

## INTRODUCING *strauss*: A FLEXIBLE SONIFICATION PYTHON PACKAGE

*James W. Trayford*

Institute of Cosmology and Gravitation  
University of Portsmouth  
Dennis Sciama Building  
Burnaby Road  
Portsmouth, United Kingdom  
james.trayford@port.ac.uk

*Chris M. Harrison*

School of Mathematics, Statistics & Physics  
Newcastle University  
Herschel Building  
Newcastle upon Tyne, United Kingdom  
christopher.harrison@newcastle.ac.uk

### ABSTRACT

We introduce *strauss* (Sonification Tools and Resources for Analysis Using Sound Synthesis) a modular, self-contained and flexible Python sonification package, operating in a free and open source (FOSS) capacity. *strauss* is intended to be a flexible tool suitable for both scientific data exploration and analysis as well as for producing sonifications that are suitable for public outreach and artistic contexts. We explain the motivations behind *strauss*, and how these lead to our design choices. We also describe the basic code structure and concepts. We then present output sonification examples, specifically: (1) multiple representations of univariate data (i.e., single data series) for data exploration; (2) how multi-variate data can be mapped onto sound to help interpret how those data variables are related and; (3) a full spatial audio example for immersive Virtual Reality. We summarise, alluding to some of the future functionality as *strauss* development accelerates.

### 1. INTRODUCTION

Sonification has been explored for the purposes of data exploration over the many decades of the ICAD community (e.g., see 2012 summary in [1]). The sonification tools that have emerged for this application come in many guises, including graphical user interfaces (e.g., *StarSound* [2], *Psst* [3]), code-based tools (e.g., *SoniPy* [4], *astronify*<sup>1</sup>), and recently online hosted tools (e.g., *Highcharts Sonification Studio*, [5], *SonoUno* [6]).

As we discuss in § 2, these tools have various goals in terms of intended users and applications. Sometimes they are domain specific, i.e., created for applying sonification within a specific scientific domain. The work presented here originated with our goal to apply sonification within the astronomy domain. Over recent years there has been an explosion of interest in developing approaches and tools for sonification in astronomy for research, science communication or education, and for improving accessibility to those who are blind or have low vision; see review in [7]. Astronomy is a domain particularly suited to exploring the

<sup>1</sup><https://astronify.readthedocs.io>



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/>

potential of sonification, due to: (1) the extremely large and often multi-dimensional datasets (e.g. hyper-spectral datacubes); (2) monitoring high-volume data for interesting and ‘transient’ events (e.g. for black holes colliding or supernovae explosions); (3) exploring time-series data for repeating features and patterns (e.g., planets transiting in front of stars[8]) and; (4) the general popularity of astronomy in the general public and media. As an example of the last point, NASA released a sonification of data from the gas around a black hole, which received international media attention and millions of hits on the Guardian’s YouTube channel<sup>2</sup> alone.

Here we introduce the design, structure and functionality of our new sonification package *strauss*<sup>3</sup> (Sonification Tools and Resources for Analysis Using Sound Synthesis). Although originally designed with astronomical data and communication in mind, we demonstrate how it has developed into a Python-based tool with utility for a wide range of applications, and is suitable for use across scientific domains.

### 2. STRAUSS MOTIVATION & PHILOSOPHY

To motivate the design choices of the *strauss* code, it is important to first motivate the needs it is aiming to fulfil, as well as placing it in the context of pre-existing tools and resources.

We posit that realising the potential of sonification for data-driven fields requires sonification techniques to become embedded into standard practice at the level of the *data analyst* - i.e. those handling and interpreting data. The data methodologies and modalities used by analysts set the perception and standards of their fields. Analysts can also represent a great range of experience levels; from students working on first projects, to world experts conducting groundbreaking research. Furthermore, the data analyst may wish to produce sonification both for their own data exploration and to communicate their research to a wide range of audiences: from the general public or school pupils through to other experts in their field. Having analysts familiar and comfortable with sonification, while thinking creatively about expression of their data through sound (beyond sonification as a passing novelty), may unlock sonification’s transformative benefits.

For familiarity and comfort to come about, we aimed to create an approach which makes everyday sonification practical, while flexible enough to handle diverse and creative applications. Furthermore, we aimed for the tool to cater to the broad experience

<sup>2</sup>[https://www.youtube.com/watch?v=\\_tXhBLg3Wng](https://www.youtube.com/watch?v=_tXhBLg3Wng)

<sup>3</sup><https://audiouniverse.org/research/strauss>

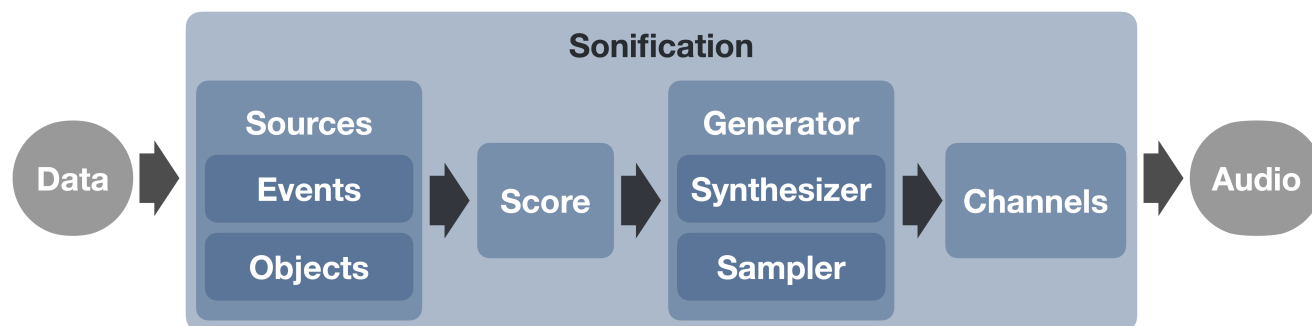


Figure 1: Schematic diagram illustrating the modular workflow of a `strauss` sonification; going all the way from the raw input data to the output audio representation. The well-defined *sub-modules* can be altered independently, to aid conceptual clarity. The remit of the modules is detailed further in § 4

and aptitude range of potential users.

