Using Hilbert Curves to Organize, Sample, and Sonify Solar Data

Abstract

How many ways can we explore the Sun? We have images in many wavelengths and squiggly lines of many parameters that we can use to characterize the Sun. We know that while the Sun is blindingly bright to the naked eye it also has regions that are dark in some wavelengths of light and bright in others. All of those classifications are based on vision. Hearing is another sense that can be used to explore solar data. Some data, such as the sunspot number or the extreme ultraviolet spectral irradiance, can be readily sonified by converting the data values to musical pitches. Images are more difficult. Using a raster scan algorithm to convert a full-disk image of the Sun to a stream of pixel values is dominated by the pattern of moving on and off the limb of the Sun. A sonification of such a raster scan will contain discontinuities at the limbs that mask the information contained in the image. As an alternative, Hilbert curves are continuous space-filling curves that map a linear variable onto the two-dimensional coordinates of an image. We have investigated using Hilbert curves as a way to sample and analyze solar images. Reading the image along a Hilbert curve keeps most neighborhoods close together as the resolution (i.e., the order of the Hilbert curve) increases. It also removes most of the detector size periodicities and may reveal larger-scale features. We present several examples of sonified solar data, including the sunspot number, a selection of extreme ultraviolet (EUV) spectral irradiances, different ways to sonify an EUV image, and a series of EUV images during a filament eruption.

I. INTRODUCTION

Sonifying a data set has the basic purposes of making data accessible to the blind and allowing the data to serve as an adjunct to other senses. It can also help all to appreciate or understand a data set in a new way. Although a one-dimensional data set can be sonified by scaling the data to pitches, image data is a more ambitious target. The data has variations in two dimensions that should be represented by the sonification, and variations seen in a series of images are even more difficult to sonify. Solar data is often in the form of images and the changes in time and space are an integral part of understanding solar variations. We will describe sonifying several solar datasets, including an exploration of ways to sonify solar images in space and time.

Composers have used many ways to create sounds and music that mimic the natural and mechanical worlds. Camille Saint-Saëns used pianos and other instruments to imitate about 14 animals in *The Carnival of the Animals*.[1] Old-time fiddle tunes use the flexibility of the combined performer and violin to imitate chickens and other natural sounds. Luigi Russolo built “intonarumori" to produce a broad spectrum of modulated, rhythmic sounds that imitated machines.[2] He also developed a graphical form of musical score to compose pieces for these devices. Others have produced music from time sequences of the natural world. A well-known example is Concret PH,[3] which was created by splicing together short, random segments of tape recordings of burning charcoal.

Analog electronic synthesizers provided another path. In their early stages they were often used to produce sound effects. As analog synthesizers became more capable, such as the Moog modular synthesizer,[4] some used them to reproduce well-known musical pieces in electronic form (e.g., *Switched-On Bach* by Wendy Carlos, 1968) while others invented new types of music (the improvisations of Keith Emerson in the works of Emerson, Lake, and Palmer.)

Few, if any, of these techniques are examples of sonifying data. Sonification can be as simple as the shrieking of a smoke alarm or as complicated as converting multi-dimensional data to an audible signal. The incessant beeps and whistles of electronic vital sign monitors in hospitals are one example where the change in a sound signals a change in health of the patient. These use the 1-D structure of sound to convey information that conditions are both normal and alarming. Whistlers are an example of how scientists sonified radio frequency data to study the ionosphere.[5]

The interactions of lightning with the electrons in the magnetosphere are heard as descending tones lasting a few seconds.

Digital electronic synthesizers give us the ability to convert any type of information from a digital representation into music.[6] One example is that a 1-D time series can be sonified by scaling the values to musical pitches, assuming a constant duration for each value, to produce a set of Musical Instrument Digital Interface (MIDI[7]) commands. A MIDI-enabled synthesizer is then used to create the musical instrument waveforms and play the commands in the MIDI file. Different time series can be combined into a sonification by using different pitch ranges or timbres (the distinctive set of tones in the selected instrument) to distinguish between them. We will use the International Sunspot Number (Version 2, S) and extreme ultraviolet (EUV) spectral irradiances from two satellites as examples of solar time series data.

Sonifying an image is different. Sound is intrinsically a 1-D format that evolves in time. A 2-D image must be converted to a 1-D series of pixel values where the order of the pixels serves as the time variable. Once the 1-D sequence exists, the pixel values are scaled to pitches, the duration is again set to a constant, and the data run through the synthesizer.

There are many ways to map a 2-D image (or higher-dimensional data) to a 1-D sequence. A raster scan is a linear reading of the image from the upper left to the lower right moving down to the next row when the current one is read, much like reading an English language document. This can be modified into a boustrophedonic algorithm where the first row is read left to right and the next right to left, continuing in this way to the end of the image. This resembles the way an ox (Greek *bous*) plows a field and hence the term. Another way is to use a space-filling curve, such as the Hilbert used here, to map the image pixels to a sequence. We will describe using Hilbert curves to convert 2-D images into 1-D sequences and converting those sequences to sound.

The sonifications of these data sets will be described:

1. International Sunspot Number (annual and monthly variations)
2. Extreme ultraviolet (EUV) spectral irradiances as a time series and a spectrum
3. A complete EUV image and seven subimages
4. A montage of EUV images showing a filament liftoff

All of the sound files are available as .MIDI and .MP3 files at https://sdo.gsfc.nasa.gov/sonify/table.html.

We start by introducing some useful musical concepts. That will be followed by a discussion of the synthesizer used and the analysis of of the 1-D data sets. The image data will be introduced, and an example using a raster scan to convert the data to 1-D will be described. We will then describe the Hilbert curves used to address the image data and present several ways to sonify the images. We discuss what can be learned from these sonifications and end with several conclusions on the utility of this method. All of the science datasets are open-source and are available at the locations listed in the Acknowledgements.

II. SONIFYING DATA

The JythonMusic software described in Manaris and Brown[8] was used to convert a data series into MIDI commands and drive a synthesizer. The concepts and terms we use to convert data to music are:

- Pitch: One of 128 frequencies (spanning 10.75 octaves of the 12-tone equal-tempered scale), from 8.18 Hz [C_{-1}] – 12.54 kHz [G_9]), with Middle C (C4, 261.63 Hz) roughly in the middle at position 60. Twenty one pitches are added below the lowest note on the piano and 19 pitches above the highest note. A range of only 128 values is small compared to the linear range of many solar and geophysical data sets. It is also small compared to the pitch discrimination of human ears. Untrained humans can discern pitch changes of $\approx 0.3\%$,[9] so roughly 43000 pitches would be necessary to resolve that frequency range. However, the MIDI standard only allows limited microtones at that spacing. Images encoded with the Joint Photographic Experts Group (JPEG[10]) algorithm have pixel values ranging from 0–255 (either in separate channels or through a color table), so we have only half of the range in pitches. Transforming data that varies by several orders of magnitude into logarithms can compress the range to small enough to sonify.

- Duration and Tempo: The duration (or length) of pitches and rests (periods of time without any sound) are specified with a floating point number that can vary from 0 (no time) to **1 (a quarter note) to** 4 (corresponds to a whole note) and longer. The tempo of a piece is the speed at which the pitches and rests are heard. Tempo is specified by the number of beats per minute (bpm); where a quarter note (QN in JythonMusic) is one beat. Durations are relative to the tempo of the piece, increasing the tempo proportionally reduces the duration of all pitches and rests. **We commonly use sixteenth notes in the image sonifications. A sixteenth note has a duration of 0.25 relative to a quarter note.**

- Loudness: The loudness (also called the dynamics or MIDI velocity) is set by an integer in the range 0 (silent) to 127 (very, very loud). As the range of sound pressure level varies from 0 dB (threshold of hearing) to 120 dB (threshold of pain), the loudness maps to a change of roughly one per dB. The response of human ears to loudness variations strongly varies from one person to another and with frequency. The least noticeable change in loudness also varies with frequency, but a reasonable value is 0.4 dB.[11] This corresponds to a 5% change in pressure and is easily accommodated by the 128 possible values. We only use loudness to weight the various datasets. It is also possible to encode information in the loudness, such as a longer duration being louder, but we do not present such cases here.
- Timbre: There are 128 possible timbres in the MIDI standard, which are refereed to as tracks. These timbres are not specified in the MIDI standard and a numbered timbre may sound different in different synthesizers. One track is devoted to percussion and uses the pitch designator to select a percussive timbre.
- Pan: Position in space is limited in this study to left-right pan. A floating point number between 0 (left) and 1 (right) determines the position, with 0.5 (centered) the default. Placing one data set in the left side and another in the right is a good way to compare two data sets. Where they agree the sounds will appear to come from the middle and otherwise they will come from separate sides.

