DIGITAL SOUND PRINT: SONIFYING ECOACOUSTIC AND SENSOR DATA FROM VATNAJÖKULL NATIONAL PARK

Miguel Angel Crozzoli

Intelligent Instruments Lab
University of Iceland
Reykjavik, Iceland
miguel@hi.is

Garrison Gerard

University of North Carolina
Pembroke, USA
garrison.gerard@uncp.edu

Thor Magnusson

Intelligent Instruments Lab
University of Iceland
Reykjavik, Iceland
thormagnusson@hi.is

ABSTRACT

The increasing volume of environmental sound data generated by Passive Acoustic Monitoring presents a challenge of scale in how to analyze and process the large volume of recordings collected. Generally, acoustic analysis is performed using visual and statistical representations. By doing so, the temporal structure of the acoustic data is not represented sonically. In this paper, we present the Digital Sound Print, a sonification experiment organized around the temporal structure of acoustic analysis from recordings in Vatnajökull National Park in Iceland. Additionally, we harmonized the recordings with weather data through imputation to speculatively extend the temporal range. This consolidated data was sonified in two layers: a synthetic texture driven by spectral analysis and sensor data as well as a granular synthesizer that reintroduces the original recordings based on event detection. Our initial experience suggests that listening to the digital sound print allows for a perception of the temporal unfolding of the acoustic environment, bringing a temporal perspective back to the listening experience.

1. INTRODUCTION

Natural environments emerge from complex ecological relations between biological, geological, and atmospheric processes. To understand and study these environments, multiple monitoring technologies are used to capture different dimensions of this complexity into data. On one hand, established remote sensing technologies such as satellites and field stations provide long-term continuous data about non-audible environmental variables. On the other hand, increasingly accessible passive acoustic monitoring (PAM) captures the acoustic dimension of these environments through recordings of biophony, geophony, and anthrophony [1]. Our work here, focuses on the sound representation of ecoacoustic data and its potential integration with sensor data through an experimental sonification design.

In ecoacoustics, PAM surveys are generally transformed into numerical and visual representation for analysis [2] [3]. Spectrograms are mostly used for visualization of this acoustic content, while machine learning algorithms classify and catalogue acoustic events such as species vocalizations [4] [5]. Furthermore, when the original sound is used within this pipeline, generally segments of audio linked to specific classification outputs are triggered, for instance the moment an identified bird call appears [6]. While this mode of accessing sound reveals specific events of interest, it does not represent the continuous temporal unfolding of the acoustic environment. Our question then was: how can we design a sonification system for ecoacoustic data that allows for the perception of its temporal structure, and what role do the actual collected sounds play in this design?

A sonic representation of environmental data that reveals meaningful temporal structures requires long-term data collections at regular intervals throughout multiple seasonal cycles. Even though PAM technology has become more accessible, it can still be difficult to record for extended periods of time, and some collections span months rather than multiple years, with exceptions such as wilding.radio's archival project [7]. Conversely, more established monitoring technology such as field stations provide continuous data at regular intervals across years. Thus, can we use these data to extend the ecoacoustic representation into periods where recordings are not available?

This paper proposes a digital sound print, an experimental sonification design to unfold the temporal structures of ecoacoustic data through a speculative approach for acoustic and sensor data harmonization and imputation. Developed through Sonifying the Arctic, an artistic research project based in Vatnajökull National Park, Iceland, the digital sound print constructs a sonic identity for environmental data points by harmonizing ecoacoustic and sensor data through temporal anchoring around solar cycles. Where acoustic data is not available, our experiment extends the sonic representation using machine learning-based imputation, generating possible acoustic profiles from sensor data. The following sections outline the context for this project, present our method to create the digital sound prints, and discuss insights from the design process.

2. CONTEXT

The Digital Sound Print was heavily informed by an approach to ecosystems and sound based in soundscape ecology. Soundscape ecology emerged from related disciplines such as acoustic ecology and landscape ecology; it is concerned primarily with the sounds of ecosystems and is situated within ecology rather than being seen as a complimentary discipline [1]. It studies the sounds of a landscape as a source of information about the ecological processes, patterns, and changes occurring within an environment. The specific kind of study that we most rely on for this project is

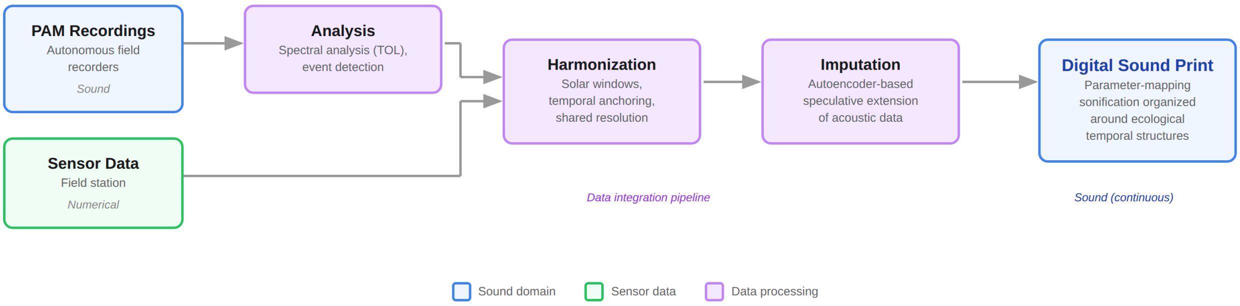

Figure 1: Digital Sound Print pipeline.