We use the description “*low barrier, high ceiling*” for such a tool; accessible and understandable for a new user to quickly perform elementary tasks, while allowing detailed technical control for the expert. We set out to create a tool that could be *easily integrated* into the analyst’s workflow, i.e. is not disruptive to use. This could include incorporating sonification as an option into the software that analysts already use. This aims to make it as straightforward as possible for analysts to start sonifying data, especially for users who are less familiar with sonification or unaware of its potential. Finally, we aim for a sonification process that is *conceptually clear*, to avoid alienating newcomers curious about sonification. This can be achieved by elements of the tool being well partitioned and structured, with a *coherent design philosophy*. Together, these criteria motivate `strauss`.

In the landscape of existing tools, we find an unserved niche for `strauss` between specialised tools for single, specific sonification applications (e.g. pitch mapping data series, inverse spectrograms of images) and highly generalised audio tools (digital audio workstations, audio synthesis environments, etc). The former typically perform one task well but offer limited flexibility (and therefore creative potential), the latter enables much finer control and broader possibilities, but typically requires users to build their own procedures with data-audio interface and mapping, separate from their usual workflow. This can be thought of as a gap between the novice being introduced to sonification and the seasoned expert, creating their own audio experiments. `strauss` is intended to bridge this gap for the typical analyst. It aims to enable the data analyst to develop from a sonification newcomer to expert, without necessarily breaking out from their normal workflow, and fostering broader experimentation with and popularity of auditory display in data-driven fields. The general functionality of `strauss` is designed to extend far beyond pre-existing tools for analysts.

We also acknowledge tools that provide their own self-contained environment and interface, also enabling visually impaired (VI) users [3, 5, 6, 9]. These are crucial for accessibility but fill a different role to `strauss` - we intend `strauss` to primarily function to popularise and facilitate sonification widely among analysts using code-based workflows, as a means to improve accessibility. Meanwhile, we do also intend to improve the accessibility of `strauss` itself (see § 5) and encourage its incorporation to provide accessible functionality for existing software.

### 3. DESIGN CHOICES

For easy integration and wide applicability, we opted to make `strauss` into a Python package. This has many benefits. Python is a widely used coding language, boasting a market share of 44% among software developers and tied for most desired language to learn for existing developers in the *Stack Overflow 2022 User Survey* [10]. Python makes external libraries very easy to incorporate, installed via a package manager (e.g. `pip`) and incorporated into code using the `import` system. For less technical users, graphical user interface (GUI) programs written in Python could simply import such a package to add sonification functionality.

To consider the actual functionality of `strauss`, a useful analogy to draw is to *visualisation* and specifically to Python *plotting* packages. In Python, this is exemplified by libraries such as `matplotlib`[11] and `yt`[12]. With the broad functionality and low-level control they afford, these libraries can be used by analysts to produce anything from plots prioritising the conveyance of technical information, to images for outreach applications prioritising aesthetics. While we should be cautious of this analogy - visualisation tools do not necessarily map directly to sonification, and sonification may be enhanced by wholly new approaches - this offers a useful template for our design approach, and we make use of this analogy throughout.

Along with this template, the `strauss` *design philosophy* benefits from a clearly delineated, modular workflow illustrated in Fig 1. An *object-oriented paradigm* (OOP) allows us to represent these modules using Python classes, such that they are entirely self-contained, and can be operated on independently, while collected together in the overarching `Sonification` class. These constituent parts of `strauss` are explained fully in § 4.

### 4. CODE STRUCTURE & FUNCTIONALITY

Given the general context and motivation of the code, we now provide a brief overview of components of a `strauss` sonification, represented by modules of the `strauss` package. A flow diagram of the different modules is shown in Fig. 1. More extensive, evolving documentation for `strauss` is hosted alongside the code<sup>4</sup>. The online documentation should provide more tech-

<sup>4</sup><https://strauss.readthedocs.io>

Source parameter	Subject	Map	Evolve
<i>elevation</i>	Channels	✓	✓
<i>azimuth</i>	Channels	✓	✓
<i>volume</i>	Sources	✓	✓
<i>pitch</i>	Sources	✓	✗
<i>time</i>	Score	✓	✗
<i>cutoff</i>	Generator	✓	✓
<i>time_evo</i>	Score	✗	✓
<i>spectrum</i>	Generator	✓	✗
<i>pitch_shift</i>	Generator	✓	✓
<i>volume_envelope/A</i>	Generator	✓	✗
<i>volume_envelope/D</i>	Generator	✓	✗
<i>volume_envelope/S</i>	Generator	✓	✗
<i>volume_envelope/R</i>	Generator	✓	✗
<i>volume_lfo/freq</i>	Generator	✓	✗
<i>volume_lfo/freq_shift</i>	Generator	✓	✓
<i>volume_lfo/amount</i>	Generator	✓	✓
<i>pitch_lfo/freq</i>	Generator	✓	✗
<i>pitch_lfo/freq_shift</i>	Generator	✓	✓
<i>pitch_lfo/amount</i>	Generator	✓	✓

Table 1: Table of Sources parameters currently available for data mapping in *strauss*, and the subject class. ‘Mappable’ properties can be mapped to persistent values per Sources. ‘Evolvable’ parameters can also be parameterised to a data series, for continuous evolution via the Objects source type.

nical and practical detail on using *strauss* and its underlying algorithms, relative to the introductory overview provided here.

