

THE SONIFICATION OF SOLAR HARMONICS (SoSH) PROJECT

Seth Shafer

Timothy Larson

Elaine diFalco

University of Nebraska at Omaha
6001 Dodge Street
Omaha, NE, 68182
sethshafer@unomaha.edu

tplarson@sun.stanford.edu

University of North Texas
1155 Union Circle
Denton, TX 76203
elainedifalco@my.unt.edu

ABSTRACT

The Sun is a resonant cavity for very low frequency acoustic waves, and just like a musical instrument, it supports a number of oscillation modes, also commonly known as harmonics. We are able to observe these harmonics by looking at how the Sun's surface oscillates in response to them. Although this data has been studied scientifically for decades, it has only rarely been sonified. The Sonification of Solar Harmonics (SoSH) Project seeks to sonify data related to the field of helioseismology and distribute tools for others to do the same. Creative applications of this research by the authors include musical compositions, installation artwork, a short documentary, and a full-dome planetarium experience.

1. INTRODUCTION

It is a poignant coincidence that acoustical physics is such an intrinsic part of our most prominent celestial object when so much of Western philosophical history connects the cosmos to sound. Our home star appears as a mass of boiling plasma, and it rings like a bell in a sandstorm, generating millions of resonant harmonic modes [1]. By applying the mathematics of spherical harmonics and fluid dynamics we are able to determine various properties of the Sun's internal structure. Although the data used is acoustic in nature, scientists have only very rarely listened to it, despite the fact that sonification of other types of solar data has yielded new scientific insights [2].

After a short introduction to the subject of helioseismology, we describe a collaborative research initiative called the Sonification of Solar Harmonics (SoSH) Project. This project seeks to transform helioseismology into a listening experience for the scientist and nonscientist alike. One of the initial outcomes from the SoSH Project is a software tool for rendering helioseismic data as audible sound. This tool is the most advanced contribution to helioseismic sonification to date and provides the first access to audio generated from the most recent data available. In addition to the SoSH Tool, we describe several creative applications derived from solar harmonics research and future directions for the project. The contents described in this paper including the SoSH Tool, datasets, and additional information can be found at <http://solar-center.stanford.edu/SoSH/>.



This work is licensed under Creative Commons Attribution – Non Commercial 4.0 International License. The full terms of the License are available at <http://creativecommons.org/licenses/by-nc/4.0>

2. A BRIEF PRIMER ON HELIOSEISMOLOGY

The study of oscillations inside the Sun is called helioseismology. In particular, here we shall consider acoustic waves. Turbulent convection near the solar surface excites sound waves, and the waves with frequencies that resonate form the harmonics. Just as the frequency of a plucked guitar string becomes higher with greater tension and lower with greater thickness, the frequencies of the Sun's harmonics enable us to infer properties of the solar interior such as its pressure and density. And just as any acoustic instrument produces a set of harmonics above a fundamental frequency that combine to create a characteristic timbre, so too does the Sun.

2.1. Spherical Harmonics

The input data for helioseismology are typically velocity images of the Sun, where each pixel gives the speed of that plasma element toward or away from the observer. It is a mathematical theorem that any such image of the Sun's surface can be expressed as a sum over spherical harmonics, which are simply a special set of functions of latitude and longitude. Each of these functions are labeled by two integers: the spherical harmonic degree l and the azimuthal order m . The degree l is ≥ 0 , and for each l , there are $2l + 1$ values of m , ranging from $-l$ to l .

One way to understand spherical harmonics is in terms of their node lines, which are the places on the sphere where the spherical harmonics have an amplitude of zero. The degree l tells how many of these node lines there are in total, and the absolute value of the order $|m|$ gives the number in longitude, so the number of node lines in latitude is $l - |m|$. Therefore a spherical harmonic with $m = 0$ has only latitudinal bands, while one with $m = l$ has only sections like an orange. A third integer, the radial order n , tells how many nodes the oscillation has along the Sun's radius. Since only the surface of the Sun is visible to us, all the values of n are present in each spherical harmonic labeled by l and m , although only some of them will be excited to any appreciable amplitude. The total mode, then, is represented as a product of a spherical harmonic and another function of radius, known as the radial eigenfunction. The radial eigenfunction depends on both n and l .

