Designing 3D-printed spectrograms for blind students

Nathan Wolek
Stetson University
nwolek@stetson.edu

Grace McEllroy
Stetson University
gmcellroy@stetson.edu

ABSTRACT

Visual representations of recorded sound waves enable computer users to make useful insights into the sound they are hearing. A three-dimensional graph known as a spectrogram can convey how frequency and amplitude are changing through time, providing listeners with opportunities to view with their eyes details that might be missed by their ears alone. However, these opportunities are not accessible to blind persons. While developing lessons about sound for a group of blind and partially-sighted students, we experimented with 3D-printing techniques to produce tactile learning aids based on spectrograms. This paper gathers background research, documents our design process, and records wisdom learned along the way so that others wishing to pursue similar experiments after us will have a head start. A final set of five 3D-printed spectrograms was designed for a lesson about common features of sound events (duration, intensity, pitch, timbre, pattern, speed). We describe the models from this set in detail and provide instructions for downloading the associated digital files so that anyone can 3D print these models for their own projects.

1. INTRODUCTION

Visual representations of recorded sound waves have become an integral part of working with digital audio on a computer. These representations are based on 2D and 3D graphs that convey moment-by-moment changes in sound's intensity and frequency over time [1]. The two most common visual representations used by computer software are known as waveforms and spectrograms. (Although some sources refer to the second item as a sonogram, we will use the term spectrogram in this paper for consistency.) For waveforms, sound is graphed with time on the x-axis and amplitude on the y-axis. Because amplitude is a measurement that can be bipolar, waveforms overall look symmetrical along their horizontal line. For spectrograms, sound is graphed with time on the x-axis, frequency on the y-axis, and intensity on the z-axis. On a 2D rendering of a spectrogram, the intensity z-axis is often transformed into color, with warmer or lighter colors representing high intensity and cooler or darker colors representing low intensity. On a 3D rendering of a spectrogram, the intensity z-axis is translated into height, with the resulting model looking reminiscent of a topographic map.

Most contemporary applications of spectrograms can trace their origins back to The Sound Spectrograph [2], which was developed by Bell Labs to analyze complex sound waves before digital computing was widely available. The 1946 paper about its development describes a device capable of producing visual representations of sound called ”spectrograms” that would capture the ”distribution of energy in both frequency and time”. It is relevant for our current study to also note that those developers tried producing solid models as a pre-cursor to their final product. Their process for creating these solid models involved graphing the amplitude output of 200 overlapping band-pass filters, then cutting out these graphs and stacking them side by side. However, the process was so labor intensive, they reported that ”production of solid models, while useful for particular purposes, is hardly a practical method for everyday needs.” Instead they turned their attention to printing these variations in shades of grey, producing spectrograms that we would recognize as very similar to those in use today.

From this work, applications in electroacoustic music and bioacoustics would develop in parallel. In electroacoustic music, Waters and Ungvary would first propose using spectrograms as a technique for visualizing the structure of compositions [3]. This tool for analysis has since become widespread and features prominently in many publications on the topic, including books edited by Licata [4] and Simoni [5]. In bioacoustics, researchers make use of spectrograms to visualize the structure of individual animal sounds [6]. Bernie Krause has frequently used spectrograms as tool to explain how overlapping elements of a natural soundscape fit together, or what he calls the acoustic niche hypothesis [7]. Even texts about birdsong for non-expert enthusiasts make frequent use of spectrograms to explain distinctions between species [8].

The benefits of having visual representations of recorded sound are many, but these benefits are inaccessible to one key population: blind persons. The two authors of this paper are currently involved with Young Sound Seekers [9], an environmental arts program created specifically for youth with visual impairments. While preparing lessons about sound for blind students, we were forced to confront the role of waveforms and spectrograms within contemporary teaching about sound art and bioacoustics. While we agree that one can certainly teach these topics without visualizations, they do provide a perspective on sound that enables analysis from outside of time and the moment-by-moment experience. We wondered if 3D printing technology could be used to develop a set of tactile spectrograms that would help us explain key features of and differences between individual sound events. By translating these visualizations into a tangible form, we hoped to develop tools that could be used by anyone to teach lessons about sound. This paper details our design process, the features of our final set of models, and plans for future work to understand their impact on learning. We begin by documenting the existing work that informed, inspired, and enabled us to start our design process.