#### 4.1. The Sources module

In *strauss*, data is represented through instances of the Sources class. Particularly, Sources specify the mapping between numerical values of variables in the data and the expressive properties of sound in the output sonification. In the *strauss* paradigm, sound emanates from these Sources, allowing us to naturally represent spatial data in conjunction with the Channels class (§ 4.4). A *strauss* sonification can comprise multiple sources, with each source able to support many variable mappings from the data. This provides a flexible system for inputting data into *strauss*, with Sources parameters used internally throughout the code when generating the sonification. In *strauss*, the user can choose to assign their data to two different types of sources class, Events and Objects, detailed below. Example applications of each type are demonstrated in § 6.

Some of the key mappable properties are the direction to the source with respect to the listener (specified via *azimuth* and *elevation* angles) the base *pitch* of the source, the relative volume of the source and the relative time in the sonification that the source is triggered. A full list of these mappable parameters can be found in Table 1. The correspondence between the mapped data and source parameter value ranges can also be specified when setting up the Source, either as a percentile (0-100%) or in data units. By default, *strauss* maps the full range of the mapped data (from 0 to 100%) to the full Sources parameter range.

**Events:** The Events child class is used to represent sources defined by some *occurrence time* and singular event properties. These can be thought of as representing individual data *points*. Some examples of data well-fit to the Events class could be, supernovae exploding, lightning strikes or mutations in a genome. Table 1 lists the currently implemented ‘mappable’ prop-

erties for Events sources, which can be parameterised once per source (e.g. mapping a single *pitch* to the maximum brightness of a set of supernova events).

**Objects:** By contrast to Events, Objects *persist through time*, with properties that *evolve*. These can be thought of as a means to represent data *series* with a single source. Some example data fitting this description could be a galaxy evolving through cosmic time, a hurricane developing in the Atlantic or a stock price changing over a year. Objects support the same ‘mappable’ properties specified once per source in Table 1, but also a subset of ‘evolvable’ sources parameters that can be evolved throughout the sonification.

#### 4.2. The Score module

With the audio Sources defined, the Score class allows us to place ‘musical’ constraints on the sound they produce to represent the underlying data. The duration of the output sonification is also specified via the Score with the timeline of the the sonification scaled to fit this duration. *strauss* primarily supports a ‘chordal’ score, where a chord or sequence of chords (i.e. notes to be played simultaneously) can be specified. The ‘pitch’ Sources property mapping (Table 1) is then used to assign Sources to notes, such that each source articulates (or ‘arpeggiates’) a note in the chord. By default this is achieved by ‘binning’ the mapped pitch parameters onto each note in ascending order, such that an approximately equal number of sources represent each note in the chord<sup>5</sup>. Where sequences of chords are specified, the chords change through the sonification, such that the note articulated by each of the Sources also depends on the mapped *time* parameter.

Single note sequences can be parameterised as a special case, using chords with a single note or series of single notes. More ‘freeform’ or *atonal* sonifications can also be created by using the ‘pitch\_shift’ Sources property applied on top of the ‘pitch’ property, allowing the sound frequency of each source to instead be freely varied from a base note over a specified range.

The Sources class is relatively underdeveloped in the current incarnation of the *strauss* code, with plans to add many more options and functions. For example, chord or key changes could be another avenue to express *metadata*, such as epochal changes in datasets. The Sources class could also be a place to extend our presetting approach (described in §4.3) to improve *strauss* as a ‘low barrier, high ceiling’ tool, such that users without much musical experience could try collections of predefined Score configurations, with simple, evocative names, choosing their preference.

#### 4.3. The Generator module

Knowing how the Sources express the data, and what they will play via the Score, it is now possible to generate the sound. The Generator class takes instruction from the two prior classes and generates audio for each individual source. This can be achieved using either the Sampler or Synthesiser child classes (along with the Spectraliser special case), detailed below.

**Sampler:** This class generates audio by triggering pre-recorded audio *samples*. A directory of audio files is used to specify which sample to use for each note of the sampler. These sam-

<sup>5</sup>i.e. for a *n*-note chord, defining bins using *n* + 1 evenly spaced *percentiles* of the *pitch*-mapped data variable.

ples are loaded into the sampler and are interpolated to allow arbitrary pitch shifting. Samples can also be looped in a number of ways to allow notes to sustain perpetually. Samples can be used to generate sonification with a more ‘natural’ sound, using familiar instruments such as bells or strings.

**Synthesiser:** This class instead generates audio additively using mathematical functions via an arbitrary number of oscillators. The `strauss` synthesizer supports a number of oscillator forms (see preset example below). The `Synthesiser` can be typically used for less organic but more flexible sonification, without requiring audio samples to be collected or cleared for use.

**Spectraliser:** A special case of the `Synthesiser`, this generator synthesizes sound from an input spectrum, via an inverse Fast Fourier Transform (IFFT), with randomised phases. The user can specify the audible frequency range that the ‘spectralised’ audio is mapped over.

These generators share a number of modulation options to affect the generated audio. In particular low frequency oscillators (LFO) affecting *volume* and *pitch* can be used to add ‘vibrato’ or ‘tremolo’ to notes. `strauss` can also shape the volume *envelope* of a note, via *attack*, *decay*, *sustain* and *release* (ADSR) parameters. Some of these parameters are mappable or evolvable in `strauss` (Table 1). These are familiar features in digital audio, with more that could be added to later versions of `strauss` [13].