Figure 1 below illustrates modes with degree $l = 5$ and all nonnegative values of m . Modes with $m < 0$ are not included because in a still image they are indistinguishable from modes with $m > 0$. As the spherical harmonics evolve in time, one would see the two signs of m rotate in opposite directions; this is discussed further below. One can see that different spherical harmonics sam-

ple different latitudes, according to the value of $|m|$ relative to the degree l . Modes with high absolute values of m have their maximum amplitude at low latitudes, whereas lower values extend to higher latitudes.

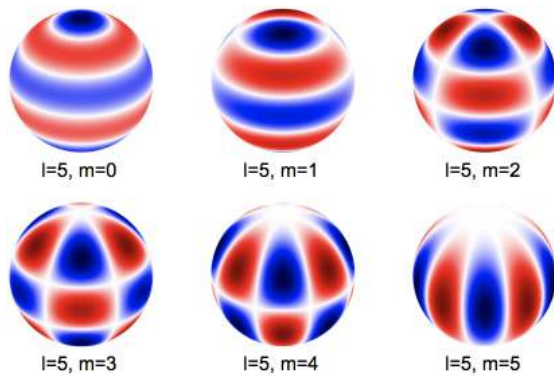


Figure 1: Surface views of the Sun showing all harmonic modes with degree $l = 5$.

2.2. Modeling the Sun’s Interior

The harmonics we are able to measure have frequencies ranging from 1000 to 5000 microhertz. For any given value of the degree l , we will find a certain range of values of the radial order n , with frequency increasing as n increases. The next two figures show interior views of the Sun for two different radial orders and degree $l = 5$. For clarity, in Figure 2 we show radial order $n = 3$, although this mode is expected to have a frequency too low to measure. Figure 3 shows radial order $n = 20$, which is easily measured, but at this frequency the nodes along the radius become so closely spaced near the surface that they are difficult to discern at this scale. It is important to realize that we are seeing modes of many different n in each spherical harmonic. For instance, for degree 5 we might measure modes with n ranging from 7 to 28, each oscillating at its own frequency, roughly 140 microhertz apart from each other.




Figure 2: Interior views of the Sun showing harmonic modes with degree $l = 5$ and radial order $n = 3$. A model predicts these modes to have a frequency of 800 microhertz.

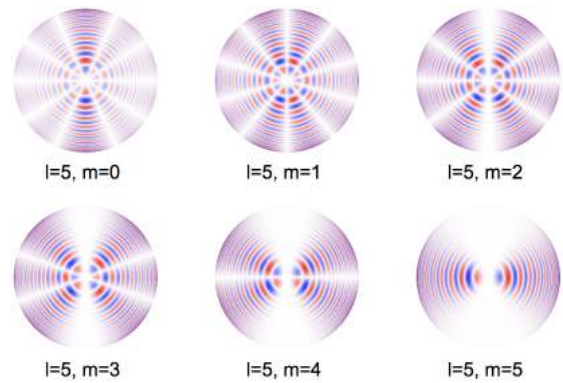


Figure 3: Interior views of the Sun showing harmonic modes with degree $l = 5$ and radial order $n = 20$. A model predicts these modes to have a frequency of 3190 microhertz.

Each mode oscillates with its characteristic frequency, and each samples different depths inside the Sun. At a given degree, high frequencies will penetrate more deeply, while low frequencies are trapped closer to the surface. Likewise, at a given frequency, high values of the degree l will be trapped near the surface, while low values will penetrate almost all the way to the core, with $l = 0$ even reaching the center. This is further illustrated in the next series of figures below, which show a selection of modes with degree $l = 25$ for the same two values of n (see Figure 4). We see that the extent in latitude of the spherical harmonic combines with the extent in radius of the radial eigenfunction to yield a mode that inhabits a particular region of the Sun.