2. BACKGROUND

2.1 Artistic applications

Some artists have seen the potential to turn visual representations of sound waves into tangible, sculptural objects. Carsten Nicolai is a German artist and musician whose work spans interactive installations and film scoring projects. In 2008, he created small twin aluminum sculptures for *yes/no*, each one based on the waveform for a single word spoken by another artist, Laurie Anderson [10]. Instead of rendering the waveform as a flat 2D cutout, Nicolai chose to spin the waveform along the time x-axis to produce a 100-centimeter-long cylindrical shape. But if you are familiar with reading waveforms, several features are instantly recognizable. The bipolar pattern of amplitude extending up and down from the center is quite clear. In addition, there are clear ripples for the pitched vowels in these words, while the noisier consonants form denser sections. He later followed this piece with two others that use similar materials, methods, and form: *zukunftsangst* [11] and *sekundenschlaf* [12].

Gilles Azzaro is a French digital artist who uses a variety of 3D printing and digital fabrication methods to transform recordings of the human voice into sculptural objects. In 2013, he used a recording from the U.S. president's State of the Union to create *Barack Obama: Next Industrial Revolution* [13]. The source was an excerpt from President Obama's 2013 State of the Union Address that made explicit reference to 3D printing as a tool that could revolutionize manufacturing [14]. Azzaro converted the audio from this 39-second excerpt into a spectrogram, then used a 3D printer to render the spectrogram into a 151-centimeter-long sculpture. The work was encased in a clear tube. The end of the tube contains a button marked start, inviting viewers to activate the sculpture. When pressed, a small laser sweeps across the surface of the spectrogram sculpture while the original audio excerpt from President Obama's speech is heard. This simple gesture helps viewers to understand the connection between the sculpture and its source. Azzaro has since produced other sculptures that follow a similar format, including one based on Neil Armstrong's famous moon landing quote [15].

2.2 Learning applications

In a standard classroom, manipulatives and other tools have been used to further students' understanding of abstract topics. Manipulatives (also known as realia) are objects that can be touched, molded, moved, or otherwise 'manipulated.' Common examples include number cubes in math [16], or anatomical models in the sciences [17]. Other tools in standard classrooms used by teachers to aid students include visuals, such as graphs or posters.

In a class with students with visual impairments, there might be a greater reliance on tactile manipulatives and technologies than visuals. Manipulatives include models with raised textures (such as 3D-printed maps) [18], tactile pictures in books, and models with braille labels [19]. Some technologies might include using braille or a related device, audio recordings, and text-to-speech software. Teachers implementing these manipulatives pointed out the importance of orienting the object for the student and connecting the model to real life [20].

Several studies found that when compared to a control group, groups using manipulatives showed an increase in understanding and test scores [17, 18, 21]. One study involving paper models, found that preserving the 3D aspect improved students' overall knowledge [17].

2.3 3D printing

References to 3D printing at ICMC have been relatively few between the years 2010 and 2018, with only 10 papers using the phrase "3D printing" and 3 more using the phrase "digital fabrication". Most of these focus on applications of printing novel instruments [22], speaker enclosures [23], or robot components [24]. During this same period, there were no references to using 3D printing technology to convert visual representations of recorded sound into a tactile form for the purposes of teaching. It seems as if this application of 3D printing has been overlooked. We believe it has great potential for teaching blind and visually-impaired students, enabling them to access information that has previously been obscured when studying recorded sound. Given the ICMA's recent conversations about diversity and accessibility issues [25], sharing information about how we might leverage 3D printing to improve accessibility is incredibly timely.

3. DESIGN PROCESS

3.1 Initial tests

In summer of 2020, we began searching for existing tools to create tactile representations of audio recordings. We soon learned that tools for converting data into the necessary files for 3D printing are now as close as your web browser. One such tool is a small, open source website called 3Dprintedsound.com [26], which quickly distinguished itself from the other options. Visitors to this website can either record sound using their built-in microphone or upload a pre-existing recording in WAV format. Once a sound has been added, users have the option to render a 3D model in one of three formats: flat waveform, circular waveform, or spectrogram. Users have several parameters that control the final results including FFT size, minimum decibels, maximum decibels, and smoothing amount. A 3D visualization of the model is presented within the browser and the mouse can be used to interactively rotate and zoom in on the model in real-time. Once users are happy with the results, they need to click a virtual button that downloads an STL file, a format commonly used in rapid prototyping and 3D printing [27].