JythonMusic is based on Java rather than C. Programs in JythonMusic are written in Python 2.7 syntax but do not have access to many of the libraries used for numerical work. As a result, data access and extraction routines were written and executed in a C-based Python environment that provided access to the NumPy library for array manipulation. The computational sequence was to read the data, extract the appropriate part, write the extracted data to a comma-separated variable (CSV) file, read that file in the JythonMusic environment, convert the data into a MIDI file, and use the JythonMusic synthesizer to play that file. **A permanent record was created by playing the MIDI commands in another synthesizer that could export the sounds to an MP3 file.**

III. SAMPLING AND SONIFYING SOLAR DATA

Several types of solar data were sonified and reported here. A summary is presented in Table I, where the source, type, and name of the corresponding MP3 file are listed. The Sec. column is the part of the paper where the data is described. A version of this table, with links to the MP3 and MIDI files, is available at https://sdo.gsfc.nasa.gov/sonify/table.html.

TABLE I. Files for each Sonified Data Set

Sec.	Source	Sonified Data	mp3 Filename
III A	SIDC	Sunspot number	TS_sunspot_annual_month.mp3
III B	EVE	EUV spectral irradiances (spectrum)	TS_EVE_sonified.mp3
III B	SEE	EUV spectral irradiances (time series)	TS_SEE_sonified.mp3
III C	AIA 193 Å	Complete image (raster)	AIA_193_full_image_sonified_raster.mp3
V	AIA 193 Å	Complete image (Hilbert)	AIA_193_full_image_sonified.mp3
V A	AIA 193 Å	Subimage 1 (Arcs)	subimage_1_x_685_y_1755.mp3
V A	AIA 193 Å	Subimage 2 (Fan)	subimage_2_x_1060_y_1120.mp3
V A	AIA 193 Å	Subimage 3 (Island)	subimage_3_x_1290_y_1690.mp3
V A	AIA 193 Å	Subimage 4 (Limb)	subimage_4_x_1800_y_992.mp3
V A	AIA 193 Å	Subimage 5 (Spot)	subimage_5_x_890_y_1035.mp3
V A	AIA 193 Å	Subimage 6 (Swirl)	subimage_6_x_750_y_1125.mp3
V A	AIA 193 Å	Subimage 7 (X)	subimage_7_x_760_y_405.mp3
V B	AIA 193 Å	Filament liftoff montage	liftoff_complete.mp3

A. International Sunspot Number

The first example is the variation of the International Sunspot Number (S) with time. The sunspot number is a weighted count of dark regions on the Sun that is often used as a long-term index of solar activity. It has been measured or derived for roughly 400 years. It is the source of much of our knowledge of the evolution of solar activity. We use Version 2 of the International Sunspot Number[12,13] between 01 Jan 1750 and 31 Dec 2018 from the Solar Influences Data analysis Center (SIDC) website, both the monthly and annually averaged values. The time dependence of S is shown in Figure 1.

After some experimentation, we determined the following sonification. The annually-averaged values were mapped to pitches between 48 and 96 of the lower timbre (PICKED_BASS) and a loudness of 125. This data was used to set the tempo (one year is one beat with a duration of a quarter note) of 400 bpm. The monthly-averaged values were mapped to pitches between 60 and 108 in the PIANO timbre and a loudness of 100. This data is played at 12 values per beat **(a group of 12 thirty-second-note triplets**), and was panned left-right with a two-year period. This allows you to hear the differences in the two signals. The slower lower voice can be audibly distinguished from the more rapidly varying higher voice.

You can listen to this sonification at
https://sdo.gsfc.nasa.gov/iposter/mp3/TS_sunspot_annual_month.mid.mp3

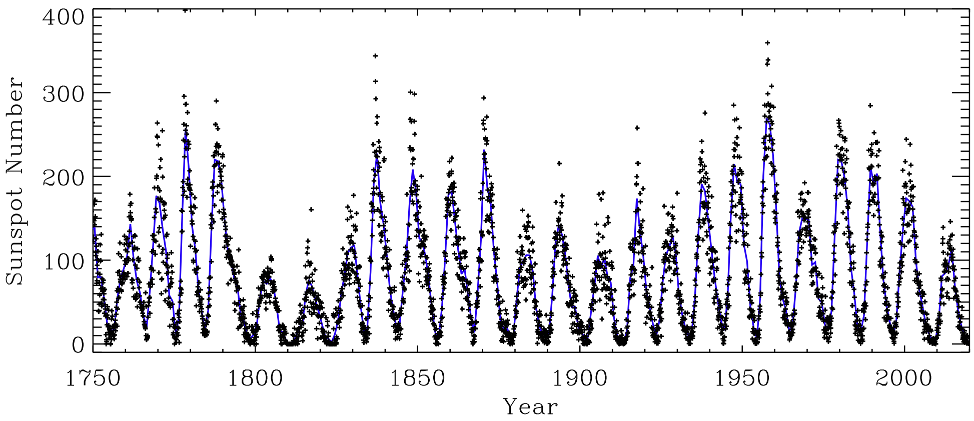

FIG. 1. Version 2 of the International Sunspot Number as a function of time since 1750. The solid blue line is the annually-averaged data (plotted at the middle of each year) and the '+' symbols show the monthly-averaged values.

B. Extreme Ultraviolet Spectral Irradiances

The next example is to sonify extreme ultraviolet (EUV) spectral irradiances from two instruments in two ways. The solar EUV spectral irradiance spans wavelengths between X-rays and the ultraviolet (roughly 10–100 nm) but is often extended to include the H I 1216 emission line (Ly-α). (Emission lines are described by the element symbol, the ion state of the element [where H I is neutral hydrogen, H II is singly-ionized hydrogen, etc.], and the wavelength of the line in Å.) This radiation is easily absorbed as it ionizes the outer electrons of many elements. This also makes it the major source of the ionosphere in the terrestrial and planetary atmospheres. The EUV emissions are also a direct measure of the magnetic field. The Sun would have considerably smaller EUV emissions if it did not have a magnetic field. The ratio of the spectral irradiance at the EUV wavelength of 30.4 nm to the peak value of a 5770 K blackbody at a visible wavelength of 500 nm is 10^{-26}. This ratio is 10^{-4} in a solar spectrum. These two properties, sensitivity to the solar magnetic field and acting as the source of the ionosphere, make measurements of the solar EUV spectral irradiance a primary goal in solar physics.

Solar EUV spectral irradiances are completely absorbed by the atmosphere and must be measured by an instrument in space. These instruments record the spectral irradiances as a function of wavelength and time. We first sonify a single spectrum from the Extreme ultraviolet Variability Experiment (EVE)[14] on NASA's Solar Dynamics Observatory (SDO).[15] EVE data is available from 5 to 105 nm from 1 May 2010 until 26 May 2014 and from 37–105 nm thereafter. Figure 2 shows that the EUV spectral irradiance has many emission lines, two of the strongest (He I 304 and C III 977) are labeled, and several roughly triangular regions of continuum emission (such as the one highlighted between 70 nm and 91 nm.) The third label points to the emission line Fe XII 193, which will be explored in later sections.

We elected to sonify the day-averaged solar EUV spectrum from EVE on 27 Feb 2014, the day of maximum sunspot number for Solar Cycle 24 (Figure 2). The log of the spectral irradiances was scaled to MIDI frequencies 36–96. That means every order of magnitude in the data spans about 1.5 octaves. The PIANO timbre was used, **each value occupies an 8^{th} note,** the tempo was set to 600 bpm, and the loudness was set to 80.