There are many potential options that can be varied to specify the sound qualities of the `Generator` for technical users. For less technical users, `strauss` uses a preset system, where sets of configuration options can be chosen from by simply specifying their names. Presets are written in YAML, a common and intuitive configuration file format, allowing nested parameters. In this way, it is envisaged that users and community members could develop their own presets for reuse or contribute them for posterity. An excerpt from the default `Synthesiser` preset can be seen below:

```
# preset name
name: "default"

# full description
description: >-
  Default preset for the synthesizer, using three saw
  wave oscillators, two of which are detuned slightly
  higher and lower respectively, with lower volumes. This
  gives a harmonically rich sound, suitable for filtering,
  with detuned unison saws removing some harshness.
# oscillator information
oscillators:
  # oscillator are denoted osc<n> with n=3 by default
  #
  # level: the intrinsic volume
  #
  # detune: the change in tuning as a percentage of the
  # input frequency
  #
  # form: the waveform, choose from:
  # ['saw', 'square', 'sine', 'tri', 'noise']
  #
  osc1:
    form: 'saw'
    level: 1.
    detune: 0.
    phase: 0
  osc2:
    form: 'saw'
    level: 0.5
    detune: -2.

...

# or 'tremolo'
volume_lfo:
```

```
use: off
wave: 'sine'
amount: 0.5
freq: 3
freq_shift: 0
phase: 'random'
A: 0.
D: 0.1
S: 1.
R: 0.
Ac: 0.
Dc: 0.
Rc: 0.
level: 1

# Master volume
volume: 1.

# Default pitch selection
pitch: 1.
```

The full version of this file is viewable via the `strauss` repository<sup>6</sup>.

#### 4.4. The Channels module

Once sound has been produced for each source, the final step is to mix the audio down into some multi-channel audio format. The `Channels` class essentially represents a bank of virtual microphones, with 3D antennae patterns, that each correspond to a channel in the output file. This allows us to naturally capture the full 3D sound field from a listening point in the `strauss` simulation, for physically accurate and realistic spatialisation. `strauss` allows fine control over how these channels are configured, but it is foreseen that most users will use a pre-configured standard *audio system*, such as ‘mono’, ‘stereo’, ‘5.1’, ‘7.1’, etc. Mixing audio signals between `Channels` happens on a source-by-source basis, where the direction to each of the sources  $j$  at a given instant  $t$  allows us to compute a coefficient  $A_j(t)$  from the antenna pattern of each channel,  $j$ , which modulates the signal,  $S$ , added to each channel,  $C$ , as  $C_j(t) = \sum_{i=0}^n A_j(t) S_i(t)$ .

These audio systems typically use a *cardioid* antenna, with the exception of the *mono* class, using a single omni-directional (spherical) pattern. A special class of *audio system* provided by `strauss` is the *ambisonic* systems. In this case, `strauss` uses a virtual microphone bank with the precise antenna pattern corresponding to spherical harmonics, up to the chosen order, in standard *ambiX* format [14]. In this way, `strauss` can natively provide ambisonic output to arbitrary order e.g. orders 1, 2, 3 yielding 4, 9 and 16 channels respectively. This is demonstrated in § 6.3.

#### 4.5. The Sonification module

The top-level `Sonification` class loads in all the above classes and produces the final sonification. Once `Sources`, `Score`, `Generator` and `Channels` classes are defined, the `Sonification` class is invoked. The `render()` method can then be run to produce the sonification. Finally, the `Sonification` class flushes the audio samples from the `Channels` class to an output audio file, or streamed audio for display in Python notebook format.

### 5. USING & DEVELOPING `strauss`

We aim to make using `strauss` as intuitive as possible, with a number of tutorial examples packaged with the code to build this

<sup>6</sup><https://github.com/james-trayford/strauss>

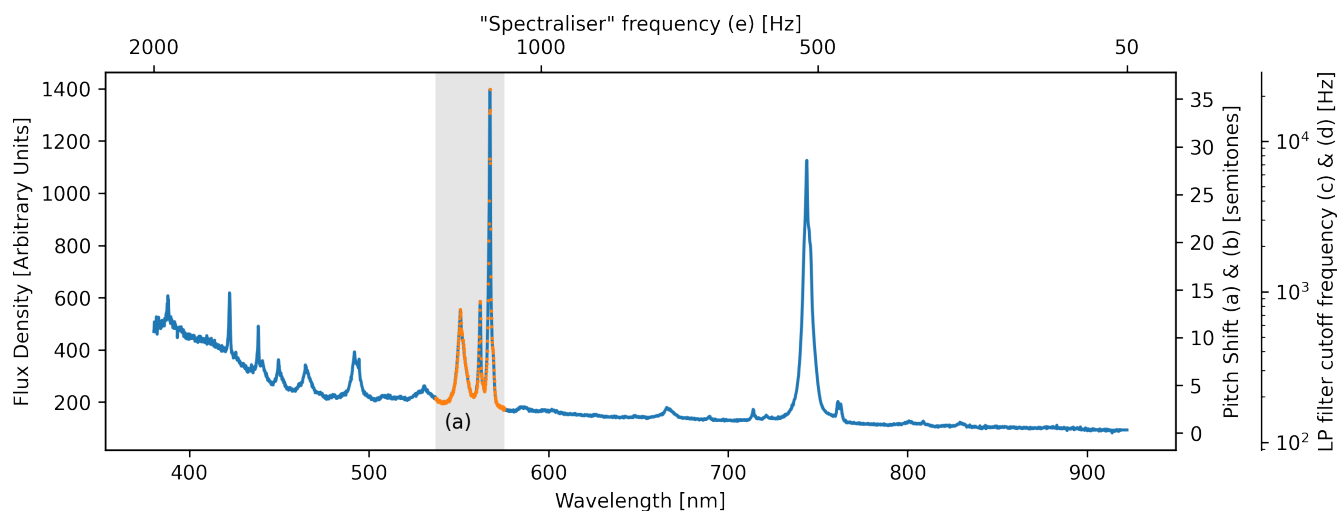


Figure 2: The example astronomical light spectrum of a galaxy (i.e., flux density, or brightness, as a function or wavelength), which has been passed through `strauss` to illustrate some of the multiple ways of sonifying *univariate data series*. The accompanying axes, to represent the different examples of mapping the data are labelled (a)-(e). The shaded region is the limited wavelength region used for example (a). Each of the methods are expanded upon in Section 6.1, where links are also provided to the video files containing both the audio and a visual representation.