3Dprintedsound.com became an essential tool for our iterative design process. Given the environmental focus of

Young Sound Seekers, we limited ourselves to field recordings as our source audio early on. We ran through trials 0 through 4 in quick succession to test various materials, forms, and lengths. One need we could foresee early on was how to ensure students would orient the models correctly. To address this we used a parameter called 'offset' to add extra thickness to the bottom, which provided space for raised text on the front of the model. We used this feature to add titles of the models in braille and traditional fonts. Through these initial trials, we made some fundamental decisions about the uniformity of the future models. First, we committed to using ABS plastic, as it is more resistant to heat damage. Second, we found that a ratio of 30 seconds of audio to 6 inches (15.3 cm) length of the model provided a good balance of detail and portability. To achieve this self-imposed standard from trial 5 onward, we used a separate program called Cura [28], which allowed us to apply some post-processing to the STL models and manipulate them to our desired proportions.

3.2 First consultation

In July 2020, we reached out to Mari, a visually-impaired adult in the local community, to evaluate the usefulness and physicality (size, complexity, color) of our trial 5 prototypes. Because pandemic restrictions prevented us from meeting in person, we mailed models to her and scheduled phone conversations to collect her input. She was our first beta-tester and helped us work on explaining the three dimensions to students unfamiliar with audio terms. Mari realized some of the peaks on the models were too fragile for tactile use, and had fallen off. Therefore, we needed to reevaluate their durability and the amount of detail we wanted to capture in print.

3.3 First field test

In December 2020, we took a few 3D-printed spectrograms to a Young Sound Seekers event for the first time. Students took turns silently feeling their way across a trial 6 model while we played the associated recording on a small speaker. Their response was very positive, but we also made some important observations. First and foremost, a single model took almost 20 minutes to travel through the group. In our current pandemic environment, it also carried with it the further complication of needing each person to use hand sanitizer before and after handling the model. We knew we would need multiple copies of a model to make it an effective part of any future lesson. Second, we realized that the printed words and braille on the front were more of a distraction than help. Only one or two students could read the braille option, while the texture of the raised words competed with the spectrogram for attention. Therefore, we decided this feature could be abandoned or simplified. Finally, it was during this test that we noticed that models did not match the usual direction of reading spectrograms. With the words facing the student, frequency decreased from foreground to background, thus putting the lowest frequencies at the far end of the y-axis.

3.4 Further refinements

In summer of 2021, we began to explore the potential of synthesized sound as an additional source for our 3D-printed spectrograms. By using simple sine wave and white noise generators, it was possible to generate sound examples that explicitly articulate specific spatial features of the tactile models. All synthesized sound was created using Ableton Live Suite 11 and its Operator device [29], primarily with the Sine Noise Attack Lead preset.

Achieving the desired results required us to take the raw STL files produced by 3Dprintedsound.com and apply some specific post-processing steps in Cura. First, we increased the FFT size on 3Dprintedsound.com to 1024 to provide more detail in the frequency dimension, but this resulted in raw STL models with approximately four to one ratio between frequency's y-axis and time's x-axis. Our solution was to use Cura to slice the models at about 8 kHz so that the models were more square. Second, we addressed the earlier concern of about directions being backward by using Cura to independently flip the y-axis. This provided models that matched the usual direction of reading spectrograms, with time increasing from left to right and frequency increasing from foreground to background. Finally, we abandoned the use of text and instead inserted a raised arrow on the front left corner of each model using Cura. This feature proved enough for our students to quickly orient each model correctly in their hands. The final dimensions of the models ended up being 5.5 inches (13.7 cm) from left to right for 30 seconds of time, and 6 inches (15.3 cm) from front to back for 8000 Hertz of frequency range.

3.5 Pilot study

The final set of models and sound recordings were designed for a November 2021 lesson about describing the features of individual sound events. Reduced listening [30] is a challenging topic to teach, so the potential for tactile learning aids that would reinforce specific concepts was exciting. The features taught as part of the lesson were: duration, intensity, pitch, timbre, pattern, speed. Together, these form the "DIP-TiPS", a mneumonic device that the students chose from several options to make the list easier to remember. In addition, we wanted to overcome the time limitations of having a single model travel through the group. For this reason, our goal was to design a lesson around eight sets of five distinct models. Each of the 40 models can take up to 10 hours to print, so with an estimated 400 hours of print time, we had to start the printing process weeks in advance. This actually allowed us to stagger the final design of each model. While multiple copies of a specific model were being printed, we were free to continue tweaking of the design of models waiting for their turn to be printed. Finally, we decide to make each model in the set a unique color so that the partially-sighted students could quickly identify the right model for each step in the lesson.