This example shows how the independent variable, in this case wavelength, does not have to

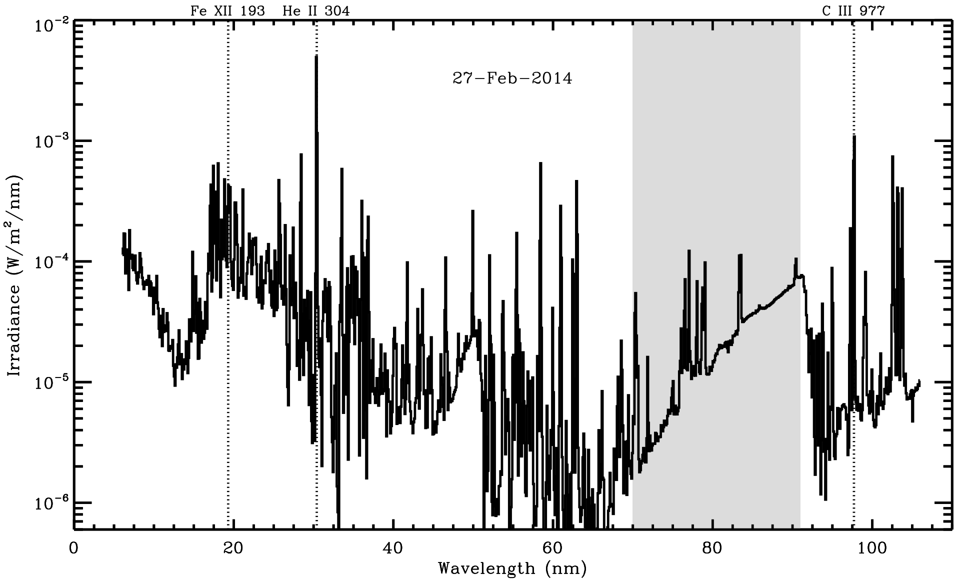

FIG. 2. A day-averaged EUV spectral irradiance for 27 Feb 2014, as measured by EVE, plotted against the wavelength in nm. The seven AIA passbands are identified with vertical dashed lines. The He II 304 Å line is the brightest in this wavelength range, with the C III 977 Å line the next brightest. The 70–90 nm continuum emission region is highlighted. The Fe XII 193 emission line will be analyzed in images below. Although the total radiant energy in this spectrum is 4.7 mW m^{-2}, about 10^{-5} times the total solar irradiance of 1361 W m^{-2}, it is responsible for much of the ionization in the thermospheres of the Earth, Venus, and Mars.

be time to sonify a data set. The independent variable must at least provide an ordering of the data set, in this case with a uniform spacing between the data points. This was judged to be the most musical example. Some of Bach's Goldberg Variations (BWV 988) sound much like this sonification. **Variation 24**, at around the 33-minute mark as played by Glenn Gould in his 1981 album of the same name, has several long chromatic runs that sound like the gradual rise of the EUV spectrum between 70–91 nm. The rapid increases in pitch of the strong spectral lines also add musical contrast to this piece.

You can listen to this sonification at
https://sdo.gsfc.nasa.gov/iposter/mp3/TS_EVE_sonified.mid.mp3

Spectral irradiances at selected wavelengths can also be extracted from the measurements as a function of time. The source of **another set of** EUV spectral irradiances is the Solar Extreme ultraviolet Experiment (SEE)[16] on NASA's Thermosphere Ionosphere Mesosphere Energetics and Dynamics (TIMED) spacecraft, **which provides daily values of the irradiances between 0.5 and 195 nm**. The spectral irradiances from 9 Feb 2002 to 11 May 2019 of several strong emission lines (He I 304, Ly-α, C IV 1548, and Fe XVI 335), along with the 0.1–7 nm soft X-ray radiometer channel, were sonified. The time dependence of these channels is shown in Figure 3. Pitches between 24 and 108 were interpolated from the log of the irradiances using the maximum and minimum of each channel as the limits. This forces the channels to have the same pitch range. The timbres were PIANO, PICKED_BASS, TROMBONE, FLUTE, and MARIMBA, respectively. **Each value occupies an eighth note, the** tempo was set to 640 bpm and the loudness was set to 100.

You can listen to this sonification at https://sdo.gsfc.nasa.gov/iposter/mp3/TS_SEE_sonified.mid.mp3

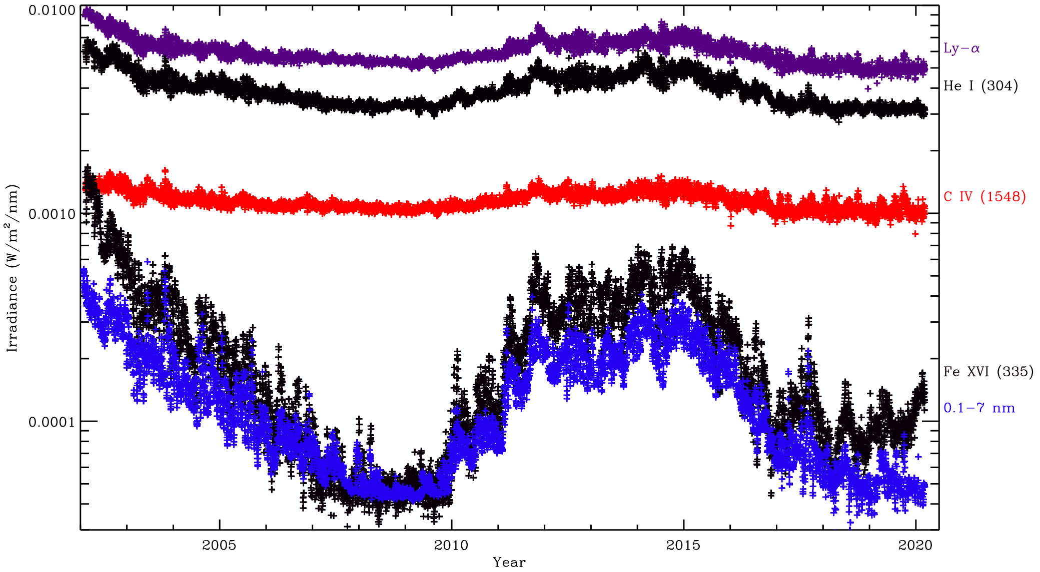

FIG. 3. The variation of selected EUV spectral irradiances from SEE with time. The selected wavelengths show different levels of solar cycle modulation. The Ly-α and 0.1–7 nm irradiances were divided by 10 and 20, respectively, before plotting.

C. Extreme Ultraviolet Images

Although measuring the EUV spectral irradiance is important, understanding those emissions requires that you also have images at those wavelengths showing how the source regions of the emissions vary in both space and time. Extreme ultraviolet images from the Atmospheric Imaging Assembly (AIA)[17] on SDO were sonified as complete images, subimages, and a time sequence of subimages. AIA provides 10 passbands: seven EUV, two ultraviolet, and one visible light. AIA 193 Å images were selected as they highlighted the desired coronal details. We will describe different ways to sonify an AIA 1**93 Å image from 23:55:53 UTC on 18 Mar 2018.**

Compared with time series data, we found that images are difficult to sonify because they are dense in information and have variations in two directions. As an example of density, a sonified 512×512 image would take almost 15 hours to listen to at a moderate tempo of 300 bpm, and a full-resolution (4096×4096) AIA image would require 40 days. Many people have a hard time remembering tone sequences and whatever is happening near the end would be disconnected from the beginning. We overcome this by either binning the image to a smaller number of pixels or selecting subimages. Based on our experiences when playing our sonifications, where we found that a person can remember tone sequences for a few minutes, we aim to create sonifications that last three minutes by binning the image to 32×32 pixels or by using much higher tempos (up to 3000 bpm). Pieces such as John Cage's *Organ*^{2}*/ASLSP* (*As Slow as Possible*) may be written for performance times of hours to years, but the density of notes is far smaller in these pieces. Only 31 notes have been sounded since a 639 year version of the piece was begun in 2001 at the Burchardikirche in Halberstadt, Germany.[18] An AIA image would sound 31 notes in the first 6.2 s at our standard tempo of 300 bpm.

AIA science data is served as monochromatic, 4k×4k, 14-bit files using the Flexible Image Transport System (FITS).[19] To make these sonifications more accessible to students, we elected to use the quicklook AIA images that are served as JPEG files created from the FITS data using a log scaling and an arbitrary color table. Concentrating on converting JPEG images allowed us to test the algorithms using images with higher contrast or more distinct features. This allows the students to sonify their favorite images. When necessary the JPEG images were converted to greyscale using the luminosity form of relative luminance to weight the individual red (R), green

(G), and blue (B) channels:

$$IM(\text{B\&W}) = 0.21R + 0.72G + 0.07B. \tag{1}$$

Although the solar images used have redundant information in the separate color channels, by applying the luminosity form to all JPEG images it is possible to analyze any image with a three-color format.

D. Raster-scan Sampling

The greyscale images must now be converted into a 1-D series for sonification. The first example is to sample them along a raster scan as described above. One initial image, binned from dimensions of 2048×2048 to 32×32, is shown in Figure 4, with the image overdrawn by a raster scan to generate the sampling curve in the lower plot.