intuition. With `strauss` targeted to introduce many of its users to its underlying concepts (i.e. those of sonification), clear tutorials are an essential element of the code development. We adopt a *Tutorial Driven Development* (TDD) paradigm; when adding new functionality we first design a model tutorial for its use and *then* implement changes in the code to ensure this works as designed. The reasoning behind this is that our intuition on how the code should work is preserved. Tutorials are available in the `examples/` subdirectory of the `strauss` repository, completing these should provide a broad introduction to the practical use of the code.

As we are aiming for a ‘*low-barrier, high-ceiling*’ tool in `strauss`, it is important to balance technical functionality with intuitive usability. One way we attempt this is via a ‘pre-setting’ approach, with pre-configured options or the constituent classes of `strauss`. This is exemplified by the `Generator` and `Channels` classes, where simple named presets can be chosen from for less technical users. The preset classes can also be modified to allow full control from a usable starting point. We plan to extend this to the `Score` and `Sources` classes, where common musical constraints and mapping choices could be chosen from. This could be considered analogous to a plotting package like `matplotlib`, where a ‘*colorbar*’ can be chosen from a pre-defined set, with evocative names (e.g. ‘*winter*’, ‘*magma*’, etc).

The question of *accessibility* for users the code is complex. In one sense the `jupyter` notebook environment aids accessibility, by lowering the barrier to use (allowing `strauss` and dependencies to be installed in-browser, using cloud resources via e.g. the `Google Colab` environment<sup>7</sup>) and simplifying elements of interface (allowing media to be embedded and played back directly in-notebook, allowing clearly formatted text via markdown). However, the `jupyter` environment has a number of issues with accessibility from a visually impaired (VI) standpoint; notebooks are not yet generically accessible with screen readers and other as-

sistive technologies [15]. For VI accessibility, future work will go into complementing `jupyter` tutorials with plain Python script counterparts (which are generically accessible), as well as adapting to the evolving accessibility status of `jupyter`.

## 6. EXAMPLE APPLICATIONS

For demonstrative purposes, we present some example outputs of `strauss`, showcasing some of the diverse possible applications and some of the key functionality. We include examples of: (1) a univariate data series using 5 different sonification approaches, applicable for research (a galaxy spectrum; § 6.1); (2) a multivariate data series using two different approaches, applicable for research (the variation of galaxy properties over time; § 6.2); and (3) an aesthetic spatial sonification application of a series of events for virtual reality (VR), applicable for outreach (the stars appearing in the night sky; § 6.3). All application are related to astronomy, but would be applicable to a variety of data with similar formats. While we comment on the utility of these examples, the intention is not to present these as the ‘*best*’ possible representations of the data or the ‘*proper*’ use of `strauss`. Instead, the idea is to convey the flexibility that `strauss` can provide, and encourage others to experiment with its features. All the media files associated with the figures and the corresponding sonifications in this work can be found at: [https://data.ncl.ac.uk/articles/media/Trayford\\_2023\\_STRAUSS\\_ICAD\\_examples/22241182](https://data.ncl.ac.uk/articles/media/Trayford_2023_STRAUSS_ICAD_examples/22241182), and we just list the relevant filenames below. a `README.txt` file also provides details of each file via the aforementioned link.

### 6.1. Research: Flexible Representation of Data Series

As a demonstration of the diverse possibilities with `strauss`, we demonstrate multiple audio representations of the single data series plotted in Fig 2. This the visible light spectrum of a galaxy

<sup>7</sup>see <https://www.audiouniverse.org/research/strauss>

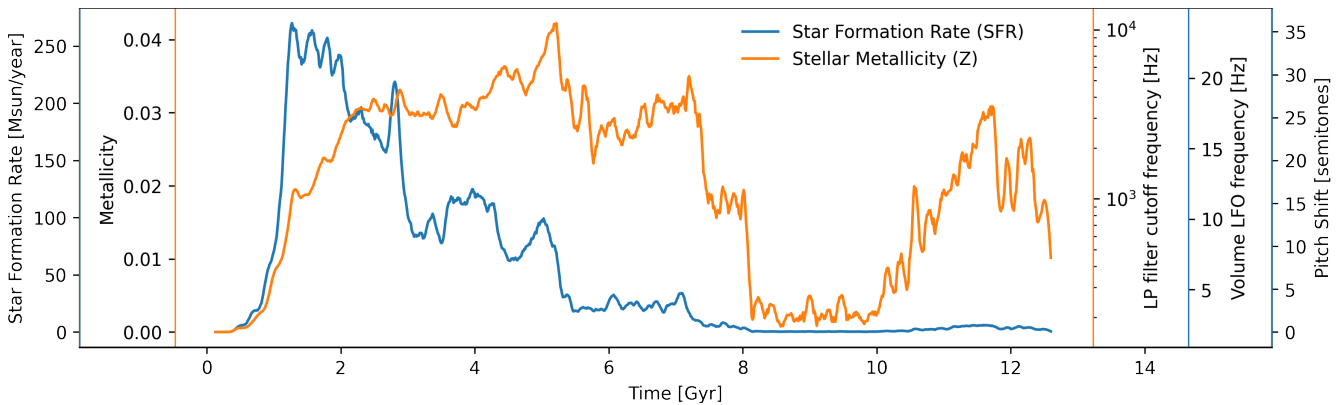


Figure 3: Example bivariate data series, including: (1) the star formation rate (SFR) and (2) the ‘metal’ mass fraction (known as ‘stellar metallicity’), of a simulated galaxy from the EAGLE simulations over 13 Gyr of cosmic time. The metallicity is mapped to a low-pass filter cut off frequency (see right axes), and we provide two different example mappings for simultaneously mapping the star formation rate to another property of the sound. These are pitch and volume LFO frequency, as represented in the axes and described in § 6.2 (where links are also provided to the video files containing both the audio and a visual representation).