Figure 4: Surface views for degree $l = 25$ and corresponding interior views for $n = 3$ and $n = 20$. The expected frequencies are 1290 and 3990 microhertz.

2.3. Typical Data Pipeline

To determine the frequencies of the Sun’s harmonics, a typical instrument might take an image once a minute for 72 days. For each image, we decompose it into its various spherical harmonic components. For each of these components, we form a timeseries of its amplitude. From the timeseries we are able to construct the power spectrum (acoustic power as a function of frequency). It is here that we are able to separate the two signs of m . Without delving into complex analysis and the theory of the Fourier transform, we simply state the two signs correspond to the positive and negative

frequency parts of the power spectrum. Furthermore, the sign of m that rotates in the same direction as the Sun will be shifted up in frequency, while the sign that goes against solar rotation will be shifted down in frequency. Because different modes sample different regions of the Sun, we are able to use their frequencies to determine solar properties as a function of both depth and latitude. For example, we are able to use the frequency splitting in m to measure internal solar rotation.

Once the power spectrum is calculated, the location of peaks will correspond to the frequencies of the modes. The height of each peak tells us the mode amplitude, and the width of the peak tells us how much the oscillation is damped. Each n will have its own peak in the power spectrum.

The relationship between l , n , and frequency is illustrated in Figure 5 below. Plotted on the left is raw power as a function of l and frequency for $m = 0$. As one can see, the modes for a given n form a ridge of power. The right panel is a scatter plot indicating which modes we have been able to successfully fit. In both plots the bottom ridge corresponds to $n = 0$.




Figure 5: Degree-frequency diagrams. Left panel shows raw power, right panel shows the modes that we are able to fit.

3. OVERVIEW OF THE SoSH PROJECT

3.1. Previous Sonification Efforts

The Sonification of Solar Harmonics (SoSH) Project is the most extensive effort toward the translation of helioseismic data into audible sound, but it is not the first. That distinction goes to Douglas Gough, who created an audio tape for public demonstration using solar data in the early 1980's [3]. More recently, Alexander Kosovichev sonified 40 days of data taken with an instrument in space [4]. Although he had the benefit of computer processing, still only a handful of audio samples were created. Nonetheless, until preliminary investigations by Larson and the current project, these few audio files were the only sonifications of helioseismic data available to the public, and they were used in a wide range of other works. [5]. Although Larson sonified data covering an entire solar cycle, it was not until the development of the SoSH Tool that one became able to sonify helioseismic data interactively.

For completeness, we note that other types of solar data have been sonified over the years, for both artistic and scientific purposes. Examples include the use of solar cycle data by Thorbjørn Lausten [6] and solar radio emissions by Thomas Ashcraft [7]. Several projects have sonified various solar wind data, notably Robert Alexander working with the University of Michigan to use ACE data [8], a UC Berkeley group using STEREO data [9], and Don Gurnett at the University of Iowa using data from the Plasma Wave Instrument onboard Voyager [10]. Finally, Chris Hayward [11] and Florian Dombois [12] conducted work similar to this project to audify geoseismic data.

3.2. Datasets for Sonification

The SoSH Project aims to sonify any solar harmonics dataset. Several such datasets are available for sonification. The first of these comes from the Michelson Doppler Imager (MDI) [13] onboard the Solar and Heliospheric Observatory (SoHO). It was in operation from May 1996 to April 2011 and was the source of data for both the audio files made by Kosovichev and the earlier work of Larson. Similar to MDI and parallel to instruments operating in space, the Global Oscillation Network Group (GONG) has operated a network of six ground-based observatories since 1995 [14]. Both GONG and MDI observe the same spectral line at a cadence of one minute.

In 2010 MDI was superseded by the Helioseismic and Magnetic Imager (HMI) [15] onboard the Solar and Heliospheric Observatory (SDO). HMI observes a different spectral line at a cadence of 45 seconds and remains in operation today. The SoSH Project represents the only sonification of HMI data to date. Although capable of doing so, we have not yet sonified any GONG data.

