal histories, transforming them into co-created soundscapes that reflect the layered experiences of their contributors.

This "capacity" opens the door to further exploration. How might future projects enhance inclusivity by integrating underrepresented voices and diverse sonic traditions? Expanding archives to reflect a wider range of auditory experiences—including contributions from non-Western contexts—could deepen our understanding of sound's role in cultural memory. Likewise, collaborations between sound artists, ethnographers, and community historians offer the potential to refine these methodologies, striking a balance between aesthetic ambitions and ethical considerations.

The potential for electroacoustic sound art to reshape perceptions of space and identity is significant, but so are the challenges. Broadening the reach of these projects while maintaining their intimate, community-centered focus will require careful and thoughtful approaches.

The convergence of sound art and ethnography reshapes how we document, interpret, and engage with the world around us. This research underscores the importance of exploring music and sound beyond aesthetic frameworks, highlighting their capacity to foster dialogue and community connection.

Yet, the questions remain open-ended. What does it mean to listen collaboratively? How might sound bridge the gap between memory and identity, the local and the universal? This research invites continued exploration into these intersections, urging sound artists, ethnographers, and scholars to embrace co-creation as both a method and a goal.

Through such inquiry, sound-based ethnography not only reimagines the boundaries of artistic practice but also asserts the vital role of listening as a mode of cultural engagement. In amplifying the voices of the past and present, it encourages us to rethink our relationships—with place, with history, and with each other. Let the dialogue continue.

References

Adams, David. "Everyday Experiences of the Modern City: Remembering the Post-War Reconstruction of Birmingham." *Planning Perspectives* 26, no. 2 (April 2011): 237–60. https://doi.org/10.1080/02665433.2011.550446.

Aduonum, Ama Oforiwaa. "Walking as Fieldwork Method in Ethnomusicology." *Ethnomusicology* 65, no. 2 (July 1, 2021): 221–58. https://doi.org/10.5406/ethnomusicology.65.2.0221.

Amelides, Panos. "Acousmatic Storytelling." *Organised Sound* 21, no. 3 (November 11, 2016): 213–21. https://doi.org/10.1017/s1355771816000182.

Bagheri Pour Fallah, Arian. "Récit Music: Musical Community in the Information Age." *Organised Sound* 28, no. 1 (February 7, 2022): 78–87. https://doi.org/10.1017/s1355771822000048.

Binnie, Gill. "It's Your History: Birmingham People's History Archive (BPHA)." *Midland History* 48, no. 1 (January 2, 2023): 109–11. https://doi.org/10.1080/0047729x.2023.2182973. "Birmingham: Introduction." Essay, n.d.

Bishop, Claire. *Artificial Hells: Participatory Art and the Politics of Spectatorship*. London: Verso, 2012.

Boersen, Ronald. "Enactive Listening: Perceptual Reflections on Soundscape Composition." *Organised Sound* 27, no. 1 (April 2022): 69–79. https://doi.org/10.1017/s135577182200019x.

Castells, Manuel. "The New Public Sphere: Global Civil Society, Communication Networks, and Global Governance." *The ANNALS of the American Academy of Political and Social Science* 616, no. 1 (March 2008): 78–93. https://doi.org/10.1177/0002716207311877.

Chan, Wun Fung. "Writing Multiculturalism? Planning for Culturally Different Identities in the City of Birmingham." *Planning Theory & Practice* 8, no. 1 (March 2007): 69–85. https://doi.org/10.1080/14649350601158154.

Chinn, Carl, and Malcolm Dick. *Birmingham: The Workshop of the World*. Liverpool: Liverpool University Press, 2016.

De Leon, Jason Patrick, and Jeffrey H. Cohen. "Object and Walking Probes in Ethnographic Interviewing." *Field Methods* 17, no. 2 (May 2005): 200–204. https://doi.org/10.1177/1525822x05274733.

Deluca, Erik, and Elana Hausknecht. "Ears to the Ground: Socially Engaged Sound Art as Learning in Process." *Organised Sound* 28, no. 1 (February 7, 2022): 3–12. https://doi.org/10.1017/s1355771822000036.

DeVito, Joseph A. *The Interpersonal Communication Book.* 16th edition. Boston, MA: Pearson, 2021.

Fox, Simon. "New Amateurs and Tricksters: A Manifesto for Music and Sound Creation." *Organised Sound* 28, no. 1 (July 19, 2022): 43–52. https://doi.org/10.1017/s1355771822000310.

Feld, Steven. "From Ethnomusicology to Echo-Muse-Ecology: Reading R. Murray Schafer in the Papua New Guinea Rainforest." Soundscape Newsletter 8 (1994): 9–13.

Feld, Steven. "Waterfalls of Song: An Acoustemology of Place Resounding in Bosavi, Papua New Guinea." In *Senses of Place*, edited by Steven Feld and Keith H. Basso, 91–135. Santa Fe, NM: School of American Research Press, 1996.

Feld, Steven. "A Rainforest Acoustemology." In The Auditory Culture Reader, edited by Michael Bull and Les Back, 223–39. Oxford: Berg, 2003.

Feld, Steven. "Acoustemology." In *Keywords in Sound*, edited by David Novak and Matt Sakakeeny, 12–21. Durham, NC: Duke University Press, 2015.

Flynn, N., & Taylor, A. P. Inside the Rust Belt: An Analysis of the Decline of the West Midlands Economy. 1: International and National Economic Conditions. *Environment and Planning A: Economy and Space* 18, no. 7 (1986), 865-900. https://doi.org/10.1068/a180865.

Gale, Richard. "The Multicultural City and the Politics of Religious Architecture: Urban Planning, Mosques and Meaning-Making in Birmingham, UK." *Built Environment* 30, no. 1 (March 1, 2004): 30–44. https://doi.org/10.2148/benv.30.1.30.54320.

Harper, Adam. *Infinite Music: Imagining the Next Millennium of Human Music-making*. Ropley: Zero, 2011.

Henry, N, C McEwan, and J Pollard. "Globalization from below: Birmingham – Postcolonial Workshop of the World?" *Area* 34, no. 2 (June 2002): 117–27. https://doi.org/10.1111/1475-4762.00064.

Hyde, Lewis. *Trickster Makes this world: How Disruptive Imagination Creates Culture.* Canongate Books, 2017.

Impey, Angela. *Song Walking: Women, Music, and Environmental justice in an African borderland*. Chicago: The University of Chicago Press, 2019.

Jack, Max. "The Crowd in Flux: Atmosphere and the Governance of Public Affects at FC Union Berlin." *Ethnomusicology* 65, no. 3 (October 1, 2021): 497–518. https://doi.org/10.5406/ethnomusicology.65.3.0497.

Kew, Tom. *The Multicultural Midlands*. Manchester: Manchester University Press, 2023.

Keylin, Vadim. "Creative Agencies in Participatory Sound Art: Two Case Studies." *Organised Sound* 28, no. 1 (February 8, 2022): 13–24. https://doi.org/10.1017/s1355771822000085.

Mauss, Marcel, Mary Douglas, and W. D. Halls. *The gift: The form and reason for exchange in archaic societies*. London: Routledge, 2002.