taken from the Sloan Digital Sky Survey (SDSS). The data is available from <http://sdss.org/>, where the object identification number of this spectrum is: 1237671260667576428. These data, show measured ‘flux density’ (i.e., energy detected per unit time, per unit area, per unit wavelength) as a function of wavelength. The data exhibit a gradually varying *continuum*, as well as narrow *emission line* peaks, which are the result of particular chemical elements and ions in the galaxy’s gas. We may represent this in many ways with *strauss*. For method (a) we represent each measured data point, within a narrow wavelength range (shaded region in figure) as individual *Events*. We represent each data point using the packaged *mallets* samples, with an event *time* mapped to wavelength and *pitch* to flux density over a 3 octave range, from an A<sub>2</sub> (110 Hz) base note (specified via the *Score*). The resulting audio file, with an animated figure has been made available in a video mp4 format (*spectrum\_realisation\_a.mp4*).

An alternative representation can be heard in example (b), for which we can instead think of the spectrum as a single evolving *Object* (*spectrum\_realisation\_b.mp4*). Here, we use the *Synthesizer* type *Generator*, mapping the pitch shift of an A5 dyad (A<sub>2</sub>, E<sub>3</sub>) to flux, smoothly varying with time to represent the varying flux with increasing wavelength. This uses the built-in *pitch\_mapper* preset to configure the *Synthesizer*. Briefly, this is an unfiltered, single-oscillator triangle wave, with a pitch shift range of 0 to 36 semitones, and all modulation and enveloping turned off. This is an example of using presets to simplify the sonification workflow in *strauss*. The *Score* class can be used to specify any choice of notes here - an A5 dyad was chosen as more pleasant than a single note, while avoiding complex harmony that could affect the pitch perception used to convey the data.

We return to our analogy of common visual representations of data for these sonification approaches; Example (a) is like a *scatter-plot* with individuated data points (see orange points in shaded region of Fig. 2), while example (b) is more like a *continuous interpolation* of the data points (represented by the blue curve in Fig. 2).

We now present examples using the same continuous,

*Object* representation of (b), but mapping flux density to the cutoff frequency of low-pass filter. This mapping is shown with the right-most axis in Fig. 2. This approach effectively varies the ‘*timbre*’ of the sound (here meaning harmonic content) to represent changing flux density, with a ‘*brighter*’ sound corresponding to a *brighter* (i.e., higher flux density) data point. This allows us to pick the carrier signal we want without varying its fundamental pitch. For example (c) we use the A5 dyad again, using the default synthesizer preset (3 detuned sawtooth oscillators) for a *tonal* representation with fixed notes that could be used in a harmonious context (*spectrum\_realisation\_c.mp4*). Alternatively, for example (d) we employ a *textural* representation, by varying the filter over a white noise carrier, using the windy built-in preset of *strauss* (*spectrum\_realisation\_d.mp4*). We have found that this type of representation can be useful to generate more *listenable* sonification of data series over extended periods [8].

Finally, example (e) is a completely different approach of sound representation (*spectrum\_realisation\_e.mp4*). Here we map the spectrum shown in Fig. 2 directly to an audible spectrum, via an inverse Fourier transform. In this regard it may be considered a *audification* of the data and all the frequency features in the spectrum are heard ‘*simultaneously*’, with the spectrum represented as one single *Event*. For this example, we also use the *strauss* in-built enveloping to apply some *attack* and *decay* to the sonification, for a more organic sound. This affords unique opportunities to explore complex, contemporary datasets with sound, such as hyperspectral data cubes, much more rapidly than is possible by scrolling through independent variable, as above [16].

## 6.2. Research: Multivariate data series

The flexible mapping capabilities of *strauss* allow us to extend the univariate data series example presented in § 6.1 to *multidimensional* or *multivariate* data. For this example, we use data from the EAGLE simulations of galaxy formation [17] - using the public EAGLE database [18]. EAGLE simulates the formation and evolution of entire populations of galaxies through cosmic time within large, cubic volumes of the Universe and using model prescrip-

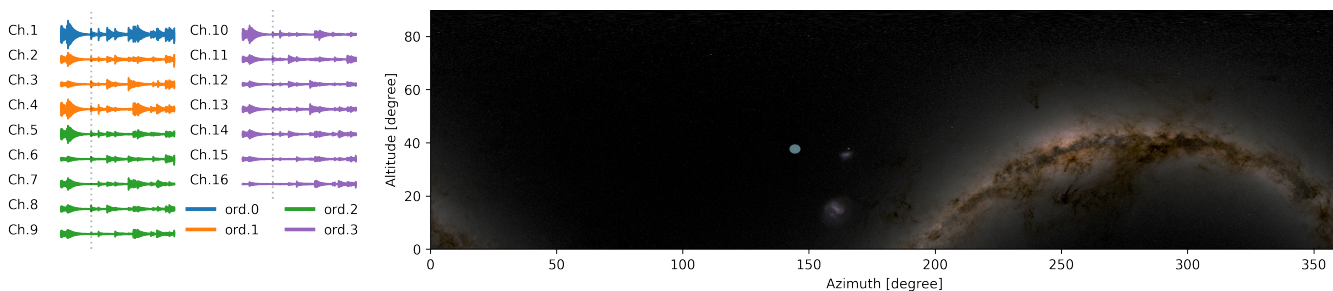


Figure 4: Illustration of an outreach application of *strauss*, creating spatial audio for the "stars appearing" virtual reality exhibit. *Left panel* shows wave-forms of the first 30s of audio in 3rd-order ambisonic format with *ambiX* [14] channel ordering. *Right panel* shows a single frame at 8s in equirectangular projection, with a single star appearing (filled circle). Dotted lines mark where this instant corresponds to a triggered sample in the corresponding audio on the left. Described in § 6.3, where links are also provided to the video files containing both the audio and a visual representation.

tions for unresolved physical processes, reproducing a number of key observation of the real universe.

