Journal SEAMUS

Spring/Fall 2018

The Society for Electro-Acoustic Music in the United States

Volume 29, Number 1–2

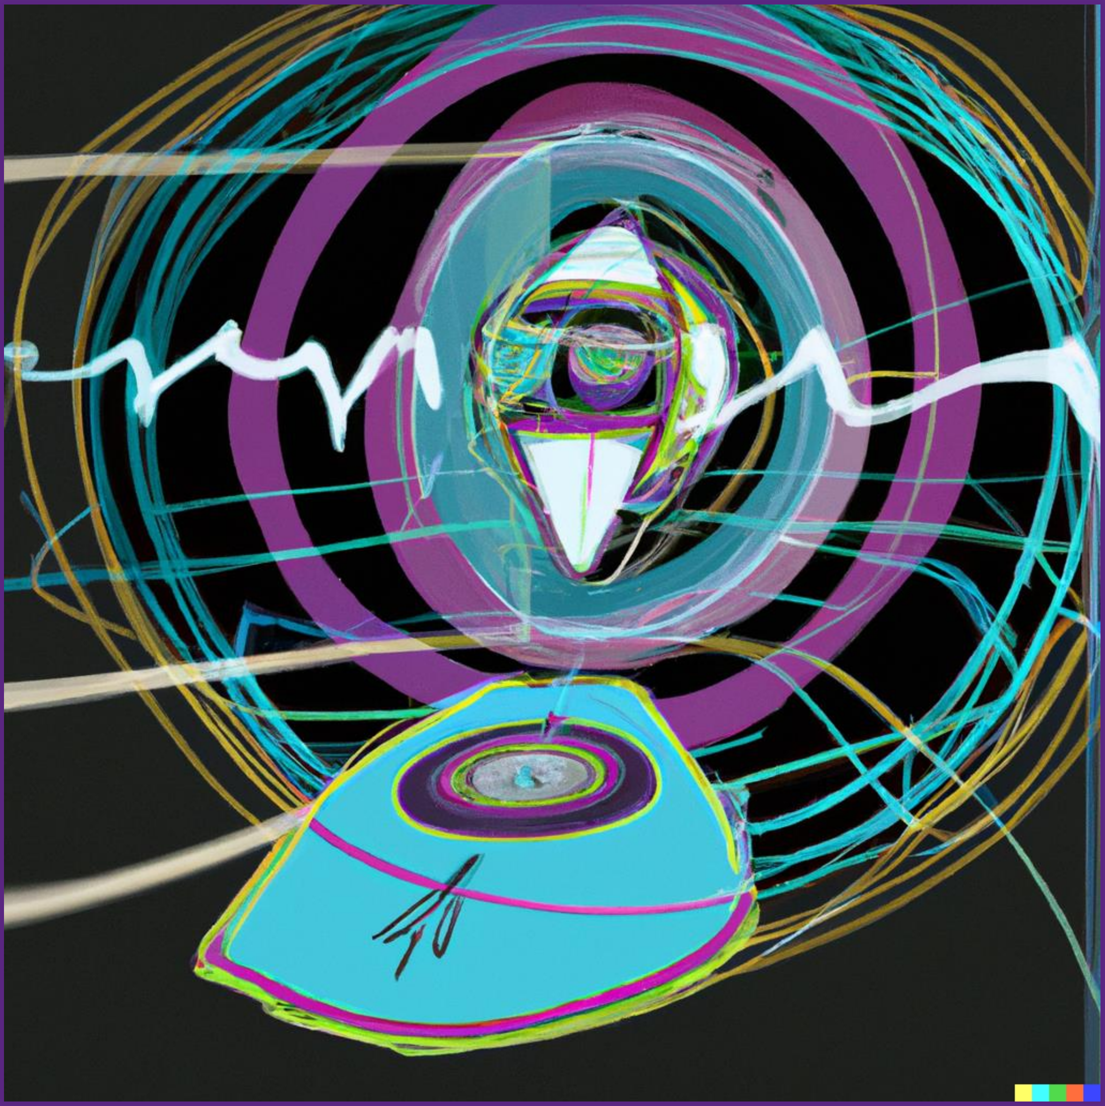

ISSN 0897-6473

Contact Information
Drake Andersen
Department of Music
Vassar College
124 Raymond Ave.
Poughkeepsie, NY 12604

journal@seamusonline.org

Journal SEAMUS
Volume 29, Number 1–2
Spring/Fall 2018

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 non-profit national 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, which 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 of 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