Moorehouse, Aaron, Harry Matthews, and Oogoo Maia. "Post-Sonic Perspectives on Socially Engaged Compositional Practices: Composing 'after Sound' and beyond Music." *Organised*

Sound 28, no. 1 (February 14, 2022): 88–96. https://doi.org/10.1017/s135577182200005x.

Myers, Kevin, and Ian Grosvenor. "Birmingham Stories: Local Histories of Migration and Settlement and the Practice of History." *Midland History* 36, no. 2 (September 2011): 149–62. https://doi.org/10.1179/004772911x13074595848915.

Norman, Katharine. "Real-World Music as Composed Listening." *Contemporary Music Review* 15, no. 1 (January 1, 1996): 1–27. https://doi.org/10.1080/07494469600640331.

Pierce, Joseph, and Mary Lawhon. "Walking as Method: Toward Methodological Forthrightness and Comparability in Urban Geographical Research." *The Professional Geographer* 67, no. 4 (July 29, 2015): 655–62. https://doi.org/10.1080/00330124.2015.1059401.

Sheller, Mimi, and John Urry. "The New Mobilities Paradigm." *Environment and Planning A: Economy and Space* 38, no. 2 (February 2006): 207–26. https://doi.org/10.1068/a37268.

Simurra, Ivan, Marcello Messina, Luzilei Aliel, and Damián Keller. "Radical Creative Semantic Anchoring: Creative-Action Metaphors and Timbral Interaction." *Organised Sound* 28, no. 1 (September 2, 2022): 64–77. https://doi.org/10.1017/s1355771822000322.

Thompson, E. P. *The Making of the English Working Class*. Harmondsworth: Penguin, 1980.

Tsuruoka, Toshihisa, Brian Ellis, and Leo Chang. "Ear Talk Project: Participatory Co-Composition on YouTube and the Web." *Organised Sound* 28, no. 1 (February 14, 2022): 110–21. https://doi.org/10.1017/s1355771822000115.

Vinen, Richard. *Second City: Birmingham and the Forging of Modern Britain*. London, UK: Allen Lane, an imprint of Penguin Books, 2022.

Voegelin, Salomé. "Sonic Memory Material as 'Pathetic Trigger.'" *Organised Sound* 11, no. 1 (March 15, 2006): 13–18. https://doi.org/10.1017/s1355771806000033.

Making Machine Learning Musical: Reflections on a Year of Teaching FluCoMa

Ted Moore
Peabody Institute
Johns Hopkins University
tmoore97@jh.edu

James Bradbury
University of Western Australia

Pierre Alexandre Tremblay
Conservatorio della Svizzera Italiana

Owen Green
Department of Music
Max Planck Institute for Empirical Aesthetics

Introduction

The Fluid Corpus Manipulation Toolkit (FluCoMa) enables techno-fluent musicians to use machine listening and machine learning in their creative practice within the familiar environments of Max, SuperCollider, and Pure Data. Housed at the University of Huddersfield's Center for Research in New Music (CeReNeM) in the United Kingdom, the project's primary development period was funded by a five-year grant (2017-2022) from the European Research Council.

While the core of the project produced code packages for all three major computer music programming environments, FluCoMa's vision and impact are broader. FluCoMa is also a collection of learning resources, code examples, commissioned artworks, musicological articles, interviews, podcasts, a philosophy about interface design for creative coding, a conversation about the future of computer music, a curriculum of machine listening and machine learning topics, a community of users around the world, and more. The extent of the FluCoMa ecosystem stems from our belief that "providing the tools" is not enough to achieve FluCoMa's mission: *enable techno-fluent musicians to use machine listening and machine learning in their creative practices*. The materials included here extend this ecosystem to encompass resources for pedagogues that might be used to teach FluCoMa in various settings. While some of the ideas and resources presented below are FluCoMa-specific, many of them are toolkit-agnostic and we hope that they can be used by anyone looking to teach machine listening and machine learning for creative music making.

To truly enable these artists, we also must provide knowledge and inspiration, all in a way that is learnable for our main user base: computer musicians. Topics in machine listening, machine learning, computational thinking, and data science are often not included in the training of electronic musicians and therefore a primary objective of FluCoMa is to build bridges of understanding from the knowledge that comes from computer music training towards a degree of fluency with these topics that enables creative music making. Moreover, the need for a critical humanities perspective on an overly STS-based teaching of these subjects has clearly emerged over the last decades, and such critical, artist-driven epistemic anchors have been documented by Snape and Born (2022) and have been disclosed as our design biases in Green et al. (2019), and reflected upon in Tremblay et al. (2022) and Green et al. (2022).

Design Goals

The learning materials included in this document have been created through participatory, iterative, and interactive design. The first draft of learning materials was developed based on feedback from composers commissioned to create music with early versions of the Toolkit. Their feedback on what was confusing about the tools, resources, examples, etc. led to a few revisions of the software interfaces themselves, as well as to produce a second draft of learning materials which were then used in over thirty workshops around the world with computer musicians from various backgrounds. The feedback from these workshop participants further refined the learning materials and have led to the broad ecosystem of learning resources now found as part of the FluCoMa Toolkit.

Tiered Learning Resources

Because different learners will desire different degrees of fluency with these topics, we have tried to tier the learning resources accordingly. The most proximal resources (such as environment-native help files) provide a working understanding of what an algorithm does alongside musical examples of how it might be used. Careful design of these entry points was a priority in the last design iteration, driven by an awareness of the copy-paste culture of a significant portion of creative coders, as experienced in our teaching as well as documented by Snape and Born (2022). The presence of this mode of coding, especially early in the learning process, heightened our responsibility towards designing inclusive, divergent, and musician-oriented help files and examples, in opposition to a more STS-driven approach.

Additional resources (such as on learn.flucoma.org) provide a deeper understanding that might satisfy one's curiosity and/or build an intuition of what is happening "under the hood," both of which can enable a more informed manipulation of the Toolkit. When appropriate, we also link to more technical and scientific resources, including white papers and more math-heavy learning sites that describes the algorithm from an engineering perspective, should the learner wish to pursue that amount of technical detail.

In some cases, paths for pursuing additional resources are more extensive and less linear. This is especially true for the more complex tools in FluCoMa, such as neural networks which have multiple dedicated web resources with varying degrees of technical information, any one of which might come after an initial introduction, but when taken all together encompass the degree of fluency we propose for our learners and users. For instance, learners who are eager to manipulate the many neural network hyper-parameters might jump to MLP Parameters, while a user who needs a little more time absorbing how a neural network works might opt for MLP Training. (Also see MLPRegressor, MLPClassifier, and Training-Testing Split.)

Tiered learning resources allow the learner to pursue knowledge as far as they deem appropriate to feed their creative practice in a given moment. Providing the learner what they *need* to know *when* they *need* to know it enables them to stay focused on a creative idea and not become overwhelmed by what could be a very large body of knowledge with a daunting learning curve. This sensitivity to the relationship between creative pursuits and technical knowledge reflects earlier findings of FluCoMa (Green at al. 2019) outlined as "Techno-Fluency" and "Divergence," which acknowledges that many people's appetite for technical matters and implementation details is contextual within the divergent priorities of a creative coding workflow. By offering signposts and links to further resources, the user knows where to keep learning if necessary in the moment, or, in the future if they decide to continue exploring.