We selected the galaxy with ID 21730536 from the largest volume simulation of EAGLE. We concentrate on two separate, but closely linked, data series that describe the star formation history of this galaxy: (1) the rate at which stars formed in the galaxy (SFR); and (2) the 'metal' mass fraction of the stars ( $Z$ ) inside the galaxy. The data are visually represented in Fig. 3. These two data series provide important insight independently, but how they vary with respect to *each other* also encodes important, complex information about how galaxies evolve over time<sup>8</sup>. Visual inspection of multiple variables at the same time, to gain insight into their relationships, is a ubiquitous problem in physical systems (typically dictated by differential equations). Therefore auditory inspection of these multiple and evolving properties is a natural avenue for consideration.

In *strauss* we can map multiple variables to different expressive properties of sound simultaneously for the same source. In Fig. 3 we treat the galaxy star formation history as a single Object source as in previous examples, with the default synthesizer preset and using an A5 dyad as carrier, following previous examples. We choose parameter mappings to have some intuitive link to the variables themselves, as indicated in the multiple  $y$ -axes of Fig. 3. We map increasing metallicity  $Z$  to the increasing low-pass filter cutoff frequency, yielding a duller sound at low  $Z$  and a more 'metallic' sound at high  $Z$  (compare *orange* axes). The star formation rate (SFR) can be thought of as a frequency at which stars form. Therefore, for realisation (a) we map increasing SFR to increasing pitch shift. As before, we produce the audio, along with an animated version of the figure in an mp4 format (*multivariate\_realisation\_a.mp4*). An alternative realisation (b), maps the SFR to a volume LFO frequency, such that more rapid pulses indicate a higher frequency of star formation (*multivariate\_realisation\_b.mp4*). *Blue* axes on the right hand side of Fig. 3 notate these SFR mappings. The efficacy of such multivariate mappings in research contexts requires evaluation, but is naturally possible with *strauss*. Such mappings can also be used in more impressionistic contexts, discussed next.

<sup>8</sup>The metals enriching each generation of stars mostly come from former generations dying off as supernovae, however these supernova can also drive gas and metals out of galaxies, and fresh gas can fall in from outside the galaxy. The balance of star formation and metal enrichment in a galaxy tells us about this *baryon cycle*.

### 6.3. Outreach: Aesthetic spatial sonification of stars in VR

In addition to research application, *strauss* can be used for more impressionistic sonifications or cases where aesthetics are a higher priority, for example when used for scientific outreach or artistic applications. An outreach application of *strauss* is exemplified by the 'stars appearing' VR sequence adapted from the *Audible Universe: Tour of the Solar System* planetarium show [19]. This example makes use of the native spatial audio output capabilities of *strauss*, as illustrated in Fig. 4. In this example, ~2500 stars are sonified as Events, mapping their brightness in the observed  $V$  band (i.e., using a camera filter which are centred around 550 nm) to occurrence time, with their base notes mapped to the stars perceived 'colour' measured as the difference in brightness of the stars in the  $B$  band (a filter centered around 450 nm) and the  $V$  bands. The intention here is to emulate the stars 'appearing' to a viewer over the course of a night. We see the brighter stars first, but as our eyes adjust, more and more dim stars appear to us. The colour mapping is also intended to be relatively natural, 'bluer' stars emitting more *high frequency* radiation are represented by *higher frequency* notes. We use the Sampler generator with Glockenspiel samples, and specify a musical choice of notes via the Score;  $D\flat_3$ ,  $G\flat_3$ ,  $A\flat_3$ ,  $E\flat_4$  &  $F_4$  - constituting a  $D\flat_6/9$  chord. *strauss* assigns an approximately equal number of stars to each note,  $D\flat_3$  is used for stars in the reddest quintile,  $F_4$  for the bluest quintile, etc. More details of the data and pitch mapping are provided in [19].

A key feature of this example is the full spatial audio. *strauss* is able to map the source of the sound accurately to their position on the sky (represented via the altitude and azimuth angles at the observing location) using an ambisonic audio representation. Fig. 4 shows the 16 channels representing 3<sup>rd</sup> order ambisonic audio in *ambiX* format - with the waveforms colour-coded by the ambisonic order they represent. Ambisonics can be naturally represented in *strauss* as a specific case of the virtual microphone system set up by the Channels class (see § 4.4), allowing an arbitrary number of sources to be represented as part of an accurately spatialised soundscape. Inspecting the waveforms in Fig. 4 shows how the audio triggered by a single star appearing contributes differently to each of the ambisonic channels. This example can be experienced in a variety of formats, including: (1) the full 3rd order audio included as a .wav file (*rendered\_stars\_ambix3.wav*); (2) a file mixed down to 8 channels for standard VR

playback support (e.g. via *Oculus TV*, *DeoVR*, etc; `stars_appearing_mw_ambix3_360_FB360.mkv`); and (3) a video that shows the first person view ('PoV'; though with sound mixed down to mono; `stars.mp4`)

## 7. SUMMARY & FUTURE WORK

We present the motivation and philosophy for the sonification Python package `strauss`, outlining the general structure of the code and demonstrating some example applications that showcase the diverse features in action. This is intended to be a broad introduction to `strauss` and its key concepts, acknowledging that the code is a work in progress, with plans for significant development. To complement this, the online `strauss` repository and documentation serve as dynamic documents, providing up-to-date technical details on the operation of the code, relevant examples, and listing planned features. Going forward, we are beginning a period of dedicated development of the code, which will allow many of the feature presented here to be developed further and many new features to be added, along with user interface and accessibility enhancements.