President	Elizabeth Hoffman	president@seamusonline.org
Vice President of Programs	Leah Reid	vp_programs@seamusonline.org
Vice President for Membership	Abby Aresty	vp_membership@seamusonline.org
Member at Large	Eli Fieldsteel	eli.fieldsteel@gmail.com
Member at Large, Outreach	Ralph Lewis	ralphlewis.music@gmail.com
Treasurer	Heather Mease	treasurer@seamusonline.org
Secretary	John Gibson	johgibso@indiana.edu
Editor, SEAMUS Newsletter	Jiayue Cecilia Wu	newsletter_editor@seamusonline.org
Editor, SEAMUS Journal	Drake Andersen	journal@seamusonline.org
Diversity Officer	Elizabeth Hinkle-Turner	dir_conferences@seamusonline.org
Director, Recordings	Scott Miller	dir_recordings@seamusonline.org
Director, Tech. Development	Becky Brown	techdev@seamusonline.org
Director, Communications	Mark Vaughn	dir_communications@seamusonline.org
Immediate Past President	Ted Coffey	ejc3q@virginia.edu

Contents

Journal SEAMUS

Volume 29, Number 1–2, 2018

From the Editor

It is my sincere pleasure to take on the role of Editor-in-Chief of Journal SEAMUS, and to announce the publication of Volume 29. I have always been impressed with the quality, rigor, and ingenuity of SEAMUS members' artistic and research endeavors—whether at conferences or in print—and I look forward to building on my predecessor Eli Stine's fabulous work in restructuring the journal, digitizing back issues, and sharing our members' scholarship and research with the SEAMUS community. I am also pleased to share my plans for the journal going forward.

First, and perhaps most obvious, is the switch to digital-only publication. Above all, this shift is intended to increase our readership, and provide new and greater access to our existing readers. This move acknowledges broader changes in the ways we read, create, and distribute research, and aligns us with peer publications who have adapted to the evolving digital landscape along similar lines. It is also my hope that by prioritizing digital publication, we can more closely and richly integrate multimedia, interactive elements, downloads, and other web-based content in future issues.

I am also pleased to report that the SEAMUS Board has approved Journal SEAMUS's transition to open access publishing. Once undertaken, this important step will have an enormous impact on both the accessibility and visibility of the compelling artistic and scholarly work being done in our community. Open access publishing also represents an affirmative step in support of SEAMUS's mission of greater inclusivity as an organization and as a steward of our musical and scholarly communities. The Board and Journal SEAMUS personnel continue to support and prepare for this transition in numerous visible and invisible ways, and we look forward to announcing a timeline soon.

I would like to take this opportunity to express my enormous gratitude to the individuals who have contributed their time and labor to Journal SEAMUS, including Digitization Manager Holli Wittman, Assistant Editors Michael Lukaszuk and Nick Hwang, and Copy Editors Nicholas Cline and Dave Mahony, as well as Board Members Abby Aresty and Becky Brown, Eli Stine, and the SEAMUS leadership.

Longtime readers will have noticed that Journal SEAMUS has experienced a bit of a backlog of late, and I am happy to share that further back issues are forthcoming throughout 2023. I am continually grateful for the extraordinary efforts and talents of our (volunteer) staff, and we all appreciate your patience as we return to normal service. To that end, I would like to invite any interested SEAMUS members to consider submitting your scholarship, articles, reviews, tutorials for the Tips and Tricks section, or other electro-acoustic music-related writing for publication in the journal. We are also interested in submissions of member-created art to be considered for future journal covers. (Inspired by this issue's content, the cover art for Volume 29 was produced with the AI tool DALL•E 2 using keywords from the article titles.) Please send all materials and inquiries to: journal@seamusonline.org

With all that out of the way, I hope you will enjoy this long-awaited double issue of Journal SEAMUS's Volume 29, featuring Caroline Louise Miller's fascinating study of electronic dance music subgenres, Sean Peuquet's critical examination of epistemological limitations in electroacoustic music, an overview of Taylor Brook's AI-powered improvisation system, and a thoughtful reflection on kinesthetic empathy in interactive and networked music by Ryan Ingebritsen, Christopher Knowlton, and John Toenjes. By highlighting the musical, the technical, and the philosophical, this issue truly showcases the extraordinary range of works, practices, and thinking that distinguish the SEAMUS community.