Music-Forward Resources

Because of the specificity of our target learner, typically, a creative coding musician, we have always tried to keep our learning materials and examples musically oriented (as can be seen below) in line with our critical musicianly-biased design ambitions. We aim to have the help files and example code make sound in a creative way. When possible, we offer pedagogical examples and thought experiments that will feel familiar and relevant to our learner such as instrument samples, drum hits, MIDI notes, synthesizer settings, measures of frequency and loudness, etc. We hope this strategy will not only explain a tool and its interface, but also provide some copy-and-paste code to get started quickly, and generally get the musical creative juices flowing while a user is engaging in the learning process.

Ecosystem of Learning Materials

Creative Coding Environment Materials

Each creative coding environment (CCE) supported by FluCoMa (Max, SuperCollider, and Pure Data) has a native system for offering reference materials. In Max and Pure Data, a "help file" provides annotated examples, while a "reference" offers additional description and detail about parameters. In SuperCollider all of this information is contained in one "help file" document. Despite this difference in interface, we have strived to keep the CCE-based FluCoMa materials similar across all three environments to hopefully foster, as per the project's original aim, discussions of creative data mining beyond the allegiances to, and affordances of, one's CCE of choice. This is enabled, in part, by the shared documents used to render the reference materials for all three CCEs and the parity of examples between them.

The FluCoMa materials provided natively in the CCE are often the learner's first engagement with our supporting materials. Therefore, the information provided is intended to provide the learner/user a working understanding of what an object does, how it might be used musically through a sound-making example, and if appropriate, how it interfaces with other FluCoMa objects. This information will hopefully provide a learner some motivation for exploring an object and the amount of knowledge necessary to do so. Each resource in the CCE contains a link to the corresponding web reference if the learner wishes to pursue a deeper understanding of the tool.

The example code provided has been created to be as similar as possible across the three CCEs while keeping idiomatic to each environment. One goal of this is to enable cross-environment communication and knowledge sharing. Users in different CCEs are able to discuss the technical and musical facets of a shared FluCoMa example. It is also possible that this allows referencing help files and example code to a classroom containing a diversity of coding environment users. Lastly, this keeps all three CCEs on an equal status, preventing any potential inference of CCE preference within the FluCoMa Toolkit.

Web Reference

Every object in FluCoMa (except a very small number of CCE-specific helper objects) has a web reference found at learn.flucoma.org. Because the reference materials that appear natively in the CCEs link to the web, the web references are considered to be a secondary resource. The goal of the web references is to offer more detailed descriptions of how an algorithm is working "under the hood."

Many of the web references have interactive explanations that allow a learner to "use" the algorithm in the browser. Not only can these be used by individual learners, we have also found that these are very useful for explanations in the course of teaching. For example, when creating a KDTree (k-dimensional tree) for the first time during a code-along class, using the interactive page helps give learners a visual sense of what is happening. This is especially true when the KDTree is being used with the plotter. Another example is the Mel-Frequency Cepstral Coefficient (MFFC) reference which has an interactive explanation that invites a learner (or a teacher during demonstration) to step through a series of interactions that build intuition about MFCCs (Figure 1).

We imagine the web references to be used as solo-learning resources, in parallel with class assignments, as teaching demonstrations, and/or as useful reminders.

Learn Articles

FluCoMa's learning website (learn.flucoma.org) also contains many articles about topics that may not fit in a single web reference page. These articles may arise as a tertiary step in a learner's path and are likely to be encountered after the CCE materials and web reference.

There are a few varieties of articles found in this category:

- explainers specific to a single FluCoMa object that offer a depth of knowledge about the internal algorithms that would be outside the scope of a web reference page, such as Audio Decomposition using BufNMF.

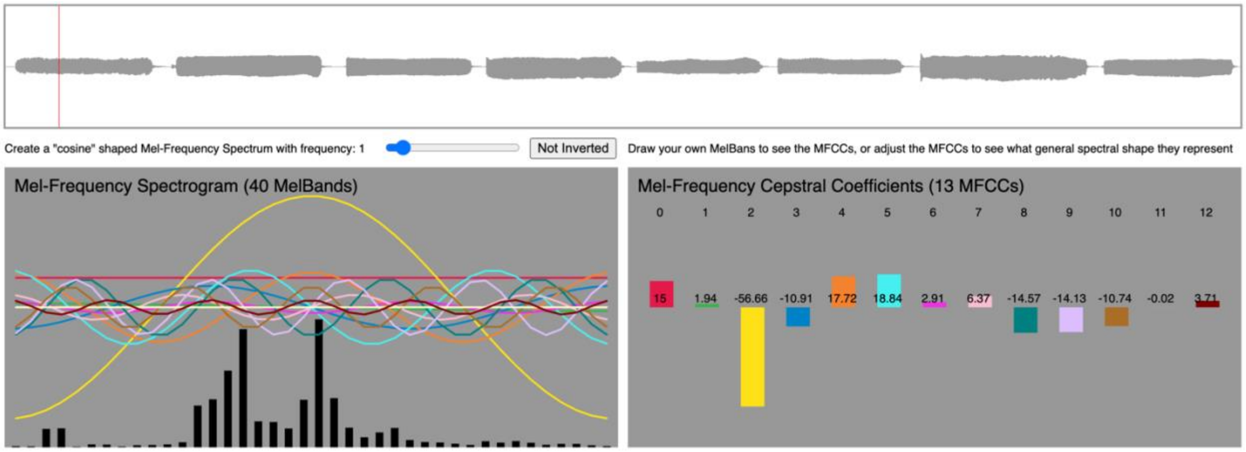

Figure 1. *Screenshot from MFCC interactive demonstration.*

- knowledge about data science that is useful for using many of the FluCoMa objects, such as Distribution and Histograms and Why Scale? Distance as Similarity
- common workflows using the toolkit, such as Batch Processing with FluCoMa

These articles are not necessarily designed to be consumed in series as part of a sequence of learning (although some could be used this way). Instead, each article is made to be approached by a learner (or guided by a teacher) at a particular point in the learning process and revisited as necessary. The idea for many of these articles arose in direct response to questions asked by workshop participants and therefore are designed to answer or provide context to common questions asked by learners. When designing a curriculum or syllabus, many of these articles would support student learning for different topics in a course.

Explore Articles

In addition to the web reference and learn articles, the website has many additional materials focused on inspiring artistic uses of the toolkit. These include example artworks, interviews with creative coders, and musicological articles that offer in-depth analysis and example patches of music made with FluCoMa. All of these can be used as context, examples, and inspiration for learners. These can also be used as entry points, as many learners will find the work produced using FluCoMa or the musical ideas expressed in these articles inspiring and motivating.

Discourse

The international community of FluCoMa users primarily communicates through the Discourse online discussion forum. Any learner (or user) of FluCoMa is invited to be a member of the Discourse as it is an excellent place to converse with like-minded artists. Learners may find the search functionality very useful to see if others have already asked and answered a question they have. The community is designed and supervised to foster positive and supportive interactions, so it is a good place to ask all kinds of questions. In addition to the threads for *Getting Started and Wayfinding* and *Usage Questions*, there are also threads for *Code Sharing*, *Learning Resources*, and *Interesting Links*, making it another place for learners to browse for examples, inspiration, and knowledge.