Some of the focus points for future development relate to the "low barrier, high ceiling" philosophy and accessibility features planned for `strauss`. In particular, the *presetting* (or *pre-configuration*) approach demonstrated in the `Generator` and `Channels` classes, could be extended to the remaining classes in the code, allowing users to choose from representative lists of (editable) configurations to try for their data, as a beginners' entry-point to using `strauss`. We also plan extensions for real-time audio rendering, and to make use of existing digital sound resources. Generally, `strauss` development will be responsive to the features and improvements that users desire, welcoming community contributions in the free, open source (FOSS) paradigm. We hope that making `strauss` available to the astronomy community, and data-centric fields more broadly, will lead innovative code applications and sonifications to naturally emerge.

## 8. ACKNOWLEDGEMENTS

JWT acknowledges support via the *STFC Early Stage Research & Development Grant*, reference ST/X004651/1. CMH acknowledges funding from an United Kingdom Research and Innovation grant (code: MR/V022830/1).

## 9. REFERENCES

- [1] N. Bearman and E. Brown, "Who's sonifying data and how are they doing it? a comparison of icad and other venues since 2009," 2012.
- [2] J. Cooke, W. Díaz-Merced, G. Foran, J. Hannam, and B. Garcia, "Exploring Data Sonification to Enable, Enhance, and Accelerate the Analysis of Big, Noisy, and Multi-Dimensional Data," in *Southern Horizons in Time-Domain Astronomy*, R. E. Griffin, Ed., vol. 339, Aug. 2019, pp. 251–256.
- [3] V. Potluri, J. Thompson, J. Devine, B. Lee, N. Morsi, P. De Halleux, S. Hodges, and J. Mankoff, "Psst: Enabling blind or visually impaired developers to author sonifications of streaming sensor data," in *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, 2022, pp. 1–13.
- [4] D. Worrall, M. Bylstra, S. Barrass, and R. Dean, "Sonipy: The design of an extendable software framework for sonification research and auditory display," in *Proc. ICAD*, 2007.
- [5] S. J. Cantrell, B. N. Walker, and Ø. Moseng, "Highcharts sonification studio: an online, open-source, extensible, and accessible data sonification tool." Georgia Institute of Technology, 2021.
- [6] G. De La Vega, L. M. Exequiel Dominguez, J. Casado, and B. García, "SonoUno web: an innovative user centred web interface," *arXiv e-prints*, p. arXiv:2302.00081, Jan. 2023.
- [7] A. Zanella, C. M. Harrison, S. Lenzi, J. Cooke, P. Damsma, and S. W. Fleming, "Sonification and sound design for astronomy research, education and public engagement," *Nature Astronomy*, vol. 6, pp. 1241–1248, Nov. 2022.
- [8] J. Tucker Brown, C. M. Harrison, A. Zanella, and J. Trayford, "Evaluating the efficacy of sonification for signal detection in univariate, evenly sampled light curves using ASTRONIFY," , vol. 516, no. 4, pp. 5674–5683, Nov. 2022.
- [9] J. Casado, B. García, P. Gandhi, and W. Díaz-Merced, "A New Approach to Sonification of Astrophysical Data: The User Centred Design of SonoUno," *arXiv e-prints*, p. arXiv:2302.00085, Jan. 2023.
- [10] "Stack overflow developer survey 2022." [Online]. Available: <https://survey.stackoverflow.co/2022/>
- [11] J. D. Hunter, "Matplotlib: A 2d graphics environment," *Computing in Science & Engineering*, vol. 9, no. 3, pp. 90–95, 2007.
- [12] M. J. Turk, B. D. Smith, J. S. Oishi, S. Skory, S. W. Skillman, T. Abel, and M. L. Norman, "yt: A Multi-code Analysis Toolkit for Astrophysical Simulation Data," *The Astrophysical Journal Supplement Series*, vol. 192, p. 9, Jan. 2011.
- [13] R. Boulanger, *The Csound book: perspectives in software synthesis, sound design, signal processing, and programming*. MIT press, 2000.
- [14] C. Nachbar, F. Zotter, E. Deleflie, and A. Sontacchi, "Ambix-a suggested ambisonics format," 07 2011.
- [15] "Jupyter accessibility." [Online]. Available: <https://jupyter-accessibility.readthedocs.io/en/latest/>
- [16] J. W. Trayford, C. M. Harrison, R. C. Hinz, M. Kavanagh Blatt, and A. Dougherty, S. snd Girdhar, "Inspecting spectra with sound: proof-of-concept & extension to datacubes," in prep.
- [17] J. Schaye, R. A. Crain, R. G. Bower, *et al.*, "The EAGLE project: simulating the evolution and assembly of galaxies and their environments," , vol. 446, no. 1, pp. 521–554, Jan. 2015.
- [18] S. McAlpine, J. C. Helly, M. Schaller, *et al.*, "The EAGLE simulations of galaxy formation: Public release of halo and galaxy catalogues," *Astronomy and Computing*, vol. 15, pp. 72–89, Apr. 2016.
- [19] C. Harrison, J. Trayford, L. Harrison, and N. Bonne, "Audio universe: tour of the solar system," *Astronomy Geophysics*, vol. 63, no. 2, pp. 2.38–2.40, 04 2022.