Similar to the sunspot series in § III A, the image was sampled in two different resolutions. The higher register was scaled from the 32×32 **binned image by mapping the pixel values between [0, 250] to pitches between [60, 120] (or C4 to C9, a span of 5 octaves). The duration was set to a sixteenth note, the loudness to 110, and the SOPRANO_SAX timbre was used. The lower register was added by mapping pixels from a** 16×16 **binned image with values between [0,250] to pitches [48, 96] (or C3 to C7, a span of 4 octaves). The duration was set to a quarter note, the loudness to 90, and the ACOUSTIC_GRAND timbre was used. The dark regions of the lower register were omitted by being set to the special variable REST. The lower register is the average value of the four pitches in the higher register in the same region of the image. The relative timing of the voices is arbitrary, but we keep the two different sequences synchronized by using a ratio of four to one.**

You can listen to this sonification at
https://sdo.gsfc.nasa.gov/iposter/mp3/AIA_193_full_image_sonified_raster.mp3

The lower curve (b) in Figure 4 shows how the raster scan is dominated by the quasi-periodic variations caused by the scan moving onto and off of the disk of the Sun. **We reduced some of noisy variations at low pixel values by replacing the dark regions with a rest, but the sonification still does not reveal much about the image other than the broad shape of the**

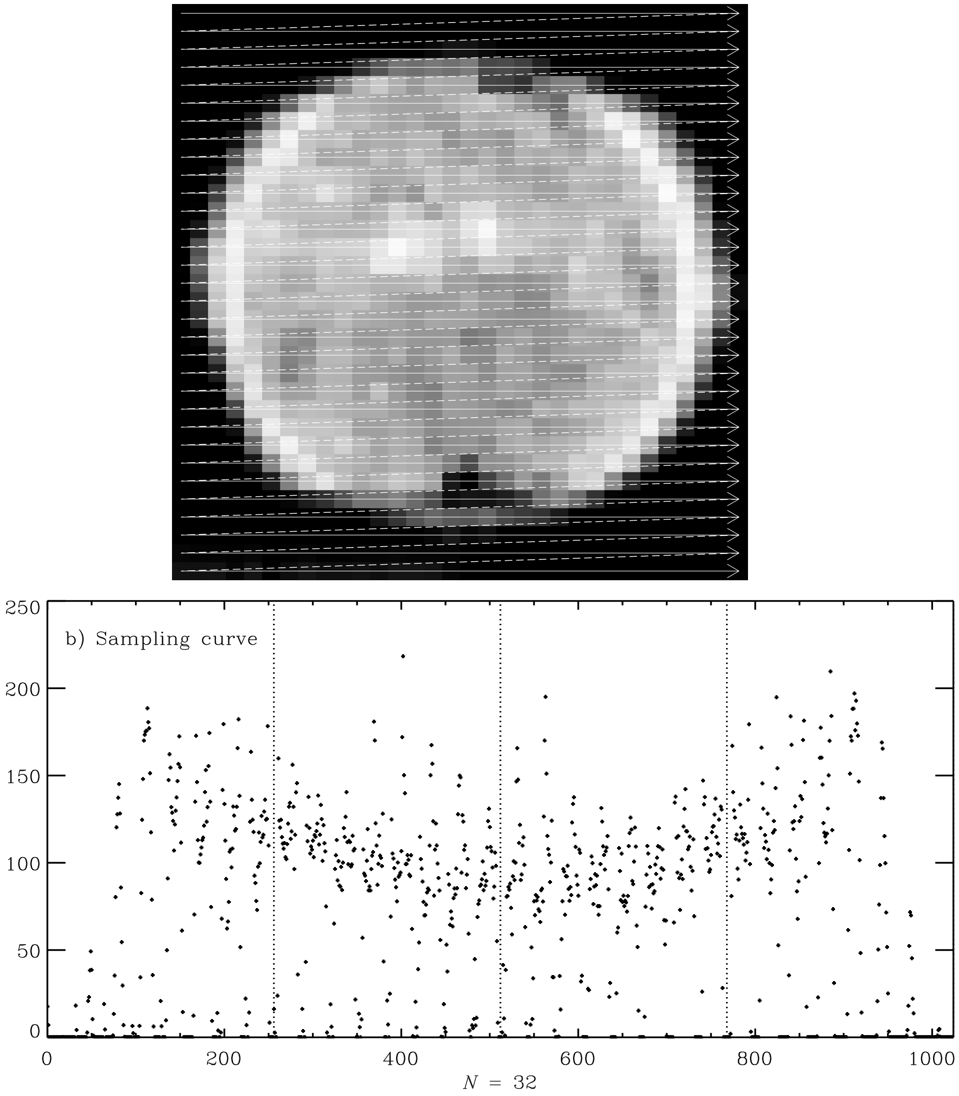

FIG. 4. A greyscale SDO/AIA 193 Å image from 18 Mar 2019 binned from 2048 × 2048 to 32 × 32. At the top is an example of how a raster scan from the top left to the lower right samples the image. The dashed lines are the return from right to left that is not used in the sampling. The resulting sampling curves for this image is shown in the lower plot (b). The vertical lines show the four horizontal strips of the image.

Sun. As a result, we explored using other methods to sample the image. The Hilbert curve was one of those methods.

IV. HILBERT CURVES

Hilbert curves are continuous space-filling curves that have been used in a surprisingly large number of disciplines. They were first described by Hilbert[20] as a simpler form of the space-filling curves of Peano[21]. A true Hilbert curve exists only as the limit of $n \to \infty$ of the n^{th} approximation to a Hilbert curve (H_n). However, the approximations are useful to provide mappings of 2-D images onto a 1-D sequence. Figure 5 shows H_n for $n = 1, 2, \ldots, 6$.

A summary of the properties of H_n:

1. There are 2^n pixels along each side of the square containing the curve
2. The Euclidean length of H_n grows exponentially with n, $2^n - 2^{-n}$
3. H_n covers a finite area as it is always bounded by the unit square
4. Two points in the image, (x_1, y_2) and (x_2, y_2), that are close together in H_n are also, with a few exceptions, close together in $H_{n'}$, $n' > n$

A Hilbert curve maps a linear variable onto the two-dimensional coordinates of an image. Its inverse is a mapping of the image coordinates onto a linear variable. This mapping property means we can use Hilbert curves to map solar images onto a linear sequence of pixel values that can then be sonified. Images tend to have dimensions that are powers of 2, so the Hilbert curves are a natural fit to addressing them.

Reading the image along a Hilbert curve has the advantage of keeping neighborhoods close together as the resolution (i.e., the length of the curve) increases. It also removes most of the detector size periodicities and actually shows the presence of longer-scale features. Because successive H_n's pass through similar neighborhoods as the resolution is refined, Hilbert curve samplings can be overplotted in time to provide contrasting versions of the image.

The neighborhood property works with other space-filling curves. Bartholdi et al.[22] describe using a Sierpinski space-filling curve to design delivery routes for Meals on Wheels. The system was simple, cheap, and paper-based. It used a manual “Rolodex" method of entering or removing addresses.

Vinoy et al.[23] and others have shown how to use Hilbert curves to construct microwave antennas. They used models and measurements of the input impedance to show that a small square

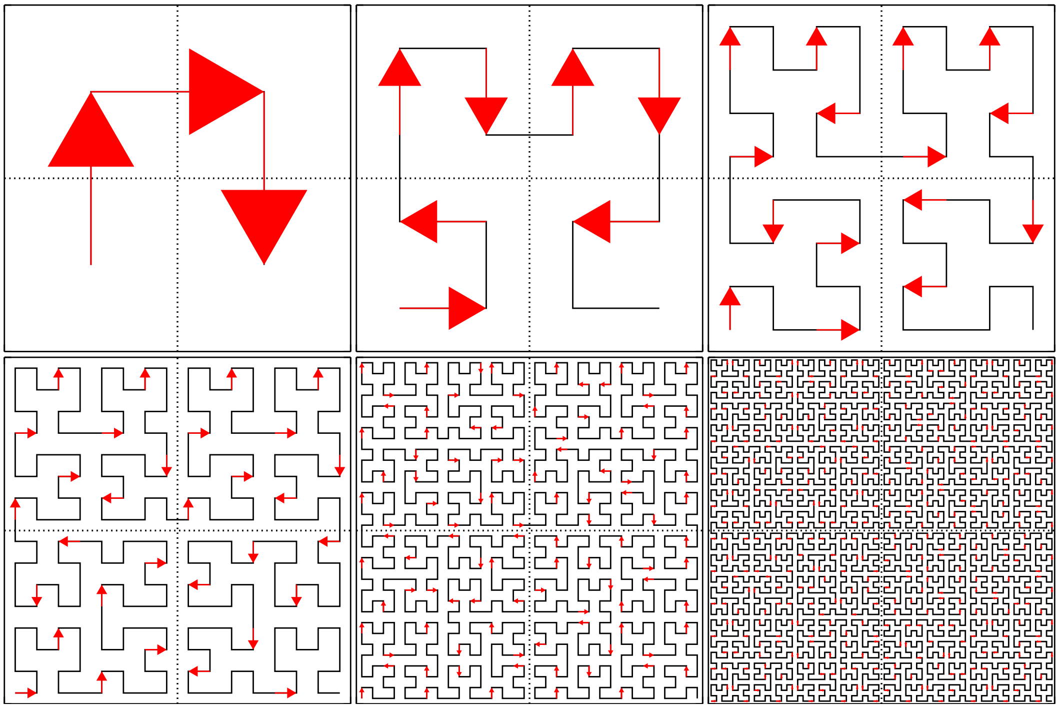

FIG. 5. The first six Hilbert curves, plotted from upper left to lower right, with arrows showing the direction of the motion into each vertex. Each subplot is drawn with axes limits of [0,1] in both directions. Among the most important properties of these curves is the single line connecting two quadrants. This can be seen by examining the dotted lines drawn to separate the quadrants. Another property is that the sampling goes around each quadrant in a similar motion (upper quadrants are sampled in a clockwise fashion and the lower quadrants in a counter-clockwise fashion.)

overlain with a conducting Hilbert curve produced an antenna whose resonance frequencies were consistent with a much longer wire antenna. They also showed how those frequencies shifted and how additional resonances were added as the order of the Hilbert curve was increased. This makes these antennas useful for mobile wireless devices.