3.3. The SoSH Tool

The SoSH Project developed a tool to input complex spherical harmonic timeseries and output sonified audio of that data. The only preprocessing needed is to convert the timeseries from FITS (Flexible Image Transport System) to WAV format. This first step represents a pure audification of the data [16]. As discussed below, the subsequent processing done by the SoSH Tool exactly parallels the steps in the scientific analysis. By using fitted mode parameters to filter and transform the data, we produce a sonification.

The strongest of the Sun's harmonics have periods of about 5 minutes, corresponding to frequencies of only about 0.003 hertz. Unfortunately, this is far below the range of human hearing, which is typically taken to be 20 – 20,000 hertz, although most people are only sensitive to a smaller range. In order to experience the sound of the Sun with our ears, these very low sounds must be scaled to the range we can hear.

3.3.1. Sample Rate Time Scaling

The most straightforward way to do so would be to use the spherical harmonic timeseries we already have in hand and speed them up. But by how much? The answer of course is arbitrary and will depend on your preference, but as long as this choice is applied consistently you will still be able to hear the real relationship between different solar tones.

Let us suppose that we want to transpose a mode in the peak power range at about 0.003 hertz up to 300 hertz; this amounts to speeding up the timeseries by a factor of 100,000. If we have 72 days of data taken once a minute, it amounts to 103,680 data points. If each data point becomes an audio sample, the sped-up timeseries would now play in just over a minute. One must also consider the sample rate. Speeding up the original sample rate of 1/60 hertz by a factor of 100,000 yields a new sample rate of 1666.67 hertz. Unfortunately, most audio players will not play any sample rate less than 8000 hertz. Assuming this sample rate, our 0.003 hertz mode on the Sun will now be transposed up to 1440 hertz and the timeseries will play in about 13 seconds.

But suppose you want to play it in a shorter time; 13 seconds is a long time to sound a single note, although you might want to do so in some circumstances. You could increase the sample

rate further still, but at some point the mode will be transposed to an uncomfortably high frequency. For example, if we sped up the sample rate so that the duration is only 1 second, the resulting frequency would be over 18000 hertz. To understand the solution to this problem, we must explore the process by which we shall isolate the modes.

3.3.2. Modal Isolation

At this point in our processing, playing an unfiltered timeseries would sound just like static, or noise. This is because very many modes are sounding simultaneously in any given timeseries, not to mention the background noise involved in our observation of the modes. Therefore, if we want to isolate a single mode, we have to do some filtering. Luckily, as mentioned above, we have already measured the frequency, amplitude, and width of many modes. We can use these fitted mode parameters to pick out the particular part of the power spectrum corresponding to a single mode, and set the rest of the power spectrum artificially to zero. Once transformed back into the time domain, we can play the filtered data back and hear only the isolated mode.

Properly speaking, we do not compute any power spectra, but rather retain both the real and imaginary parts of the Fourier transform. The input timeseries are also complex because each of them contains both signs of the azimuthal order m (the timeseries for $m = 0$ has only a real part). Because we also perform the transform over the full length of the timeseries rather than windowing, we preserve the maximum amount of phase information present in the original signal. Although of course we cannot hear the phase, it may nonetheless become important when adding different modes together. Finally, we note that the output is strictly real because it now contains only a single m .

3.3.3. Pitch Shifting and Time Stretching

Since we are selecting only a narrow range of frequencies, we have the freedom to shift the entire power spectrum down in frequency before we transform back to timeseries. This timeseries will play in the same amount of time as before, but the frequencies in it will be transposed down by the same factor that we shifted the power spectrum. For every power of 2 shifted down, the tone will drop by one octave. This allows for the resulting audio to be set to any arbitrary pitch, solving the problem with uncomfortably high frequency playback.

As long as you use the same sample rate and downshift factor when you sonify every mode, the frequency relationships between them will be preserved. In other words, you will hear the same musical intervals that exist on the Sun.

Conversely, the duration of the resulting sonification can be altered by modifying the sample rate, although one must also multiply the downshift factor by the same amount if they desire to keep the same total transposition factor. Any arbitrary sample rate is permissible. The sample rate modification happens after the inverse transform and uses bilinear interpolation to maintain an internal sample rate of 44,100 hertz. When writing the output to a WAV file, however, the chosen sample rate is written into the WAV header.