Drake Andersen, Editor-in-Chief

P.S. In my Tips and Tricks article, you will notice that I have made extensive use of in-text hyperlinks in an effort to more fully take advantage of Journal SEAMUS's new digitally native format—a trend I hope will continue in future submissions!

Tips and Tricks

Gesture Mapping with a Neural Network in Max

Drake Andersen
Vassar College
dandersen@vassar.edu

Many of Journal SEAMUS's readers are familiar with machine learning and artificial intelligence tools for music making in one form or another. Indeed, Taylor Brook's excellent article in this volume offers an under-the-hood look at one such sophisticated system. However, I wanted to reach out to those who might be interested in using these tools but are unsure of where to start by offering a brief tutorial using the popular software Max. Click here to download the patch bundle.

While machine learning and artificial intelligence continue to be discussed in the media as tools for creating content—whether ChatGPT, Bard, DALL-E, or MusicLM—this usage is really only scratching the surface of what is possible. Rather than focusing on content creation, I want to discuss how to use the fundamental component of these systems—a neural network—to recognize human gestures in the context of a live electroacoustic performance. The method I describe below is simple by design, prioritizing reinforcing basic concepts through transparency and straightforward coding over efficiency and the use of cutting-edge tools.

Gesture recognition—the interpretation of human gestures by computers—is an important topic for the application of machine learning methods. Of course, what constitutes a "gesture" can vary widely depending on the context. But for the sake of argument, let us imagine a live electroacoustic performance in which the computer recognizes the performer's physical gestures and maps them onto sounds, whether triggering samples or sonic transformations.

Before we get started, I want to emphasize that I am by no means an expert in this area, though I have used and developed machine learning tools in my own musical practice for a number of years. I would like to call out a few resources that have been enormously helpful to me, in the hopes that they might also provide useful context for interested readers. Some of the best conceptual (and mathematical) resources I've come across include Andrew Ng's video series on machine learning and 3Blue1Brown's Youtube channel. I also found it useful to dive into previous research using machine learning for music. For example, the work of David Cope and Roger Dannenberg (here and here) is especially helpful in framing and evaluating research questions pertaining specifically to music. This paper surveying AI methods in algorithmic composition provides a nice overview of the field as well.

The basic idea for this patch is that each gesture will be represented as a list of numbers. When training the neural net, the user will specify a gesture label, which will then be associated with the particular sequence. I will use eight distinct labels, though this is easily modifiable. Over time, the idea is that the neural network can generalize what is distinctive about each gesture. Once the training is complete, the neural network will be able to categorize incoming gestures by assigning them existing labels.

There are many different ways of defining gestures. For example, gestures can be dynamic (moving) or static (not moving). I'm especially interested in dynamic gestures, since their beginnings and endings can be recognized to automatically tell the computer when a gesture is starting or ending. In this patch, I'll be imagining the movement of a hand, fingertip, or other single point in a vertically oriented, two-dimensional plane. (Imagine tracing a shape on a steamed-up bathroom mirror.)

To keep things simple, I will divide up the two-dimensional plane into eight columns and characterize each gesture according to the average values in these columns. The user will sweep their hand back and forth across the plane horizontally, with the variation in vertical position defining each gesture.

To get started, make sure you have the ml.star library by Benjamin Day Smith downloaded in the Package Manager (File → Show Package Manager). We'll use a very simple type of artificial neural network called a multilayer perceptron. A perceptron is a function that attempts to represent the relationship between its input and output by weighting corresponding nodes. While all perceptrons contain an input and output layer (the data to be evaluated and the evaluation, respectively), multilayer perceptrons are unique in that they contain one or more hidden layers that are weighted during the training phase. We'll use a multiplayer perception object called [ml.mlp] included in the ml.star package.