4. FINAL MODELS

The first model focuses on the first two items in our list of features, duration and intensity (see Figure 1). A sine wave oscillator is used to create a series of sound events with a consistent pitch and combined with a short burst of noise on the attack that adds emphasis to the beginning of each event. The consistent pitch means that these sound events

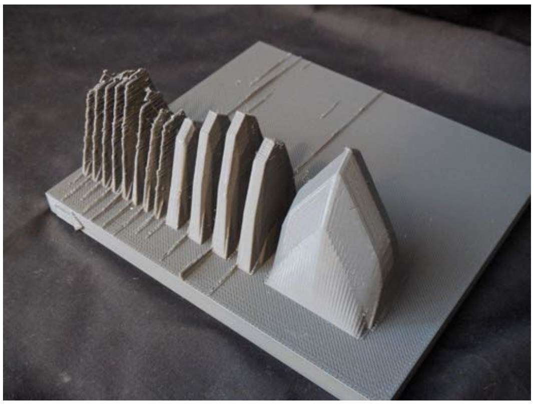

Figure 1. The duration-intensity model.

form a horizontal line of tall peaks across the foreground of the model, with the noisy attack creating vertical ridge at the left edge of each event. The events are designed to form three groups, with each cluster having a consistent duration. The eight events in the first group have the shortest duration, so their peaks form thin triangles that are flexible and actually rather fragile. The gaps between them only measure about a half centimeter, so they are not wide enough to fit fingers between the peaks. Their intensity gradually increases from the first to the fifth peak, which is tallest of the group, then recedes to the eighth peak in the group. The four events in the second group are twice the duration of the events in the first group, so they are thicker and more stable. They feature the same rising and falling pattern of intensity changes, with the third peak in the group of four being the tallest. Here the gap between these four peaks is almost a centimeter, so it may be possible for smaller fingers to fit between the peaks. The final cluster is actually a single sustained tone with a duration equal to the earlier groups that crescendos and decrescendos over its duration. Because of its longer duration, this event forms the largest and most stable peak on the entire model. The shape is somewhat reminiscent of an arrowhead, with the tip capturing the moment of highest intensity for this sound event in the right corner of the model's foreground.

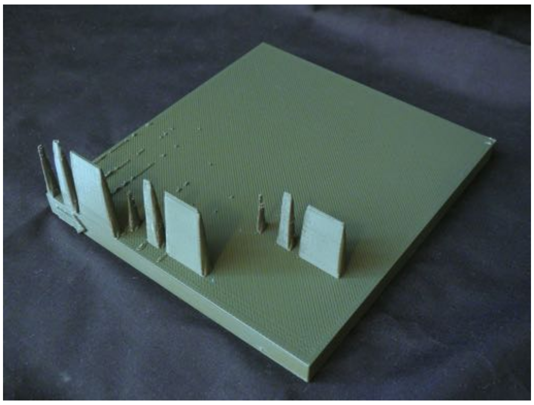

Figure 2. The duration-pitch model.

The second model concentrates on duration and pitch (see Figure 2). The sine oscillator is used here to articulate a single rhythmic pattern three times, with each rhythmic group repeating the same short-medium-long pattern. The first peak in the pattern is about the diameter of a toothpick. This is followed by a second peak that is about a half centimeter wide and then a third peak that is just less than two centimeters wide. Each group articulates this pattern on a single pitch. The lowest pitch is found in the left corner of the foreground, with the next two iterations of the pattern increasing two octaves from the previous one. This results in three discontinuous horizontal ridges that gradually get farther from the foreground. Because the frequency scaling of the spectrogram is linear and these changes in pitch are exponential, the distance from the foreground is noticeably different for the final group. When this model is used while listening to the two-octave changes in the source recording, it can provide a tangible demonstration of how the spaces between octaves are not linear.