3.3.4. Using the SoSH Tool

The SoSH Tool is an open-source system implemented in the visual programming language Pure Data (PD). Upon launching the

demo patch “modefilter_standalone.pd” (see Figure 6), you are required to specify a path to your local data directory. The patch will search this directory for ASCII files containing the fitted mode parameters, as well as the WAV files for whatever spherical harmonics you wish to sonify.

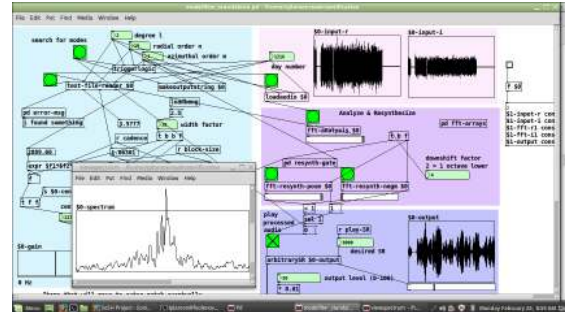


Figure 6: A screen shot of the core processing engine inside the SoSH Tool.

With these files in place, the dataset is now accessible by the day number corresponding to the first day of the timeseries, the spherical harmonic degree l , the radial order n , and the azimuthal order m . By default, the tool is set to use MDI data. If it is successfully able to load the necessary audio files, it will automatically trigger FFT analysis of them. Next, the fitted mode parameters corresponding to the inputs are used to generate a gain array in which full output is allowed for a frequency interval centered on the mode frequency and of width five times the mode width. Full attenuation is set for all other bins. This default width of the pass band is the same as that used for the fitting, but it can be increased multiplicatively in the software by specifying the “width factor.”

Once the gain array is generated, the FFT data is multiplied by it and resynthesized. If a downshift factor is specified, the FFT will be shifted down by this amount before the inverse transform.

Beyond the basic sonification of single modes, several extensions to the SoSH Tool enable more advanced functions. One extension processes several modes and additionally outputs their results into a single array (see Figure 7). Another extension turns the tools into a kind of multitimbral sampler that can trigger different modes assigned to different MIDI pitches. Yet another extension reads a score from a text file to produce a sequence of modes.

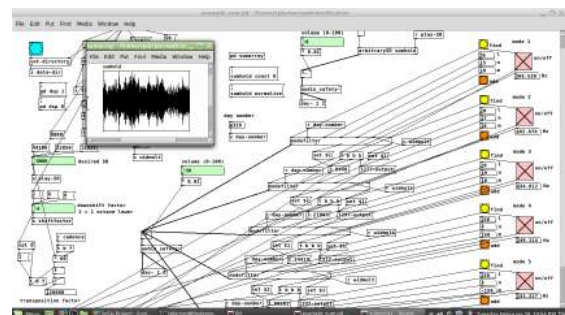


Figure 7: A screen shot of an extension to the SoSH Tool that additively combines modes into a single output array.

The SoSH Tool, datasets, audio samples of sonified data, and

much more can be downloaded from the SoSH website [17].

3.4. Audible Science

The physical phenomenon most accessible to hearing is solar rotation. In just a single timeseries of moderate m , the frequency splitting between the two signs results in a clearly audible beat frequency when they are played together. In a similar fashion, one may also easily hear the so-called small frequency separation, which is sensitive to conditions in the solar core [18]. Both types of frequency difference are known to vary with the solar cycle. However, if we apply an absolute frequency shift rather than a relative one, one may hear the effect of the solar cycle on a single mode by juxtaposing timeseries from two different epochs.

4. CREATIVE APPLICATIONS

While the SoSH Tool is clearly a potential benefit to the scientific community, it also generates possibilities for creative output. Music compositions, installation artwork, a documentary, and a full-dome planetarium experience are detailed below.