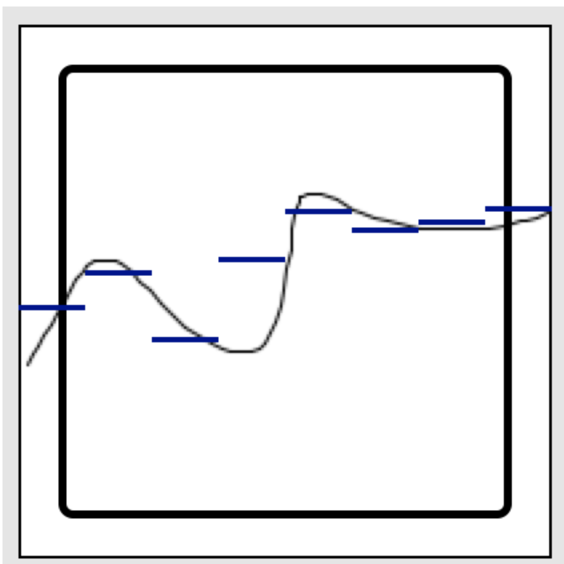

Fig. 1. Visual representation of a gesture

Figure 1 illustrates a typical gesture as represented through the system, with the continuous human-generated gesture given by the black line (an [lcd] object), and the mean-based simplification given by the eight blue bars (a [multislider]).

Gesture recognition characteristically involves a lot of pre-processing because you have to transform human gestures—which are typically complex and time-based—into an input format that a neural network can recognize and work with. In this case, the eight vertical columns are the eight data points that make up the input layer of the neural network. The output layer will have eight points as well, but these will represent the eight possible gesture categories. These categories, it should be emphasized, will be completely user-defined: whatever the user labels "Gesture 1" will become gesture 1, etc. (The fact that I have eight labels for gestures and eight columns of gesture data is entirely coincidental—these do not have to be the same number!)

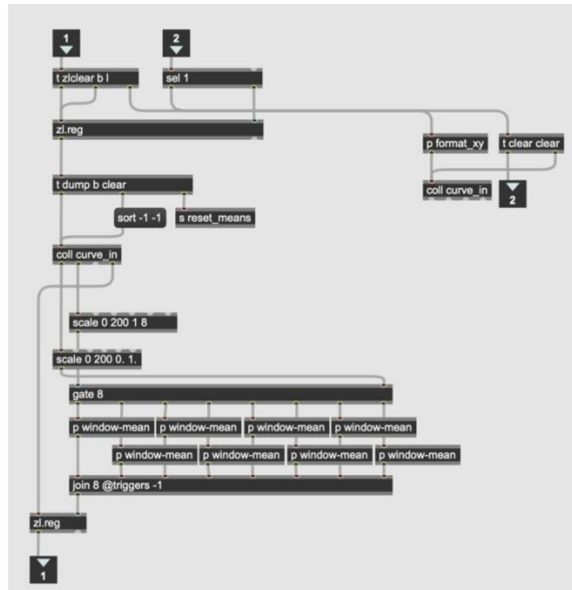

Fig. 2. Pre-processing section of patch

There were a handful of other minor processing details. For example, the [lcd] object used for drawing the line counts y values from the top down. I flipped this by running the xy coordinate list output through a [vexpr] object with a simple expression. I also had to calculate the mean of each segment separately, which involved sorting the values from the "dump" output of the coll object storing the line's coordinates into eight bins of equal size. I ended up solving this by scaling the x values to the range of one to eight and using them as the control for an 8-outlet gate, with each output leading to a separate [mean] object. Figure 2 depicts some of this pre-processing.

Figure 3 offers a view of how the various subunits of the patch are connected. The workflow begins when the user presses "Train" button, which initiates the training phase and specifies its duration (either in epochs [i.e. rounds] or until a minimal error value is reached). Next, the user chooses a training label to associate