Inter-Connectivity of Resources

As described above, the FluCoMa learning resources are generally tiered to offer learners the degree of detail needed at a given moment to pursue a creative idea. One way in which our tiered approach is executed is cross referencing the different learning resources. The help files in each creative coding environment (CCE) link to their respective web reference, from which a learner can be pointed to many more resources (Figure 2). The web reference, learn articles, and

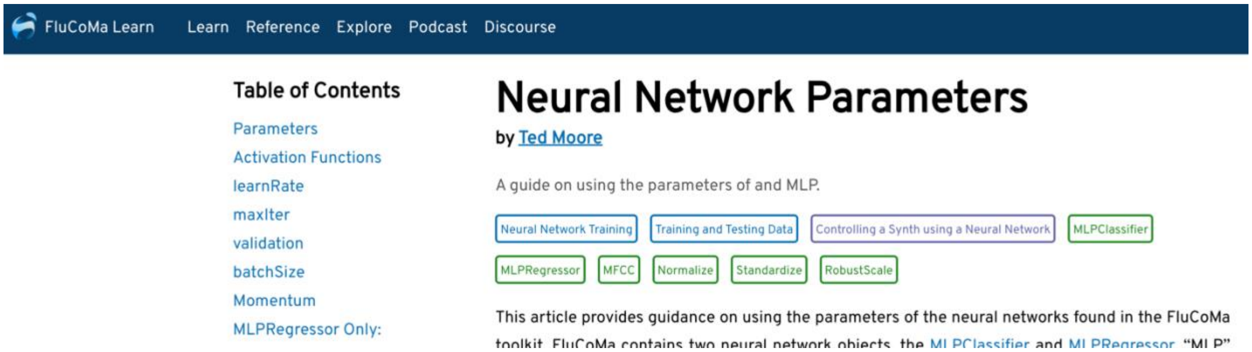

Figure 2. *Screenshot from "Neural Network Parameters" page showing links to other related pages as colored boxes.*

explore articles all cross-link with each other so that a learner reading a web reference might discover an Explore Article about a musician's use of an object, and from there discover a Learn Article to help them pursue the creative use they just learned about, etc. Different learners will need different degrees of technical specificity, inspiration, and modes of engagement at different times. We hope that setting someone "loose" on the website will enable them to find uses of FluCoMa that are meaningful to them as well as the knowledge needed to support them.

Example Lesson Plan: MLPRegressor

After teaching numerous workshops we have identified a few class plans that work well to get new learners excited and making sound while laying a foundation of facility with FluCoMa to support further activities and/or self-guided learning. The lesson plan summarized here introducing neural networks with the MLPRegressor is often a first instruction to FluCoMa.

Watch a Video Tutorial of this Lesson Plan:

- MLPRegressor in Max
- MLPRegressor in SuperCollider

The first activity that we often engage learners with is building a neural network that performs regression to control a synthesizer with ten control parameters from a space of only two parameters. The lesson takes anywhere from 40-90 minutes depending on the class of learners. This activity gets learners making sound quickly and uses part of the toolkit that is often quite exciting for newcomers to machine learning (neural networks). This activity has been strongly influenced by Rebecca Fiebrink's Wekinator example and philosophy of "small data is beautiful data."

Here is a brief outline of the lesson plan:

1. Share a real world example (including watching a performance excerpt) of why someone might want to use a system like this.
2. Using a slides presentation step through how we will be collecting training data and training the neural network, including some intuition about how the training process works.
3. Open up the CCE of choice and demonstrate a completed version (Max, SuperCollider) of the instrument we're about to code.
4. Code the instrument together, as a code-along, starting from a "starter patch" (Max, SuperCollider) that has a few key items already in place:
    - a synthesizer to control.
    - a 2D control space to use as input to the neural network.
    - a MLPRegressor object with many arguments already specified.
5. Let the learners play with the instrument (and augment it in their own way).

The starter patch is important here so that we don't spend too much time doing CCE-specific boiler plate code but instead get right into using FluCoMa. It also ensures that learners have a synthesizer to make sound with right away when the code-along is complete. Because there are many arguments to the MLPRegressor object and each of them can require a fair amount of explanation to use well–and in coordination with each other–we've chosen to provide the arguments to the MLPRegressor programmed into the starter patch. During the lesson we tell the learners that these arguments can be explored further in a future lesson and/or in the learn article.

The extensions of this activity are to:

1. Practice training the neural network.
    - Clearing the neural network and retraining
    - Training the neural network to a different degree to see if it is more (or less, or differently) musically expressive.
    - Delete the input data and choose new data points to pair with the synthesis parameters.
    - Delete all the data and create a whole new training.
2. Attach a different sound-making algorithm to the output of the neural network.
    - Granular synthesis / sample playback
    - Frequency modulation
    - A VST plugin
    - We have often encouraged learners to bring a sound-making algorithm of their design to the workshop to connect as a next step to this activity.
3. Attach a different type of controller to the input of the neural network. This might be something like:
    - Multiple parameters on TouchOSC
    - MIDI controller
    - Leap motion
    - Wearable device
    - Pixel information from a camera (perhaps using Jitter in Max)

Not only does this activity quickly provide learners with a machine learning instrument that is very extensible, it also introduces some of the key elements of FluCoMa:

1. DataSets
2. Buffer interfacing
    - `fluid.buf2list~` and `fluid.list2buf~` for Max
    - `FluidBufToKr` and `FluidKrToBuf` for SuperCollider
    - `array set` and `array get` in Pure Data
3. Using small, personalized, artist-created DataSets
4. Aesthetic evaluation of results
5. Iterative trial-and-error workflows with machine learning algorithms

Classifier Extension

After completing the MLPRegressor activity, one common extension is to do an activity training the MLPClassifier to distinguish between two timbres. Depending on the group of learners and how much time is available, sometimes this would only include opening up the classifier demonstration file (Max, SuperCollider) and doing a quick training and testing, along the way relating it to what was just done with the MLPRegressor activity. If time allowed and there is interest, we performed a more involved activity similar to the MLPRegressor: demoing the code and then building it together.

Watch a Video Tutorial of this Lesson Plan:

- MLPClassifier in SuperCollider
- MLPClassifier in Max

1. Share a real world example (including watching a performance excerpt, this one has a before and after) of why someone might want to use a system like this.
2. Using a slides presentation step through how we will be collecting training data and training the neural network,

including some intuition about how the training process works.

3. Open up the CCE of choice and demonstrate a completed version (Max, SuperCollider) of what we're about to code.
4. Code-along, starting from a "starter patch" (Max, SuperCollider) that has a few key items already in place:
    - The sound files containing timbres that will be used as training and testing data.
    - A MLPClassifier object with many arguments already specified (for the same reason described above).
5. Let the learners explore classifying some of their own sounds and test the effectiveness of different audio descriptors.

Common Learning Challenges and Strategies

Over the course of teaching many workshops, we observed some common challenges for FluCoMa learners. Below are a few of the challenges we found and some strategies for approaching them pedagogically.

New Ways of Using Buffers

FluCoMa uses buffers to store all kinds of data, not just audio. This may be new for learners who are used to using buffers *only* for holding audio, and may even conflate the two as a single concept ("buffer == audio"). This becomes increasingly complicated when we begin to manipulate the data in buffers as arrays and matrices.