4.1. Music Composition

Music composition is one creative outcome from the SoSH Project. diFalco’s composition *Helios* (2018), from a set of works collectively titled *Cosmophonía*, is written for orchestra and imagines a journey directly through the center of the Sun beginning at the furthest reaches of the heliosphere [19]. The composition applies the acoustic phenomenon of helioseismology by scoring the natural harmonic series as it relates to the position inside or outside of the Sun. In addition, the primary pitch material is derived from a solar flare recorded on 23 November 1998 by Peter Messmer et al. at the Institute of Astronomy in Switzerland (see Figure 8) [20].

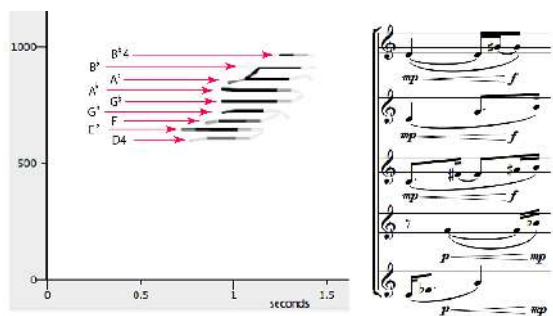


Figure 8: Analysis and orchestration of a solar event used in *Helios* (2018) by Elaine diFalco.

Many possibilities exist to directly implement the audio output from the SoSH Tool into compositions employing electronic means. As noted above, this audio can be manipulated to map to any pitch or stretched/compressed to any duration. One particularly exciting application is using the sonified data at low rate (perhaps even the real-time rate of one data point per minute) to drive other processes like synthesis parameters, sample playback, tempo rate, dynamic contour, etc. In this way, the sonified data can act as an LFO to shape larger-scale parameters.

The authors are collaboratively composing a new work for orchestra and electronics using the SoSH Tool to generate both the acoustic score and the accompanying electronics. Given the large datasets available, the performance will sonify an entire solar cycle to create a kind of “solar concerto.”

4.2. Installation Artwork

The application of the SoSH Project to installation based work is particularly attractive. Shafer’s multichannel projector installation *Sol Invictus* (2015) uses the raw visual data from the Atmospheric Imaging Assembly (AIA) onboard SDO to display the complex convection happening in the chromosphere (see Figure 9) [21]. The gigantic scale of the solar forces is subverted by the minuscule dimension of the images generated by the five micro projectors.



Figure 9: A still from the multichannel projector installation *Sol Invictus* (2015) by Seth Shafer.

4.3. Documentary and Intermedia

The process of creating immersive experiences to understand the abstract datasets of helioseismology is the subject of diFalco’s documentary *Immersing* (2018). The film discusses the topic of helioseismology, the complexity of the dataset, and the intermedia practice of creating immersive experiences.



Figure 10: Still from the documentary *Immersing* (2018) by Elaine diFalco.

Another important outcome from the SoSH Project is the SoshPy Visualization Package. Written in Python, this software package was used to generate Figures 1–4 of this paper. It can also plot all three components of the vector velocity associated with any given mode, as well as its energy density. More interestingly, it is also able to plot sums of modes and create animations.

Plans for a large-scale intermedia experience using the SoSH Tool are underway. The full-dome projection theater, commonly used in planetariums, is an ideal platform for the SoSH Project due to the immersive possibilities and its attraction to the science-seeking public. Due to the shared use of spherical harmonics for representation, Ambisonic panning derived from the data will be used to direct spatial placement in the full-dome [22]. Further details of this creative activity will emerge as the project progresses.

5. CONCLUSIONS AND FUTURE RESEARCH

The SoSH Project designed a software tool that sonifies helioseismic data from the largest and most recent datasets available. The project is a significant update to and expansion upon earlier efforts to sonify data related to the surface and interior of the Sun.

Current efforts are underway to create an interface between the SoSH Tool and the SoshPy Visualization Package. Both programs will ultimately read from the same “score”; a SoSH extension will create the soundtrack for the SoshPy animations. Going the other direction, we will use SoshPy to trace a trajectory through the Sun and then output the “score” corresponding to the relative amplitudes of modes we specify.