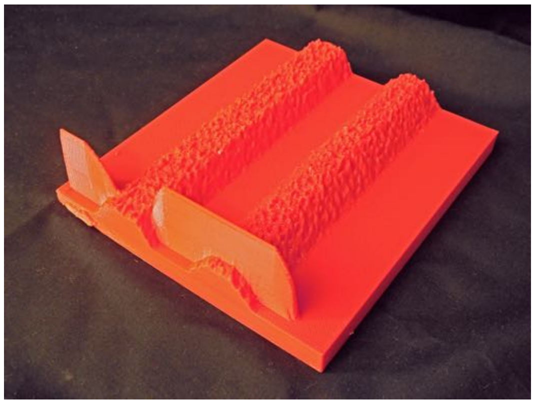

Figure 3. The timbre model.

The third model is designed to focus on timbre and demonstrate differences between tonal or noisy sound events (see Figure 3). It does this by cross-fading between a sine wave oscillator and white noise generator. Because of its low pitch, the sine wave oscillator forms a thin, smooth, horizontal ridge near the front of the spectrogram. This contrasts easily with the wide, bumpy, vertical ridges created by the white noise fading in and then out. The bumpy texture of the white noise provides a clear contrast with the flat background of the digital silence in this spectrogram and the smooth shapes found in the other models. Although they sound perceptually about the same intensity, the sine wave's horizontal ridge is almost twice as tall as the white noise's vertical ridge. This feature helps to demonstrate that more focused pitches need to be more intense to balance with broadband noise. Because the sine wave fades out before the first white noise entrance and fades back in as it exits, there is a gap in its ridge. This contrasts with the second entrance of the white noise where the sine wave persists, causing the taller, pitched, horizontal ridge to seemingly slice through the shorter, noisy, vertical ridge.

The fourth model demonstrates changes in both speed and patterns (see Figure 4). The sine wave oscillator is used again to create brief sound events that individually form peaks on the model similar to the shape of a toothpick, and together these successive events form a horizontal line near the front of the model. The events have a pitch pattern that alternates between tones that are a half step apart, repeating a high-low-high pattern that is often used

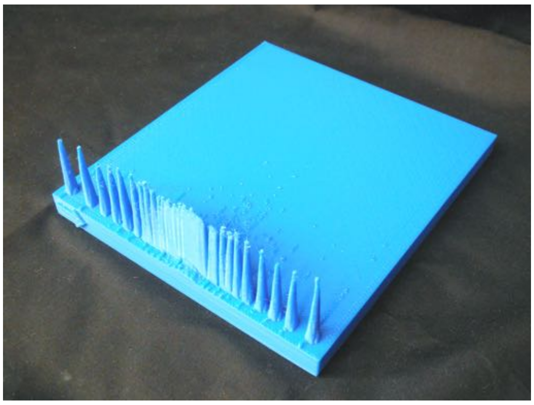

Figure 4. The speed-pattern model.

in demonstrations of auditory scene analysis [31]. The distance between pitches is not enough to create vertical space between events, but it does create a subtle wobble when running fingers horizontally across the line of events. The speed of their repetition begins slow, increases to its fastest pace at the mid-point of the model, then decelerates back to their original pace before the end. This change in their pacing creates gaps between events at the left and right extremes of the line, but the peaks in the middle are so close together that they essentially fuse together. The rhythmic pattern was designed to be irregular, something that makes the gaps irregular where they are present. This contrasts with the regular pitch pattern heard in the model, as well as a contrast with the more regular rhythmic patterns in models one and two.

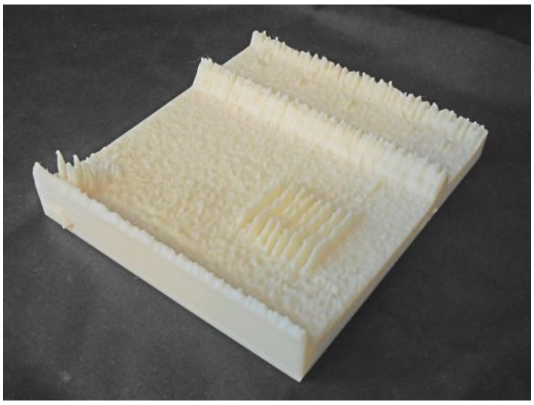

Figure 5. The cricket-owl-frog model.