Seeger and Widmayer[24] describe using space-filling curves to access multi-dimensional datasets with a 1-D addressing scheme. The 1-D curve imposes an order on the data access that is difficult to implement using a multi-dimensional access polynomial. Morton[25] describes using the space-filling Z-order curves to access a file address database. Like the Hilbert curve, Z-order curves preserve the locality of most of the points being mapped.

Multi-dimensional Fourier integrals (as well as others) can be reduced to a 1-D form by mapping the coordinates onto a space-filling curve, essentially converting the integral into a Lebesque integral.[26]

V. EXTREME ULTRAVIOLET IMAGES SAMPLED ALONG HILBERT CURVES

A 32×32 image can also be sampled using an $n = 5$ Hilbert curve. This is shown in the top panel of Figure 6. Similar to the raster scan method in § III D and the sunspot series in § III A, the image was sampled in two different resolutions that are then played together. The higher register was scaled from the 32×32 binned image by mapping the pixel values between [0,250] to pitches between [60, 120] (or C4 to C9, a span of 5 octaves). The duration was set to a sixteenth note, the loudness to 110, and the SOPRANO_SAX timbre was used. The lower register was added by mapping pixels from a 16×16 binned image with values between [0,250] to pitches [48, 96] (or C3 to C7, a span of 4 octaves). The duration was set to a quarter note, the loudness to 90, and the ACOUSTIC_GRAND timbre was used. The lower register is the average value of the four pitches in the higher register in the same region of the image and the two registers are synchronized.

You can listen to this sonification at
https://sdo.gsfc.nasa.gov/iposter/mp3/whole_AIA_193_full_image_sonified.mid.mp3

The difference between the sampling along a Hilbert curve and a raster scan can be seen by comparing the lower curves in Figure 6 and Figure 4. The lower curve in Figure 6 shows the Hilbert curve sampling localizes the off-disk portions of the image along the curve and hence in time in the sonified version, while the lower curve in Figure 4 shows the raster scan has strong modulations of the signal by the shape of the Sun.

A. Using Subimages to Emphasize Features in Extreme Ultraviolet Images

The AIA 193 Å image in Figure 6 is vastly undersampled. The **number of pixels in a sonified image scales as 2^{2n},** where n is the order of the Hilbert curve used to sample the image. One way to increase the accuracy of the sampling while keeping a reasonable length in the sonification is to sub-sample the image. Seven 64×64 subimages of the 2048×2048 2019 Mar 18 image are shown in Figure 7, numbered to agree with Table I.

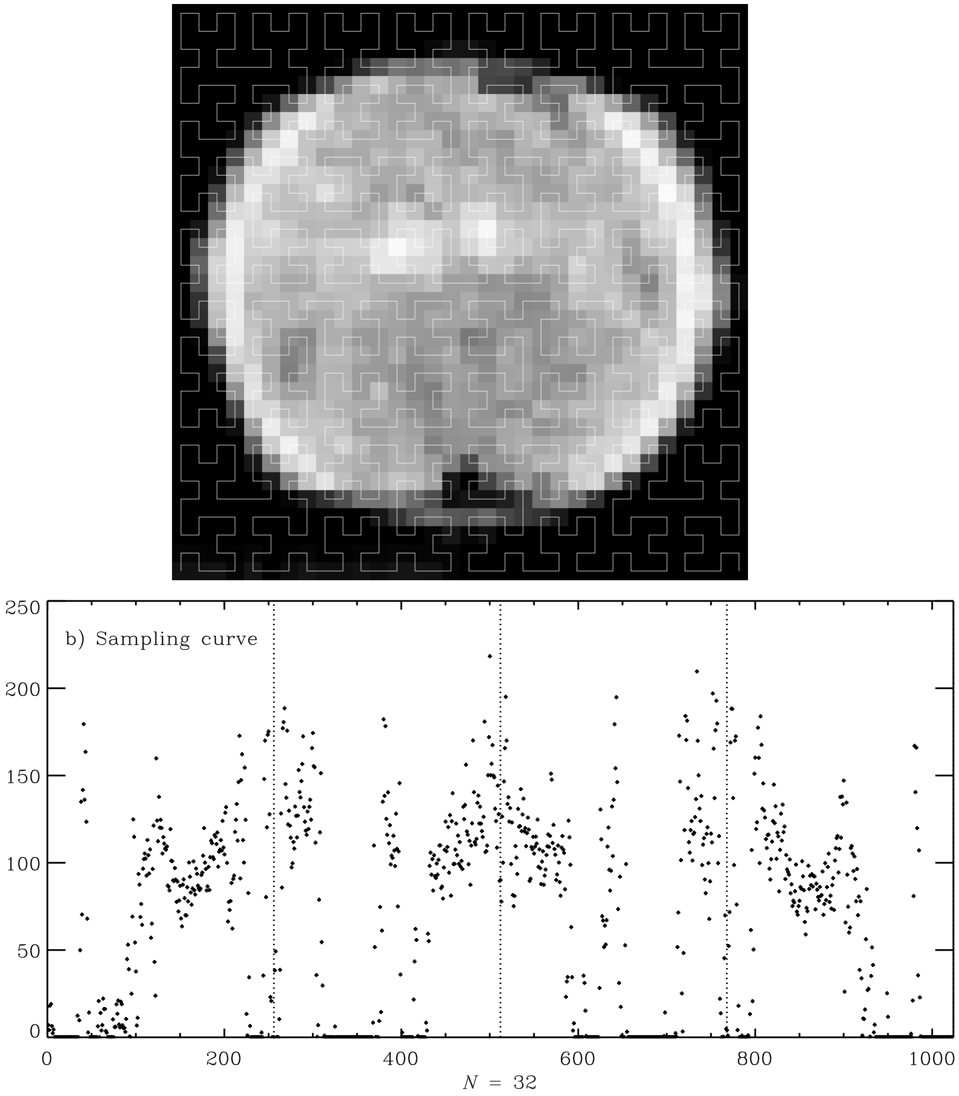

FIG. 6. The top panel has an $n = 5$ Hilbert curve (H_5) is drawn over the greyscale SDO/AIA 193 Å image from 18 Mar 2019 binned from 2048×2048 to 32×32. The lower plot (b) shows the resulting sampling curve. Each pixel in the image is assigned to a point in the curve. The centers of square pixels are located where the curve has a right angle bend, at the halfway mark of straight segments that are two units long, or two centers proportionally spaced along the straight segments that are three units long. The vertical lines show the four quadrants of the image.