Initial Encounter

The first moment at which a learner is asked to view a buffer in a new way is often when we allocate a buffer to hold a data point. If the data point has just 2 dimensions (such as with the MLPRegressor activity), we will allocate the buffer with only 2 frames (in Max: @samps 2; in SuperCollider: Buffer.alloc(s,2)), at which point we reflect on the concept of a buffer being a container of values that often holds samples we play back at a rate of 44,100 per second. However, these values could actually represent anything and playing them at a specific rate through a digital-to-analog converter is just way of using values.

Holding Analyses

Once we start writing audio analyses into buffers (with the feature argument), learners often have a hard time keeping track of the structure of the buffers (What do the channels represent? How many are there? What do the frames represent? How many are there?). We found that offering CCE-agnostic charts of the "shape" of the buffer is very helpful for giving learners a mental model (Figure 3).

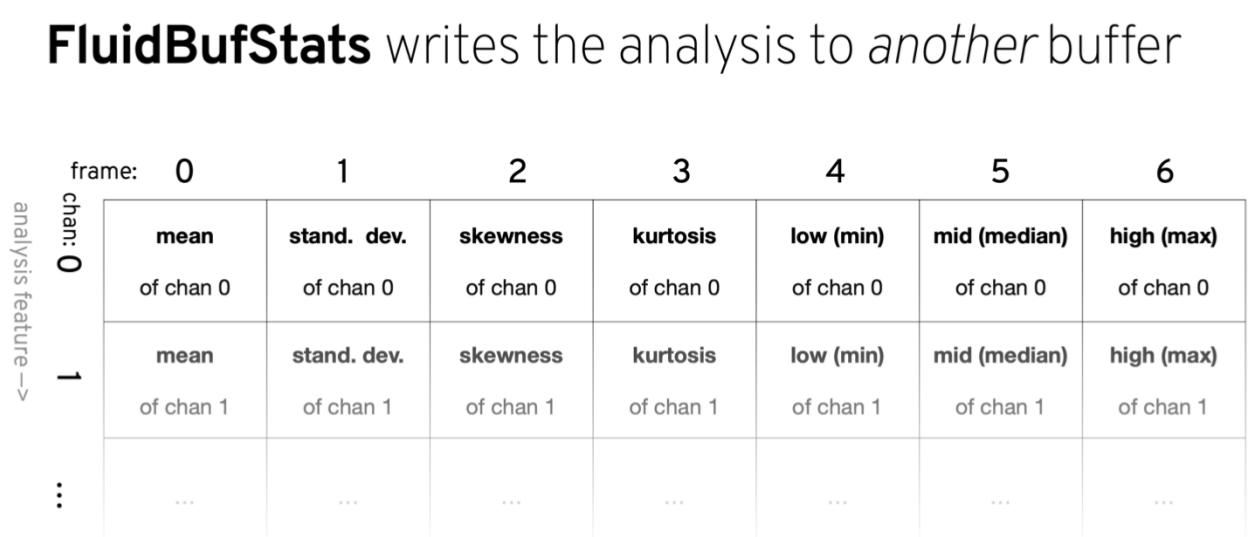

Figure 3. *Chart demonstrating what the channels and frames of a particular buffer represent.*

It's also useful to point out that for buffers that hold audio analyses, the frames (or what we sometimes refer to as the "x axis" in reference to the charts above) is still a time series, just like audio is, but now it's not a time series of voltages (as in audio), it's a time series of descriptors (such as spectral centroid).

Because each frame represents an FFT frame from the STFT (short-term Fourier transform) analysis, the sample rate would not be a usual 44,100 samples per second, but a much lower rate of frames per second (FFT frames per second). For example, if an audio buffer with a sample rate is 44,100 Hz is analyzed with a hopSize of 512 samples, the features buffer that the analyses get written into will have a sample rate of 86.1328125 frames per second (44100 / 512). In SuperCollider and Max, FluCoMa buffer processors (such as the audio descriptor analyzers) set the *sample rate* of these buffers appropriately (Pure Data arrays don't hold metadata). If the values in that buffer are read back at that rate, they will correspond in time (be synchronized with) to the audio on which the analysis was based. Pointing this out to learners helps them remember and conceptualize the relationship between the source audio being analyzed, the STFT process, and the resulting time series of descriptor values.

Padding: Where is My Data

Once the relationship between audio sampling rates and the sample rate of a descriptor's time-series is clearer, another challenge for learners is matching indices of analyses in a descriptor's time-series buffer with a given moment in the source buffer. This confusion can increase once all the parameters of windowed and padded Fourier analysis are properly explained. Again, an interactive-graphic illustration was developed to help learners build understanding and intuition, this time created using Max.

Manipulating and Copying Data

Often it is necessary to manipulate the data in a buffer, such as picking out values from certain channels and/or frames and copying them to another buffer. In order to provide some "test and check" interactivity to build fluency with these operations, the appropriate web references have interactive GUIs for practicing (Figure 4). These include:

- BufSelect
- BufSelectEvery
- BufFlatten
- BufCompose
- BufScale

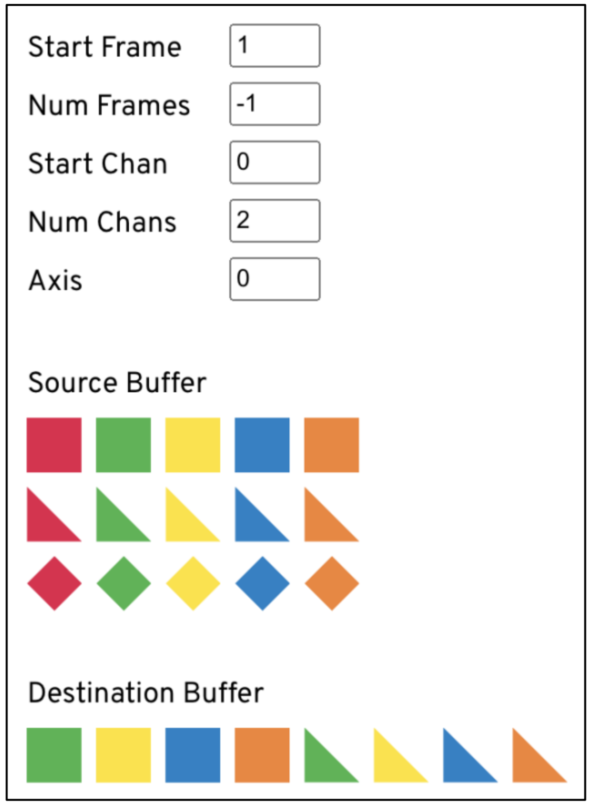

Figure 4. *Interactive web GUI for "testing and checking" the transformations of BufFlatten.*

Why Buffers

It also may be of interest for learners to hear the explanation of *why* buffers are used in this way. FluCoMa uses buffers in this way for a few reasons:

- The notion of "buffer" is shared across all three CCEs that FluCoMa supports. This allows for shared syntax and usage of objects across all three environments.
- The ease of exporting buffers as 32-bit float PCM files allows for interfacing

with applications and languages outside FluCoMa.
- In all three CCEs, buffers are accessed at the lower levels of code allowing for:
    - Faster processing by interfacing directly with the C++ code.
    - Simpler implementation of functions across all three CCEs because all three environments share the same FluCoMa core C++ code for audio analysis, buffer processing, and algorithms.
- Having data in buffers allows for it to be more flexibly accessed and used in other parts of the CCE. For example, because the MLPRegressor writes predictions into a buffer, it's possible to be predicting wavetable shapes directly into a buffer that is simultaneously being read out of.