The final model was based on a short field recording that we had used throughout the design process (see Figure 5). Recorded overnight inside the Canaveral National Seashore on 21 May 2020, it features crickets, an owl, and a frog, all within the span of 30 seconds. The differences in the pitch, timbre, and pattern of these sound events make the recording a good spectrogram for application toward the end of the lesson. The cricket stridulations have an overlapping pattern that merges together into a band at the highest pitch among the three sound events, forming two ridges near the top of the spectrogram. The owl has the lowest pitch, with its sound event captured in three peaks near the left corner of the foreground. Because the speed of the owl's vocalizations increases, the middle peak is clearly closer to the last than the first. The frog vocalizations have a pitch that is between the owl and the cricket sound events and sit in the right half of the spectrogram. Because these frog vocalizations have a regular pattern and a timbre that is noisier than a single pitch, they form a series of vertical ridges about 4 cm long. Finally, because this is an acoustic recording instead of synthesized source, the extra noise gives the spectrogram's background a bumpy texture that was missing from the other models.

5. FUTURE WORK

The students who participate in Young Sound Seekers programming had a consistently positive response to our 3D-printed spectrograms. Each time we brought them models during the design process, we would hear comments about how "cool" it was to have a tactile representation of sound. While this subjective feedback is gratifying, we are already working to obtain more objective feedback about the impact our models have on learning. The DIP-TiPS lesson presented in November 2021 served as the pilot study for an experimental procedure designed to measure the impact of these 3D-printed spectrograms through pre- and post-tests. During the next few months, we already have plans to refine the procedure and run the study with both sighted and visually-impaired participants.

We would also like to invite others to use these models in their own teaching and research. To encourage the use of these models by others, we are making the necessary files available for download so that others can 3D print copies for their own projects. For each model, the STL files, associated audio recordings, photos, and text descriptions can be obtained under an MIT License from the following address:

https://github.com/nwolek/3d-printed-spectrograms

Acknowledgments

This project was developed as part of Young Sound Seekers, a partnership between Atlantic Center for the Arts and Stetson University that develops programming for blind and partially sighted youth at Canaveral National Seashore in Florida USA. Young Sound Seekers receives financial and logistical support from the National Park Service through its Natural Sounds and Night Skies Division. Early stages of this research were supported by the Stetson Undergraduate Research Experience (SURE) program. Throughout the development process, 3D printing services were provided by the duPont-Ball Library's Innovation Lab.

6. REFERENCES

[1] N. Adams, "Visualization of Musical Signals," in *Analytical Methods of Electroacoustic Music*, M. Simoni, Ed. New York, New York: Routledge, 2006, pp. 13–28.

[2] W. Koenig, H. Dunn, and L. Lacy, "The Sound Spectrograph," *The Journal of the Acoustical Society of America*, vol. 18, no. 1, pp. 19–49, 1946.

[3] S. Waters and T. Ungvary, "The Sonogram: A tool for visual documentation of musical structure," in *Proceedings of the 1990 International Computer Music*

Conference, Univ. of Glasgow, Scotland, 1990, pp. 159–162.

[4] T. Licata, Ed., *Electroacoustic Music - Analytical Perspectives*. Westport, Connecticut USA: Greenwood Press, 2002.

[5] M. Simoni, Ed., *Analytical Methods of Electroacoustic Music*. New York, NY, USA: Routledge, 2006.

[6] C. Elemans, K. Heeck, and M. Muller, "Spectrogram analysis of animal sound production," *Bioacoustics : the International Journal of Animal Sound and its Recording 18 (2008) 2*, vol. 18, Jan. 2008.

[7] B. Krause, "The Niche Hypothesis: A virtual symphony of animal sounds, the origins of musical expression and the health of habitats," *The Soundscape Newsletter*, no. 06, Jun. 1993. [Online]. Available: https://www.researchgate.net/publication/295609070_The_niche_hypothesis

[8] A. Thomas, *RSPB Guide to Birdsong*, ser. RSPB. London UK: Bloomsbury Publishing, 2019.