Finally, further creative applications of the SoSH Tool are in development. The full-dome experience described above is in early planning stages and will benefit from community feedback on the project at large and the SoSH Tool specifically.

We hope that the accessibility of the data and now tools to easily sonify that data will lead to new scientific insights related to our home star. In addition, as we attempt to make artwork from these fascinating natural phenomenon, we encourage others to also generate new creative work inspired directly or indirectly from this project.

6. ACKNOWLEDGMENT

The authors would like to thank Jesper Schou for helpful discussions.

7. REFERENCES

- [1] A. Graps. (2009) Helioseismology. [Online]. Available: <http://soi.stanford.edu/results/heliowhat.html>
- [2] R. L. Alexander et al., “Audification as a diagnostic tool for exploratory heliospheric data analysis,” *The 17th International Conference on Auditory Display*, pp. 1–4, 2011.
- [3] D. O. Gough, personal communication, 2018-08-31.
- [4] A. Kosovichev. (2008) The singing sun. [Online]. Available: <http://solar-center.stanford.edu/singing/singing.html>
- [5] D. Scherrer. (2010) The sun in music. [Online]. Available: <http://solar-center.stanford.edu/art/music.html>
- [6] T. Lausten. (2000) Sol: Data of two suncycles. [Online]. Available: <http://www.sol-sol.de>
- [7] T. Ashcraft. (2010) Heliotown–radio sun. [Online]. Available: http://www.heliotown.com/Radio_Sun_Introduction.html
- [8] R. L. Alexander et al., “Sonification of ace level 2 solar wind data,” *The 16th International Conference on Auditory Display*, pp. 39–40, 2010.
- [9] J. Luhmann et al. (2010) Uc berkeley–sounds of space. [Online]. Available: http://cse.ssl.berkeley.edu/stereo_solarwind/sounds.html
- [10] D. Gurnett. (2018) The university of iowa–space sounds. [Online]. Available: <http://www-pw.physics.uiowa.edu/space-audio/index.html>
- [11] C. Hayward, “Listening to the earth sing,” in *Auditory Display: Sonification, Audification, and Auditory Interfaces*, G. Kramer, Ed. CRC Press, 1994, pp. 369–405.
- [12] F. Dombois et al., “Sonifyer: A concept, a software, a platform,” *Proceedings of the 2008 International Conference on Auditory Display*, pp. 1–4, 2008.
- [13] P. H. Scherrer et al. (2011) The michelson doppler imager. [Online]. Available: <http://soi.stanford.edu/>
- [14] J. W. Harvey et al. (1996) The global oscillation network group (gong) project. [Online]. Available: <https://gong.nso.edu/>
- [15] J. Schou et al. (2018) Helioseismic and magnetic imager. [Online]. Available: <http://hmi.stanford.edu/>
- [16] F. Dombois and G. Eckel, “Audification,” in *The Sonification Handbook*, Thomas Hermann et al., Ed. Logos Publishing House, 2011, pp. 301–324.
- [17] T. Larson, S. Shafer, and E. diFalco. (2018) Sonification of solar harmonics project. [Online]. Available: <http://solar-center.stanford.edu/SoSH/>
- [18] W. J. Chaplin, *Music of the Sun: The Story of Helioseismology*. Oneworld Publications, 2006.
- [19] E. diFalco, “*Cosmophonia: Musical expressions of astronomy and cosmology*,” Masters Thesis, University of North Texas, 2018.
- [20] P. Messmer. (1998) Acoustic recording of a type iii solar burst recorded on november 23, 1998 between 1140-2280 mhz. [Online]. Available: <http://www.astrosurf.com/luxorion/Radio/solar-burst-type3-23nov98-messmer.wav>
- [21] S. Shafer. (2015) *Sol Invictus*. [Online]. Available: <http://sethshafer.com/solinvictus.html>
- [22] M. A. Gerzon, “Ambisonics in multichannel broadcasting and video,” *J. Audio Eng. Soc.*, vol. 33, no. 11, pp. 859–871, 1985.