Fourier Transform and STFT

As with many audio tasks, the STFT is central to much of how a learner interacts with FluCoMa. Many times learners can do exciting things, learn a lot about the toolkit, and make great music without reflecting on the Fourier transforms happening "under the hood." We have found however, that for many of the algorithms in FluCoMa (such as AudioTransport, Sines, and many more), adjusting STFT settings (windowSize, hopSize, and fftSize) has an important aesthetic impact on the results, and therefore we suggest that it is important to understand what impact these parameters have.

There exists a wealth of resources on the internet for learners to build fluency with the Fourier Transform, so we didn't feel the need to recreate these learning tools. We have however, curated a small set of musician-oriented ideas that go just deep enough for learners to gain access to more musical expression from these FluCoMa tools.

- Fourier Transform
- BufSTFT

Stateful Objects

Many of the FluCoMa data objects hold some state. For example after calling fit on a Normalize object, it holds the minimum and maximum value of each dimension in the fitted DataSet so that it can scale future transform calls appropriately. MLP objects hold the state of the MLP internal parameters (the state of the trained model). This interface design is based on many of the data processing objects in the Python sci-kit learn package. We found that some FluCoMa learners find it challenging to conceptualize or remember that certain objects are holding a state that they will need to call upon later. One feature that may help with conceptualizing objects in this way is the option to *name* objects in Max and Pure Data (SuperCollider natively uses variable names to identify objects).

Named objects may help learners remember that certain objects hold state because they have a sense of it being a non-generic, task-specific object, such as a Normalize object called norm-pre-pca. This gives it a special sense of purpose and an indicator of what state it holds and where in data processing one would call upon that state.

Advanced Neural Networks

Learners often follow up the MLPRegressor Lesson Plan with questions about the hyper-parameters (which FluCoMa calls parameters or arguments to keep in line with the CCE language musicians are used to seeing) of the MLP. In shorter workshops (two days or fewer) we have felt that it's not enough time to delve into this with enough depth to make it well understood and *useable* for the participants, so we've directed them towards our web resources on the topic. In longer workshops (three days or more) we have taken time later in the week (day three or four) after the introductory activity to unpack many more ideas and strategies about the MLP objects.

Web resources:

- Neural Network Training is an overview of how neural networks "learn." It is intended for those who would benefit from gaining a little more intuition about what is going on "under the hood" or for learners that have a little more curiosity they want to satisfy. Much of what this article expresses is included in the introductory activity.

- Neural Network Parameters goes through each parameter in the MLP objects and gives a more thorough description of what it controls, why one might adjust it, and what a generally reasonable starting place is. This is often where we direct learners who ask about these parameters when we don't have time to unpack them during a workshop.
- Training and Testing Data describes why it might be important to validate the results of a trained MLP. It explains why one would go about validating a model, what to look out for, what certain results might mean, and what one might do to improve a model.

Teaching Materials

The sequence of explanation that we've used for both the MLPRegressor and MLPClassifier seems to work quite well for giving learners intuition about the training process of an MLP. These explanations can be seen at the beginning of the MLPRegressor and MLPClassifier tutorial videos.

As stated above, when we have had time in workshops, we've allocated time to explain the parameters in more detail. A very useful site for explaining and playing with the momentum parameter can be found on distill.

One might notice in many of the resources above that there are node-and-edge graphs of MLP architectures. We found that these are very useful for learners to visualize and concretize a few facets of MLPs: (1) "feed-forward", (2) "back-propagation", (3) "fully-connected layers", (4) numbers of hidden layers and nodes, (5) total number of parameters in an architecture, and more. We also learned that it is important to have a visual representation of the architecture that is *actually used* in the activity. We use a very basic graphviz script to generate these graphs (Figure 5).

Navigating Human and Machine Assumptions

One of the most challenging conceptual hurdles for FluCoMa learners is to reconcile the differences between the way machines and humans listen. What perceptually might seem obvious to a human listener can be very challenging for a machine to discern. Newcomers to machine listening often set out to perform a task making many assumptions about how a system will work, what data they will use, and how they will compute a result, not realizing that what they *think* they're telling the machine or asking it to compute is quite different from what it will give back in return.

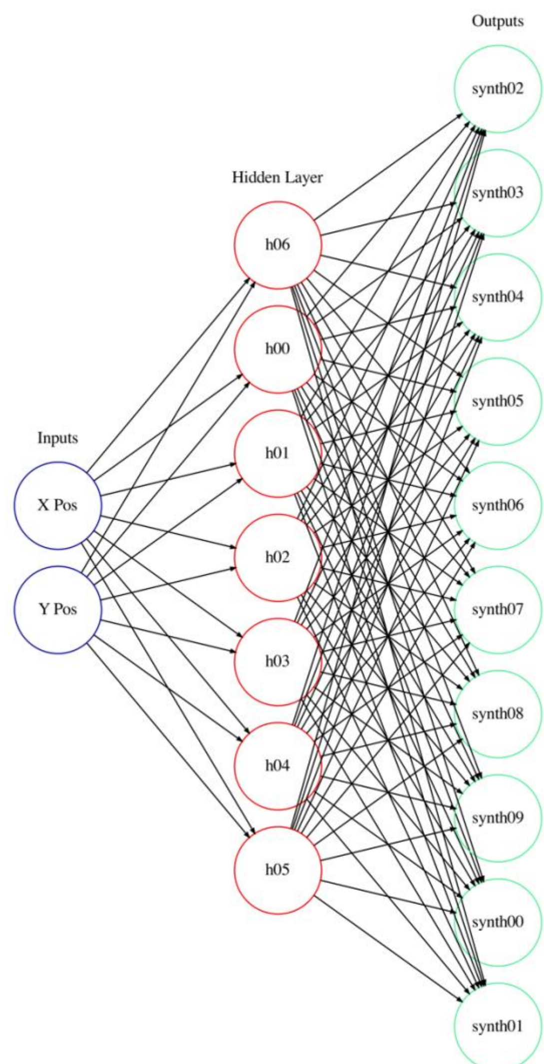

Figure 5. *An example graph visualizing the structure of a neural network. This particular architecture is used in the Example Lesson Plan: MLPRegressor.*

Machine Listening: Pitch

One activity that has been quite successful is a simple "listening test." We ask a class of learners to sing the pitch of a sound file. Most listeners will sing the right pitch class of the tone but down a few octaves from the actual frequency. The first thing to point out to the group is that they were only singing the pitch from the part of the sound file that was *most* pitched. They didn't even attempt to sing the "pitch" during the scratchy parts. We then look at a chart showing the result of a Pitch analysis on the buffer (Figure 6). One can see where the pitch is stable (the parts that the listeners sung), but also that there is a lot more "pitch" analysis there. The machine listens to all of it (and reports back on all of it). This is because as humans we're constantly engaged in multi-modal listening, switching between different ways of perceiving sound, depending on which seems appropriate or useful at a given moment. When the sound file is making scratchy sounds, we as humans don't even register it as a "pitch" to sing, but a machine does. We also highlight the second dimension, the pitch confidence, and how it could be used as a descriptor in itself.