Each subimage was sonified by being sampled along a Hilbert curve. The full-resolution pixels were first binned to 32×32, sampled with an $n = 5$ Hilbert curve. The tones were produced by

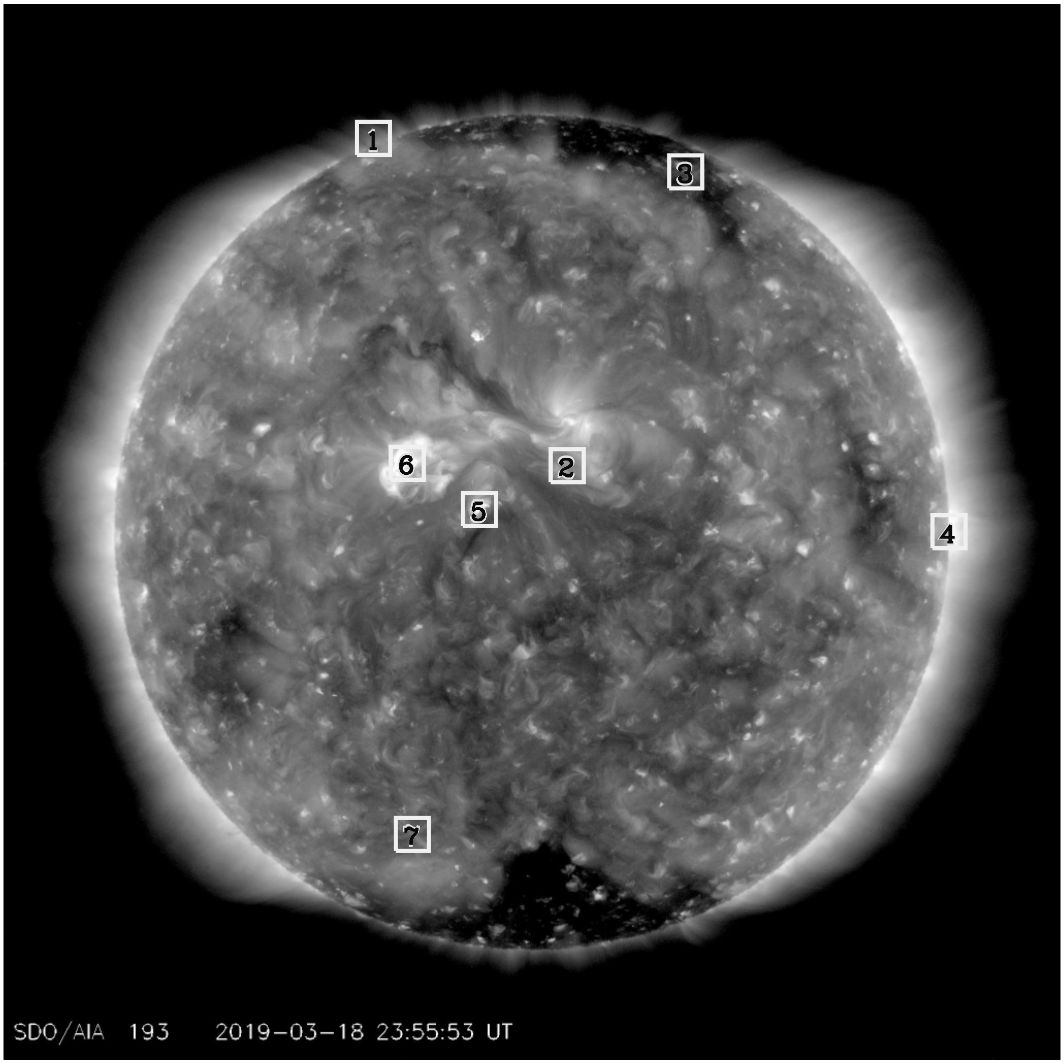

FIG. 7. A greyscale SDO/AIA 193 Å image from 18 Mar 2019. This image was used as an example for sonifying still images in § III D and § V. The boxes mark the locations of the examples in Table I.

mapping pixel values between [0, 255] to pitches between [60, 120] (or C4 to C9). The duration was set to a sixteenth note, the tempo to 300 bpm, the loudness to 110, and the SOPRANO_SAX timbre was used. A second voice was added by mapping the pixels from a 16×16 binned image with values between [0,255] to pitches [48, 84] (or C3 to C6, a span of 3 octaves). The duration was set to a quarter note, the loudness to 75, and the ACOUSTIC_GRAND timbre was used.

You can listen to this sonifications by accessing the clickable image at `https://sdo.gsfc.nasa.gov/iposter/`.

B. Filament Liftoff Sequence in Extreme Ultraviolet Images

The final example is sonifying a series of images from AIA on SDO. **The goal is to determine whether the time sequence in the images can be heard.** We selected the filament liftoff of 2010 Mar 10–12 as an example (Figure 8). Eight subimages that included the filament liftoff were extracted, binned to 32×32, sampled along an $n = 5$ Hilbert curve and sonified. A short chorus and ending cadence were written. The piece was made by inserting the subimages in turn, separated by a chorus and ending with the cadence, thus creating a single time series of pitches.

The early results for this sequence were not very successful. We then tried several ways to improve the sonification. First, the length of the individual frames was reduced by including only the lower-left and upper-left quadrants of those subimages. This corresponds to the first half of the sequence sampled by the Hilbert curve. When this did not produce a satisfactory result, we selected only those images with a noticeable difference. This produced seven images that emphasized the variation but were unevenly spaced in time. Finally, the pixels in this sequence were converted to tones by subtracting the average of each image from the sampled data, mapping the resulting values from [-60, 60] to pitches [36, 96] (or C2 to C7, a span of 5 octaves). The duration was set to a sixteenth note, the loudness to 110, and the PIANO timbre was used. Only the final attempt that includes all of these steps is presented here.

You can listen to this sonification at
https://sdo.gsfc.nasa.gov/iposter/mp3/liftoff_complete.mid.mp3

This was the least satisfying sonification because the changes in time were subtle and difficult to resolve. We have been investigating other ways to show the movement of material through both space and time. The subtraction of the mean was **one** example of one such a technique. By removing the average any overall brightening or darkening of the region did not dominate the change in time. Another possibility is to sonify the running difference images that AIA produces. **We have sonified shapes moving through space to study this effect. Part of the issue is the large number of redundant pixels that do not significantly change value in time. Sonifying a sequence in time remains an area of active research.**

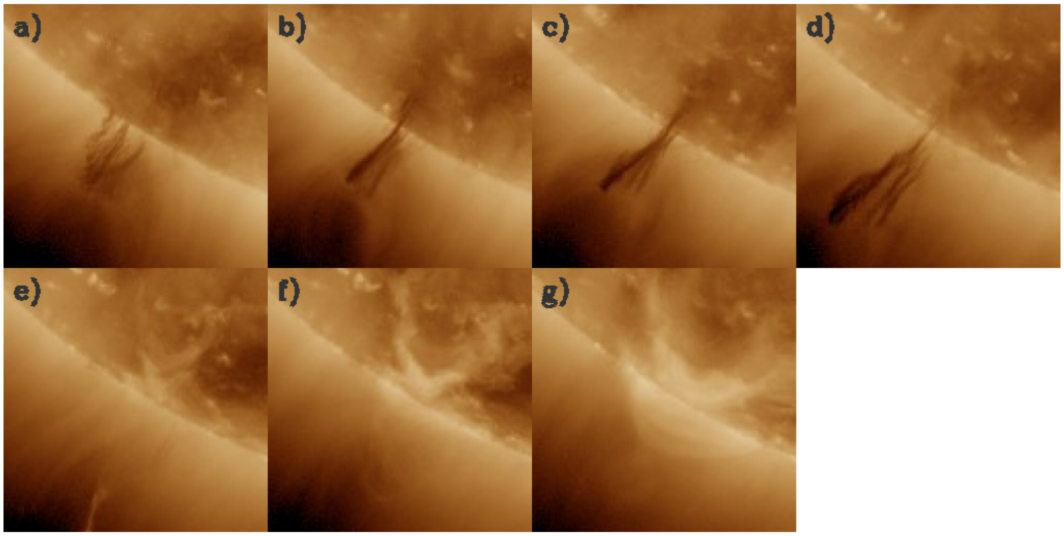

FIG. 8. Montage of the final seven solar images showing a filament liftoff. Starting from the upper left, the images were recorded at a) 2012-03-10 02:27:20, b) 2012-03-11 03:27:44, c) 2012-03-11 17:59:08, d) 2012-03-11 23:29:08, e) 2012-03-12 01:29:20, f) 2012-03-12 02:28:56, g) 2012-03-12 04:27:56, and h) 2012-03-12 06:29:56, respectively. (All times are UTC).

VI. DISCUSSION OF SONIFIED DATA

Based on our experiments, percussive sounds, such as PIANO and PICKED_BASS, seem to work better for sonifying data. Percussive timbres securely place the sound on the beat and produce interesting changes as the tempo increases. A timbre with a noticeable rise or decay time tends to sound muddy as the tempo is increased.

Our attempts to create a beat and melody by playing two versions of averaged data, such as the annual vs. the monthly values of S, were not a complete success. We continue to explore how to make the sonified data sound more like music and less mechanical.

Although sonified data does not sound like most types of music, at least some pieces of classical music has similar qualities. Bach's Goldberg Variations (BWV 988) sounds much like the image sonifications described above. As we note above, the chromatic runs in **Variation 24**, at around the 33-minute mark as played by Glenn Gould in his 1981 album of the same name, sounds quite similar to the EVE spectrum .