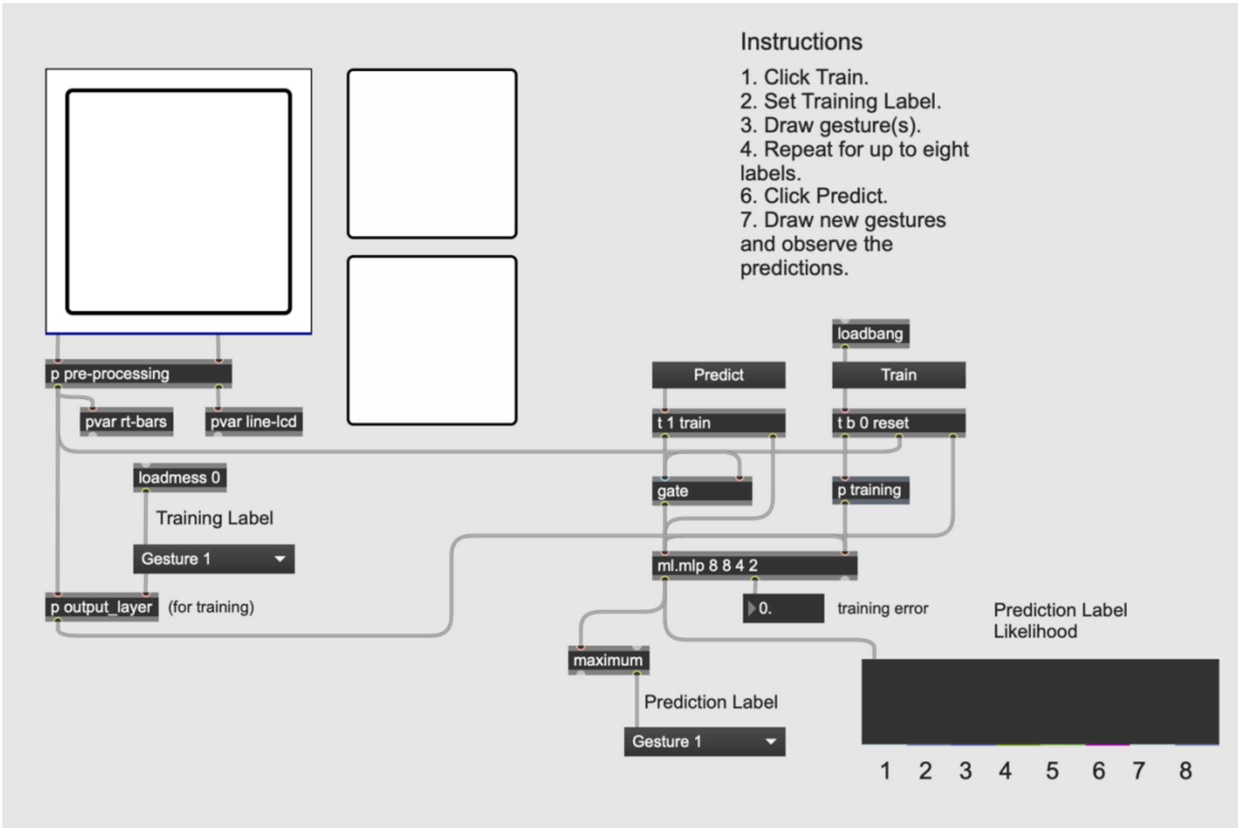

Fig. 3. Overview of patch

with the gesture they are about to draw. After drawing the gesture, they can repeat this process for up to eight labels. Once the user is done training the model, they can click the "Predict" button and draw gestures to see the neural network's prediction. When in prediction mode, the patch gives not only the most likely label for a gesture, but also the likelihood for each label (as a multicolored [multislider] object).

To integrate the patch into a real-world system, simply insert your own source of x-y coordinates into the patch—either into the [lcd] object or directly into the pre-processing sub-patch. A click-based gating system is used to define individual incoming gestures within the pre-processing sub-patch, but you can also populate the "curve_in" [coll] object directly.

The final version of this patch, while functional and easy to use, also has plenty of room for improvement. For example, the inner workings of the neural net remain hidden to the user so as not to clutter the interface, but this also prevents the user from adjusting the inner structure to produce better predictions. The patch would also benefit from some gesture "filtering," by which gestures are not recognized unless they pass across all eight columns. This will become especially important in real-world applications.

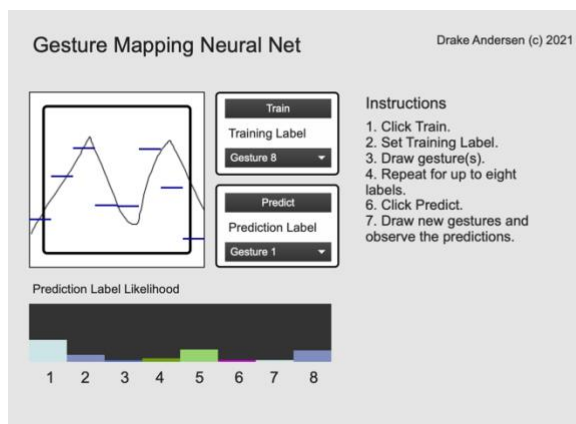

Fig. 4. User interface

Explore the patch bundle yourself. The patch bundle is available for download at: https://drakeandersen.com/wp-content/uploads/2022/01/gesture-_map_-mlp.zip

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

Department of Music
Vassar College
124 Raymond Ave.
Poughkeepsie, NY 12604

www.seamusonline.org