Another useful outcome of this activity can be acknowledged when discussing Distance as Similarity. Most listeners will sing the correct pitch class but down a few octaves. For these listeners, being 12 half steps lower (a distance of 12) is closer than being 1 half step lower (a distance of one). This is again a recognition that what humans might assume to be *similar* a machine might not. Recognizing these misalignments offers a great opportunity to reflect on how one might bring the machine's perception closer to our own, for example, using the confidence measure from a Pitch analysis may help focus the machine on the *more* pitched moments and a Chroma analysis would ignore the octave.

Statistics

Many workflows in FluCoMa require the use of statical summaries of audio descriptor time series, real time audio analyses, or whole DataSets. It is important for learners to develop a sense of what tools are available and why one might reach for one statistical summary rather than another.

BufStats is perhaps the most commonly used object in FluCoMa and therefore has a somewhat involved reference page. Many of the statistics available might be familiar to learners (mean, median, minimum, and maximum), while others might be new (standard deviation, skewness, kurtosis, derivatives, etc.). During the introductory tutorials, many moments arise that are useful for reflecting on the statistical analyses being used and how they affect the sonic results being heard. A few examples are reflecting on what it means to have an average spectral centroid of a sound slice and using the maximum value of an analysis rather than the mean.

BufStats has many more features that are somewhat less explored and probably not appropriate for learners just getting acquainted with FluCoMa. There are few Learn Articles that cover these topics in musicianly ways including Weighting Stats and Outliers. It is also important to convey to learners, as is stated on the BufStats page:

> *While it can be difficult to discern how to use some of these analyses practically (i.e., what does the kurtosis of the first derivative of*

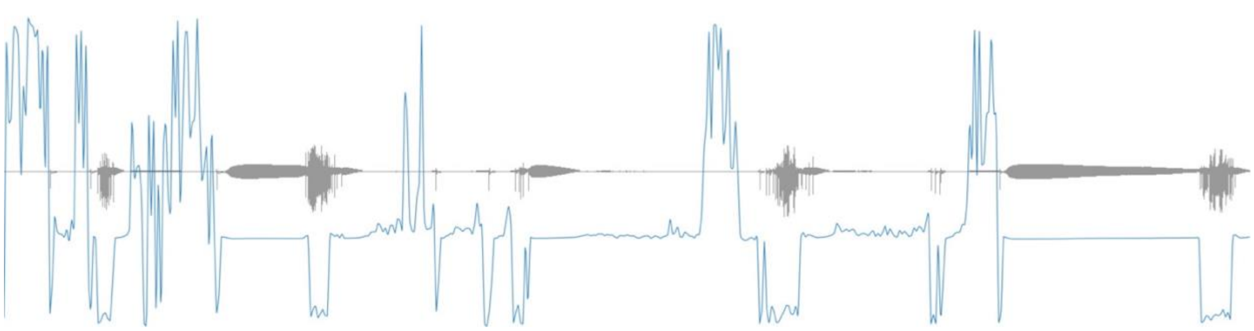

Figure 6. *Pitch analysis of a buffer that has some pure tone and some scratchy timbres.*

spectral centroid indicate musically?), these statistical summaries can sometimes represent differences between analyses that dimensionality reduction and machine learning algorithms can pick up on. Including these statistical descriptions in training or analysis may provide better distinction between data points.

"Know your Data"

As with all data science and machine learning, understanding what data represents, what it can tell you (and more importantly what it *can't*), and what transformations *do* to data is essential. It is continuously important for FluCoMa learners to reflect on their data. There are a few visualization tools that are very important for users to get comfortable with including a two-dimensional plotter and a multichannel waveform and time-series viewer. While teaching, these tools should be used whenever possible to help learners understand the data processing that is happening and get them in the habit of visually checking on their data regularly to build understanding of what their data represents (Figures 7 and 8).

Many machine learning algorithms make various assumptions about data (similar to how humans make assumptions about sound and machine listening). One of these assumptions is that the dimensions in a DataSet are identically distributed, often Gaussian distributed. It can be important to know how one's data is distributed and it is possible to check on using a histogram. Our Distribution and Histograms page gives some examples of different kinds of distributions, what they mean, and some example code to check on a distribution using a histogram.

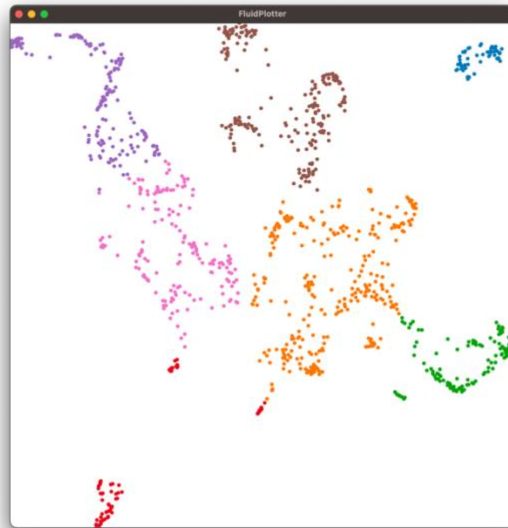

Figure 7. *Screenshot of FluidPlotter used in SuperCollider.*

Scalers and Distance as Similarity

Another important concept for learners to understand is how measures of distance impact perceptions and assumptions about similarity. Once the machine has "listened" and a statistical summary has been computed, a next step is often to compare data points by computing the distance between them (most of the FluCoMa tools use Euclidean Distance). Computing distance makes questions about ranges and scaling relevant, such as how a mismatch of scale may overly weight the importance of dimensions that have larger ranges.

This is a great opportunity to compare scalers available in FluCoMa. One way of clearly demonstrating that different scalers will have different sonic results (and that those sonic results are not always *predictable*) is to choose a single point in a DataSet (such as one sound slice) and find what the nearest neighbor is with (1) no scaling, (2) Normalize, (3) Standardize, and (4) RobustScale. (This is essentially what the sequence of images does in the Comparing Scalers page.) Doing these comparisons in real time while hearing the sonic differences can help concretize the importance for learners.

It can often be important for learners to keep track of which dimensions might be logarithmic or linear and know how those differences could affect measures of distance in relation to human perception. One concrete example we provide is on our scaling page under the heading Linear vs Logarithmic Scales where we state:

For example, frequency analyses might be provided in hertz (which is on a linear scale),

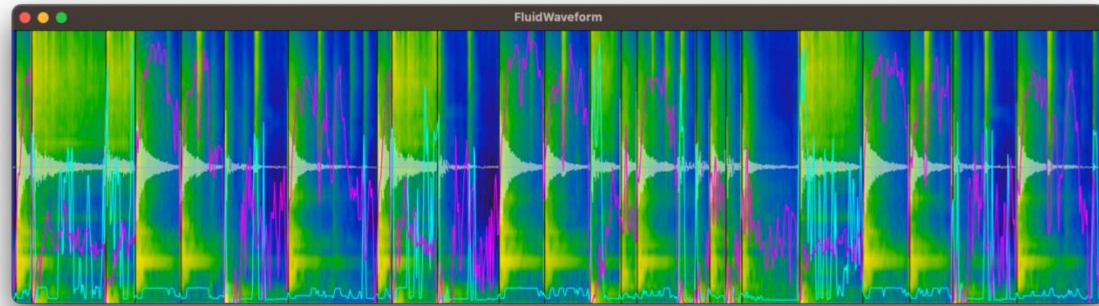

Figure 8. *Screenshot of FluidWaveform used in SuperCollider.*

> *however this doesn't reflect how humans actually perceive pitch distance. For a more perceptually relevant scale it is displayed logarithmically in pitch space (perhaps labeled as MIDI notes or semitones). If measuring in hertz, the distance from the A4 down one octave to A3 is 220 hertz, while the distance from A4 up one octave to A5 is 440 hertz—twice as far even though we perceive them to both be one octave! Measuring these distances in semitones will reflect the way we perceive them: both exhibit a distance of 12.*

FluCoMa analyses which report units in hertz, such as pitch and spectral centroid, have an argument called unit which specifies if the hertz value should be returned in hertz or MIDI notes.