Sonifying data streams is a flexible area of study. To skip the image analysis step, you can

load the MIDI files provided for each of the sonifications into any compatible synthesizer. This will immediately reveal that different synthesizers assign different timbres to each numbered track, so the files will sound different in each synthesizer. You can also change the timbre of a part in the synthesizer, providing another level of experimentation. Other sound font files can also be used with the synthesizers, including the JythonMusic synthesizer, again providing another area to explore.

This also explains why the provided MP3 files do not always match what was heard when the JythonMusic synthesizer is used. You cannot create an MP3 file directly from the JythonMusic synthesizer. You can capture the sounds in either a recorder or software such as Audacity while the MIDI commands are executed. Or you can load the file containing the commands into another synthesizer that has export capability. The MP3 files provided here were created by opening the MIDI files in GarageBand, a proprietary program from Apple, and exporting the MP3 files.

You can also use other programs to generate the MIDI file from a dataset. For example, Lilypond[27] is a music engraving program that can also produce a MIDI file that is playable in a MIDI-capable synthesizer. You also get a beautiful score of the piece as a bonus. Similar to the JythonMusic workflow, the data file was opened in Python, the data was scaled to pitches and those pitches were written in Lilypond syntax to a Lilypond-readable text file. An example of a score is shown in Figure 9. Strong spectral lines can be seen in measures 31 and 35.

Mapping data to variations in pitch may not be the optimum solution for sonifying data. A large value of a dataset may be better represented by changes in the volume, emphasizing the strength of the larger value. We did some experiments on such variations and found that the limited ability of humans to sense changes in loudness and to remember a baseline level of loudness over an entire piece made this less effective at sonifying data. Sonifying the data using a constant pitch with variable loudness also led to annoyance caused by the unchanging pitch.

Other examples of sonifying solar data include solar oscillations,[28,29] solar wind data,[30,31] and an interactive image to music experience.[32] The first three examples are for 1-D time series while the fourth uses the motion of a person to sample an image. Others have produced a sonified solar system.[33] The image sonifications described herein may be one of the few examples of such a project.

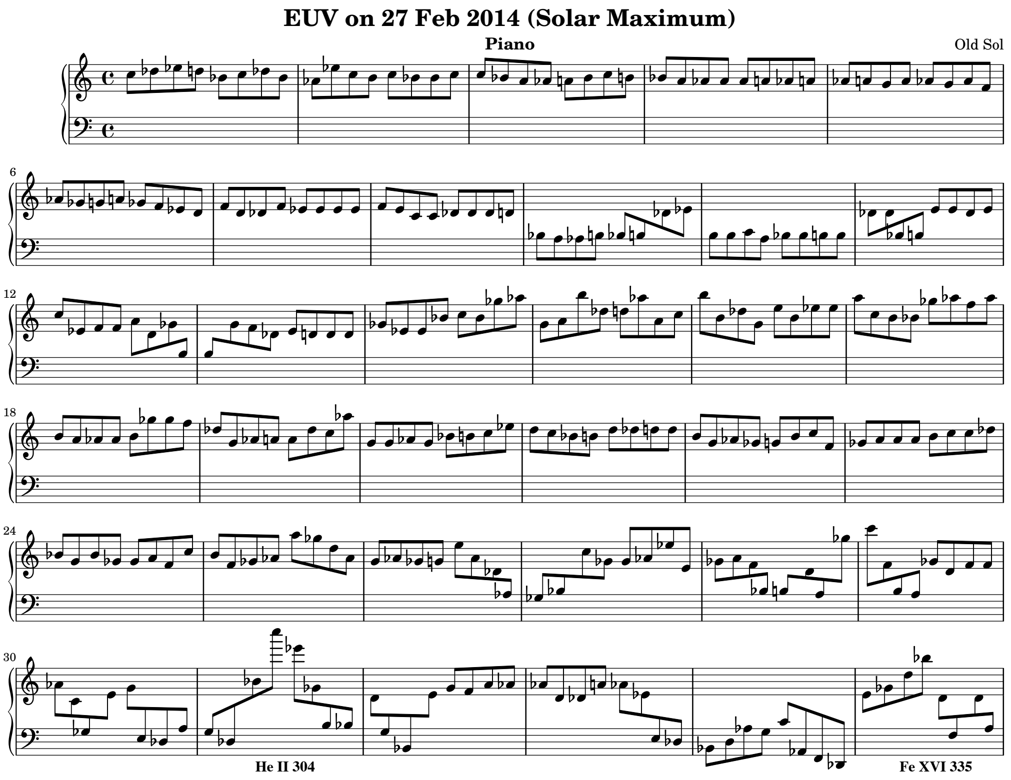

FIG. 9. The first page of a piano score of the EVE spectrum in Figure 2 created by Lilypond. The He II 304 Å line can be seen in measure 31 and the Fe XVI 335 Å line in measure 35. The scaling to pitch is different than the sonified example to better fit on the staves.

VII. CONCLUSIONS

We have sonified solar data as time series, an EUV spectrum, a time series of EUV spectral irradiances, EUV images with various techniques, and a time sequence of EUV images. The EUV spectrum showed that the independent variable does not have to be time. We demonstrated that using a Hilbert curve to address a solar image gives a sonification that shows more of the image variations and less of the shape of the Sun.

One shortcoming of the Hilbert curve sampling method is the separation of two regions near the limb. In these examples, images are sampled by a curve that crosses from the upper left quadrant to the upper right near the equator. This means the northern polar region is sampled in two distinct areas far from one another. The two lower quadrants do not have a direct connection and the southern polar region is also divided into two distinct regions, one at the beginning of the series and the other at the end. This can be remedied by rotating the Hilbert curve (or the image) 90° in either direction, which moves the connection between quadrants to the poles and keeps those regions in a smaller neighborhood while dividing the equatorial limb sectors into disparate parts of the sampling curve.

Other techniques can be used to sonify solar images. Coincident images observed in different wavelengths of light can be sampled and placed in different timbres or pan positions. Once the next solar maximum passes, another EVE spectrum could be used to play against the solar maximum spectrum illustrated here. Higher-order Hilbert curves can be constructed to sample a series of images. This would keep points within a neighborhood in both space and time. **Software that directly produces sounds rather than adhering to the MIDI standard might create sonifications that better represented the data. This could overcome the limited number of pitches available in the MIDI standard.**

Sonifying solar images is a way to explore the interface between tempo and pitch. Increasing the tempo to 3000 bpm (or 50 Hz) allows you to investigate whether an extremely rapid tempo results in an envelope with the individual pitches providing an amplitude modulation of that envelope. Frequencies of 15–30 Hz (900–1800 bpm) are near the limit of pitch discrimination.[34] The difference between the buzz saw of the raster scan image (Sec. III C) and the smoother sound of Hilbert curve sampling of Sec. V is one example of how the envelope makes a big difference in the perception of the data.

Listening to the Sun allows people to enjoy our closest star in a new direction. This does not apply only to the blind, most people can hear the variations of the Sun. With time these techniques will also allow people to more fully explore images as well.

VIII. QUESTIONS AND OTHER PROJECTS

Many projects can come from data-driven sonifications. There are also many ways to do those sonifications. We selected the JythonMusic synthesizer because we could load any data we wanted into the program. Once the MIDI exists it can be loaded into any compatible synthesizer for playback or experimenting.

Here are some suggestions that can motivate students to listen to their data:

1. **A simple way to sonify an image is to put the sampled pixels into a sound file, such as a WAV file and played at the CD sample rate of 44100 samples per second. This "audification" of an image does not require a synthesizer. The file can be opened in most media players and listened to. Compare the audified image with the sonified image and describe the differences, aside from the speed of the audified image.**
2. Can you find ways to vary the tempo of the music to represent variations in a data set? Scientific data tends to have even spacing and the simplest way to sonify the data is to maintain an even tempo. You can use the JythonMusic routine `Mod.tiePitches` to tie together identical notes to add some variety to the rhythmic spacing. Another routine, `Mod.accent` allows you to accent a beat, which also provides some rhythmic texture to the music.
3. Can loudness be used to emphasize important features in a log-scaled variable? Comparing the score of the spectrum in Figure 9 with the physical data in Figure 2, we can see that a few emission lines outshine much of that spectral region but that dominance is not reflected in the sonification. Perhaps increasing the loudness of the strong emission lines would better illustrate this dominance.
4. Three-color AIA images are created by putting coincident images in different wavelengths into individual color channels. These can be sonified by assigning a voice and pan position to each of the channels that will audibly emphasize the differences in the channels.
5. A wavelet analysis of a time series can be used to isolate persistent from ephemeral frequencies. Can a wavelet spectrum be sonified to show the persistent frequencies as droning notes and ephemeral events as more rapid variations?
6. Can other instruments be played against the synthesizer output? The sonified data has no explicit key, so improvised solos and rhythms can be played along with the sonified data.