[9] Atlantic Center for the Arts, "Young Sound Seekers." [Online]. Available: https://atlanticcenterforthearts.org/youngsoundseekers/

[10] C. Nicolai, "yes/no," 2008. [Online]. Available: http://www.carstennicolai.de/?c=works&w=yes_no

[11] ——, "zukunftsangst," 2015. [Online]. Available: https://www.carstennicolai.de/?c=works&w=zukunftsangst

[12] ——, "sekundenschlaf," 2018. [Online]. Available: https://www.carstennicolai.de/?c=works&w=sekundenschlaf

[13] G. Azzaro, "BARACK OBAMA : NEXT INDUSTRIAL REVOLUTION," 2013. [Online]. Available: http://www.gillesazzaro.com/pages/en/printing3D.htm

[14] B. Obama, "Remarks by the President in the State of the Union Address," U.S. Capitol, Feb. 2013. [Online]. Available: https://obamawhitehouse.archives.gov/the-press-office/2013/02/12/remarks-president-state-union-address

[15] G. Azzaro, "NEIL ARMSTRONG," 2019. [Online]. Available: http://www.gillesazzaro.com/pages/en/armstrong.htm

[16] J. P. Jones and M. Tiller, "Using Concrete Manipulatives in Mathematical Instruction." *Dimensions of Early Childhood*, vol. 45, no. 1, pp. 18 – 23, 2017.

[17] S. J. Rehorek, P. G. Falso, and J. R. Siebert, "Using Paper Models to Teach Basic Concepts of the Human Musculoskeletal System." *HAPS Educator*, vol. 23, no. 3, pp. 488 – 498, 2019.

[18] M. E. Brittell, A. K. Lobben, and M. M. Lawrence, "Usability Evaluation of Tactile Map Symbols across Three Production Technologies." *Journal of Visual Impairment & Blindness*, vol. 112, no. 6, pp. 745 – 758, 2018.

[19] J. Martin, "Teaching Basic Cryptography Concepts Using Braille and Large Print Manipulatives." *Journal of Science Education for Students with Disabilities*, vol. 22, no. 1, 2019. [Online]. Available: https://scholarworks.rit.edu/jsesd/vol22/iss1/6

[20] L. P. Rosenblum, L. Cheng, and C. R. Beal, "Teachers of Students with Visual Impairments Share Experiences and Advice for Supporting Students in Understanding Graphics." *Journal of Visual Impairment & Blindness*, vol. 112, no. 5, pp. 475 – 487, 2018.

[21] K. Larkin, "Mathematics Education and Manipulatives: Which, When, How?." *Australian Primary Mathematics Classroom*, vol. 21, no. 1, pp. 12 – 17, 2016.

[22] S. Barrass and T. Barrass, "CYMATIC SYNTHESIS OF A SERIES OF BELLS," in *Proceedings of the 2013 International Computer Music Conference*, Perth, Australia, 2013, pp. 18–23.

[23] A. R. Jensenius, K. Glette, R. I. Godøy, M. Høvin, K. Nymoen, S. A. Skogstad, and J. Torresen, "FOURMS, UNIVERSITY OF OSLO – LAB REPORT," in *Proceedings of the 2010 International Computer Music Conference*, New York, USA, 2010, pp. 290–293.

[24] J. Murphy, A. Kapur, and D. A. Carnegie, "SWIVEL: ANALYSIS AND SYSTEMS OVERVIEW OF A NEW ROBOTIC GUITAR," in *Proceedings of the 2013 International Computer Music Conference*, Perth, Australia, 2013, pp. 445–448.

[25] N. Wolek and A. Slater, "Can we align our research and shared values to improve accessibility?" *array: The journal of the ICMA*, pp. 27–32, 2021. [Online]. Available: https://journals.qucosa.de/array/article/view/3271/3090

[26] T. Toplak, "3DprintedSound," Mar. 2019. [Online]. Available: https://github.com/TimToplak/3DprintedSound

[27] Library of Congress, "STL (STereoLithography) File Format Family," Sep. 2020. [Online]. Available: https://www.loc.gov/preservation/digital/formats/fdd/fdd000504.shtml

[28] Ultimaker, "Cura," Nov. 2021. [Online]. Available: https://github.com/Ultimaker/Cura

[29] Ableton Inc., "Live Instrument Reference: Operator." [Online]. Available: https://www.ableton.com/en/manual/live-instrument-reference/#26-6-operator

[30] P. Schaeffer, "Acousmatics," in *Audio Culture: Readings in Modern Music*, 1st ed., C. Cox and D. Warner, Eds. New York, New York: Continuum, 2004, pp. 76–81, translated from French by Daniel W. Smith; original publication 1966.

[31] A. Bregman, "Audio demonstrations of auditory scene analysis." [Online]. Available: http://webpages.mcgill.ca/staff/Group2/abregm1/web/downloadsintro.htm