Human vs. Machine Assumptions

One more concrete example of how human and machine assumptions differ comes from a learner who has having trouble getting KMeans to cluster data points in the way they thought it should. The learner wanted KMeans to cluster points on a 2D plot according to the clusters that are easily visually identifiable by a human. KMeans was clustering it much differently, including leaving many clusters empty. This was solved by demonstrating (including with the original data) how the learner could seed KMeans using a "human in the loop" approach to direct its processing with information also inferred by a human.

De-Myth-ifying Machine Learning

Sometimes we have questions from learners that sound something like, "I want X to do Y. How can FluCoMa do this?" At this point, our pedagogical step is to break down the goal into smaller and more specific questions and tasks that we can start approaching together with the learner. This process often reveals the assumptions that the learner might be making about how audio analyses work, or what a machine will be able to perceive, or how long an analysis or algorithm might take, and all of the tradeoffs involved in making decisions about the process. Sometimes the question transforms from "How can FluCoMa do this?" to "Can FluCoMa do this?" at which point perhaps there is a different tool that we can point them to, or help them realize that their goal is too lofty–that it stems from a belief that "AI can do anything" or "throw it at a neural network and it'll figure it out." Usually, this process enriches the learner's ideas about what is possible with FluCoMa (even if it's not necessarily what they hoped) and provides a lot of possibilities for investigation.

Another de-myth-ification that has occurred is when learners will assume that the machine learning *is* performing some magic when it is not. This is often in the form of learners not *validating* or testing the machine learning model or the results of their algorithm. The first disclaimer to make is that, as artists, we're interested in artistically compelling experiences, so regardless of what the algorithm is or isn't doing, if the user thinks it sounds good, we encourage them to keep it. It can also be important to test the systems that we build to see if they are doing what we think they're doing. This can be beneficial for a few reasons:

- Validation can reveal our assumptions and/or misunderstandings about *how* things work, providing opportunities to deepen our knowledge and skills.
- Validation can offer ways to improve our system to get even closer to our desired outcome, such as realizing that a system is overfit if generalization is what is desired.
- Validation can reveal nuances in the system that might offer more paths of exploration and creativity.

Framing validation with these benefits in mind can help encourage learners to put in the extra work that it takes.

Relevance to Contemporary Society

We believe that learning about data, data science, and machine learning through FluCoMa can be used as a lens to consider how these tools operate in contemporary society, in particular to uphold inequalities, injustices, and hegemonies. By gaining fluency and understanding with these algorithms, one can come to understand what these algorithms are good for, what they are not good for, how they go wrong, and the relationship between data, algorithms, and the humans that use them. The skills and knowledge, both explicit and tacit, that working with FluCoMa fosters can be used to reflect on many of the AI events and concerns that are constantly appearing in the news. Pedagogues might draw on these events to use as discussion topics where a classroom of learners can collectively reflect on contemporary topics using the experience and understanding built through FluCoMa.

Below is a list of books (in no particular order) that provide many examples of contemporary technologies negatively impacting marginalized communities. Many additionally offer directions for how to approach and use data science ethically. We recommend selecting a book or selected readings from these books to augment learning. Each of these is written for different audiences, so selecting which is best is at the instructor/learner's discretion.

- *Data Feminism* by Catherine D'Ignazio and Lauren F. Klein.
- *Weapons of Math Destruction* by Cathy O'Neil
- *Hello World* by Hannah Fry
- *Revolutionary Mathematics* by Justin Joque
- *Blockchain Chicken Farm: And Other Stories of Tech in China's Countryside* by Xiaowei Wang
- *The Alignment Problem* by Brian Christian

Conclusion

FluCoMa has now reached many thousands of people in the form of artists, students, pedagogues, and audiences. More examples, explainers, and divergent uses continue to surface from the user base, expanding the vision of the project. When traveling to concerts and conferences in the field, we often encounter users who are eager to share with us their work that uses FluCoMa, demonstrating the project's success at building community around a shared excitement of using these tools for music making.

The curriculum of learning materials has enabled users to learn data science and machine learning principles and also has propelled them to learn more beyond what is included in the FluCoMa Toolkit and learning resources. Some users have requested extensions to FluCoMa that would enable more complex or idiosyncratic machine learning systems. These requests demonstrate that our presentation of foundational knowledge, paired with modeling divergent and creative uses rooted in artistic practice, has permeated to our user base who is now imagining and pursuing their own creative and divergent artistic ideas using machine learning and data science tools. We hope that these ways of thinking will resonate forward from FluCoMa as new tools and artistic expressions emerge.

All the products of the FluCoMa Research Project are open source and we welcome the raising of errors, omissions, suggestions, and contributions. We hope that the positive impact of FluCoMa will continue through the addition of music, code, and ideas from a growing community.

References

Fiebrink, Rebecca. "Machine Learning as Creative Design Tool." Presentation at the Centre for Research in New Music (2019). Archived on YouTube: https://www.youtube.com/watch?v=Qo6n8MuEgdQ.

Fiebrink, Rebecca, and Perry R Cook. "The wekinator: a system for real-time, interactive machine learning in music." In *Proceedings of The Eleventh International Society for Music Information Retrieval Conference*, ISMIR 2010, Utrecht.

Green, Owen, Pierre Alexandre Tremblay, Ted Moore, James Bradbury, Jacob Hart, Alex Harker, and Gerard Roma. "Architecture about dancing: Creating a cross environment, cross domain framework for creative coding musicians." In *33rd Annual Workshop of the Psychology of Programming Interest Group (PPIG)*, 14–24, 2022.

Green, Owen, Pierre Alexandre Tremblay, and Gerard Roma. "Interdisciplinary research as musical experimentation: A case study in musicianly approaches to sound corpora." In *Electroacoustic Music Studies Network Conference: Electroacoustic Music: Is it Still a Form of Experimental Music?* Zenodo, 2019.

Snape, Joe, and Georgina Born. "Max, Music Software, and the Mutual Mediation of Aesthetics and Digital Technologies." *Music and Digital Media: A planetary anthropology* (2022): 220-266.

Tremblay, Pierre Alexandre, Gerard Roma, and Owen Green. "Enabling programmatic data mining as musicking: the fluid corpus manipulation toolkit." *Computer Music Journal*, 45(2):9–23, 2022.

Drake Andersen
Editor-in-Chief, Journal SEAMUS
journal@seamusonline.org

Department of Music
Central Connecticut State University
1615 Stanley St.
New Britain, CT 06050

www.seamusonline.org