ACKNOWLEDGMENTS

Version 4.6 of the JythonMusic software was downloaded from `https://jythonmusic.me`. All of

the data used in this research is available as continually updated files from publicly-accessible sites. The monthly averaged (SN_m_tot_V2.0.csv) and the annually averaged (SN_y_tot_V2.0.csv) International Sunspot Number (Version 2) data were obtained from the Solar Influences Data Center (http://sidc.oma.be/silso/datafiles). Daily averaged SEE measurements were obtained as the SEE Level 3 Merged NetCDF file at http://lasp.colorado.edu/data/timed_see/level3/latest_see_L3_merged.ncdf. Daily averaged EVE measurements were obtained the EVE Level 3 Merged NetCDF file at http://lasp.colorado.edu/eve/data_access/evewebdataproducts/merged/EVE_L3_merged_1a_2019135_006.ncdf. AIA images were obtained as JPEGs from the SDO website https://SDO.gsfc.nasa.gov.

---

[1] Sabina Teller Ratner. *Camille Saint-Saëns, 1835-1921: A Thematic Catalogue of His Complete Works*, volume 1. Oxford Univ. Press, New York, 2002. pp. 185–192.

[2] Dieter Daniels. Luigi russolo «intonarumori», 2020. URL http://www.medienkunstnetz.de/works/intonarumori/audio/1/.

[3] Iannis Xenakis. Electro-Acoustic Music. Vinyl LP, Nonesuch H-71246, 1970.

[4] Roger Luther. Moog Archives. URL http://moogarchives.com.

[5] Robert A. Helliwell. *Whistlers and Related Ionospheric Phenomena.* Dover Publications, Inc., 2006. Originally published by Stanford University Press, Stanford, California (1965).

[6] John H. Flowers. Thirteen years of reflection on auditory graphing: Promises, pitfalls, and potential new directions. In *Proceedings of ICAD 05-Eleventh Meeting of the International Conference on Auditory Display, Limerick, Ireland, July 6-9, 2005*, pages 406–409, 2005.

[7] The MIDI Association. Midi association, 2020. URL https://www.midi.org.

[8] Bill Manaris and Andrew R. Brown. *Making Music with Computers: Creative Programming in Python.* Taylor and Francis Group, LLC, Boca Raton, Florida, 2014.

[9] John Backus. *The Acoustical Foundations of Music.* W. W. Norton & Company, New York, 1969. p. 113.

[10] Frederik Temmermans. JPEG, 2020. URL https://jpeg.org/index.html.

[11] John Backus. *The Acoustical Foundations of Music.* W. W. Norton & Company, New York, 1969. The

Fletcher-Munson curves in this work have been revised and updated in ISO 226:2003 but the conclusions needed here remain valid.

[12] F. Clette, L. Svalgaard, J. M. Vaquero, and E. W. Cliver. Revisiting the Sunspot Number. A 400-Year Perspective on the Solar Cycle. *Space Sci. Rev.*, 186:35–103, December 2014. doi:10.1007/s11214-014-0074-2.

[13] Frédéric Clette and Laure Lefèvre. The new sunspot number: Assembling all corrections. *Solar Phys.*, 291:2629–2651, 2016. doi:10.1007/s11207-016-1014-y.

[14] T. N. Woods, F. G. Eparvier, R. Hock, A. R. Jones, D. Woodraska, D. Judge, L. Didkovsky, J. Lean, J. Mariska, H. Warren, D. McMullin, P. Chamberlin, G. Berthiaume, S. Bailey, T. Fuller-Rowell, J. Sojka, W. K. Tobiska, and R. Viereck. Extreme Ultraviolet Variability Experiment (EVE) on the Solar Dynamics Observatory (SDO): Overview of Science Objectives, Instrument Design, Data Products, and Model Developments. *Solar Phys.*, 275:115–143, January 2012. doi:10.1007/s11207-009-9487-6.

[15] W. D. Pesnell, B. J. Thompson, and P. C. Chamberlin. The Solar Dynamics Observatory (SDO). *Solar Phys.*, 275:3–15, January 2012. doi:10.1007/s11207-011-9841-3.

[16] T. Woods, S. Bailey, F. Eparvier, G. Lawrence, J. Lean, B. McClintock, R. Roble, G. Rottman, S. Solomon, and W. Tobiska. TIMED Solar EUV experiment. *Physics and Chemistry of the Earth C*, 25:393–396, 2000. doi:10.1016/S1464-1917(00)00040-4.

[17] J. R. Lemen, A. M. Title, D. J. Akin, P. F. Boerner, C. Chou, J. F. Drake, D. W. Duncan, C. G. Edwards, F. M. Friedlaender, G. F. Heyman, N. E. Hurlburt, N. L. Katz, G. D. Kushner, M. Levay, R. W. Lindgren, D. P. Mathur, E. L. McFeaters, S. Mitchell, R. A. Rehse, C. J. Schrijver, L. A. Springer, R. A. Stern, T. D. Tarbell, J.-P. Wuelser, C. J. Wolfson, C. Yanari, J. A. Bookbinder, P. N. Cheimets, D. Caldwell, E. E. Deluca, R. Gates, L. Golub, S. Park, W. A. Podgorski, R. I. Bush, P. H. Scherrer, M. A. Gummin, P. Smith, G. Auker, P. Jerram, P. Pool, R. Soufli, D. L. Windt, S. Beardsley, M. Clapp, J. Lang, and N. Waltham. The Atmospheric Imaging Assembly (AIA) on the Solar Dynamics Observatory (SDO). *Solar Phys.*, 275:17–40, January 2012. doi:10.1007/s11207-011-9776-8.

[18] John-Cage-Orgel-Kunst-Projekt. Organ^2/aslsp. URL https://www.aslsp.org/de/klangwechsel.html.

[19] W. D. Pence, L. Chiappetti, C. G. Page, R. A. Shaw, and E. Stobie. Definition of the Flexible Image Transport System (FITS), version 3.0. *A. & Ap.*, 524:A42, December 2010. doi:10.1051/0004-6361/201015362.

[20] D. Hilbert. Über die stetige abbildung einer linie auf ein flächenstück. *Mathematische Annalen*, 38: 459–460, 1891.

[21] G. Peano. Sur une courbe, qui remplit toute une aire plane. *Mathematische Annalen*, 36:157–160, 1890.

[22] John J. Bartholdi, Loren K. Platzman, R. Lee Collins, and William H. Warden. A minimal technology routing system for Meals on Wheels. *Interfaces*, 13(3):1–8, 1983.

[23] K. J. Vinoy, K. A. Jose, V. K. Varadan, and V. V. Varadan. Hilbert curve fractal antenna: A small resonant antenna for VHF/UHF applications. *Microwave Opt Technol Lett*, 29(4):215–219, 2001.

[24] Bernhard Seeger and Peter Widmayer. Geographic information systems. In Sartaj Sahni and Dinesh P. Mehta, editors, *Handbook of Data Structures and Applications*, chapter 56. CRC Press, Boca Raton, Florida, 2nd edition, 2018.

[25] G. M. Morton. A computer oriented geodetic data base; and a new technique in file sequencing. Technical report, IBM Ltd., Ottawa, Canada, 1966.

[26] Norbert Wiener. *The Fourier Integral and Certain of its Applications*. Dover, New York, 1933.

[27] Lilypond. Lilypond ... music notation for everyone, 2020. URL http://lilypond.org.

[28] Alexander G. Kosovichev. Solar Sounds, 1997. URL http://soi.stanford.edu/results/sounds.html.

[29] Tim Larson. SoSH Project: Sonification of Solar Harmonics, 2020. URL http://solar-center.stanford.edu/sosh/.

[30] Robert L. Alexander. Solar wind sonification, 2011. URL http://www.srl.caltech.edu/ACE/ACENews/ACENews144.html.

[31] Andrea effe Rao. Sounds from the Sun - Data Sonification - Thesis, 2016. URL https://www.behance.net/gallery/35831845/Sounds-from-the-Sun-Data-Sonification-Thesis.

[32] Marty Quinn. “Walk on the Sun”: An interactive image sonification exhibit. *AI & SOCIETY*, 27(2): 303–305, 2012. doi:10.1007/s00146-011-0355-1. URL https://doi.org/10.1007/s00146-011-0355-1.

[33] Michael Quinton, Iain McGregor, and David Benyon. Sonifying the solar system. In *The 22nd International Conference on Auditory Display (ICAD-2016)*, pages 28–35, 07 2016. doi:10.21785/icad2016.003.

[34] John Backus. *The Acoustical Foundations of Music*. W. W. Norton & Company, New York, 1969. Ch. 7.