Passive Acoustic Monitoring (PAM), which use autonomous field recorders to capture recordings over long periods of time. PAM surveys have been used for a variety of purposes from establishing acoustic baselines to tracking biodiversity changes across an area [8].

Due to the sheer volume of recordings involved with PAM, most common analytical methods focus on visual and statistical representations such as spectrograms and sound event classification. For instance, the United States National Parks Service's process for determining the median loudness level without anthropogenic noise relies on human annotation of spectrograms to count aircraft and other motor noise [9]. The use of acoustic indices is becoming increasingly common for processing large amounts of recordings or even estimating biodiversity from PAM recordings [10]. These analysis tools allow researchers to distill large numbers of recordings down into a single data point that represents an acoustic parameter such as the ratio of anthropogenic to biophonic sound [11]. While incredibly useful, these indices have acknowledged flaws and are heavily dependent upon the researchers who use them having an understanding of what the acoustic index in question is actually calculating [12]. Within these methods, the temporal information/structure that sound carries as a perceptual source is absent. Within this context of extensive abstraction, we argue that there is space for the use of the recorded sound itself.

The sonification community has addressed environmental data sonification across scientific and artistic intents. Lindborg et al. [13] survey 32 climate data projects and notice that the main approach is to map environmental numerical variables to sound synthesis parameters, with only a few integrating recorded sounds into their design. Within this landscape, Zelada and Çamcı's "Unnatural Nature" [14] exemplifies this main approach, sonifying environmental variables such as temperature and CO2 through synthetic sounds that recreate a natural soundscape through artificial means without incorporating environmental recordings. Other projects integrate recorded sounds in various ways. Lindborg's "Stairway to Helheim" [15] combines recorded sound objects with sound synthesis in a sonification of 138 years of historical weather data. Elmquist et al. Birdsongification project [16] uses recorded bird songs as auditory icons to represent species in a biology visualization, with additional biological metadata conveyed through parameter mapping sonification. Ekeberg's "Ingenmannsland" [17] uses field recordings of Norwegian coastal rainforest as primary sonic material that is algorithmically processed and progressively replaced with synthetic sounds to reflect the scale of habitat loss. While these projects use sensor data and recorded sounds as part of their mapping or perceptual sound design, they do not use environmental sound data analysis to drive the temporal unfolding of the sonification.

Closer to our project, Spence and Ballora's Reef REM Ember piece [18] uses acoustic analysis of marine reef soundscapes as data input for sonification, with transient detection of the recordings driving percussion layers that represent the rhythmic complexity of the acoustic environment. The authors describe their process as sonifying sound and frame it as an indirect measurement of the soundscape's complexity. Their project is relevant to our work as it establishes a way of using acoustic analysis to generate sonification from environmental recordings following data's temporal unfolding. However, their analysis focuses on rhythmic structures rather than spectral content, and the sonification output consists of discrete percussive events rather than a continuous sonic representation as in this project.

Our proposed project builds on the context outlined above and addresses the question of how to use recorded environmental sound data analysis for a continuous sonification that follows the temporal unfolding of the acoustic environment. Through developing this experiment, we further explore how to combine our sound data with sensor data from a field station at the same location, and how to expand the sound data representation across periods where recordings are not available based on this combination. The following sections describe this experiment in detail.

3. DIGITAL SOUND PRINT

The digital sound print is a sonification experiment that explores the sonic representation of the acoustic state of an ecosystem. It takes as its input PAM recordings and extends the temporal coverage of the dataset with weather data, combining both sources into a sonic result that preserves a temporal structure that can be aurally explored faster than listening to the original recordings. Our digital sound print aims to construct a continuous sonic representation of environmental recordings that follows the temporal unfolding of the acoustic environment. This sound continuity means an uninterrupted, sustained, and evolving texture that sonifies the temporal aggregation of the acoustic data, allowing the listener to hear how acoustic structures unfold as they navigate through time. Spectral and event detection analyses of our site-specific recordings were used as our main data source for the sonification, and we further expanded this data by integrating sensor data from a field station at the same location. Both sources were combined using solar windows as temporal anchors and aggregating statistically within these windows towards the same temporal resolution. This harmonization revealed the temporal imbalance between both sources, with sensor data covering a significantly longer period than the acoustic recordings. This led us to implement an autoencoder-based imputation strategy to speculatively extend our sonic representation into periods where no sound data was recorded (see figure 1 for an overview of this pipeline). With the resulting dataset we created a parameter-based mapping sonification organized around ecological temporal structures.

3.1. Data Collection and Analysis

The data for this project combines two sources collected in Vatnajökull National Park in Iceland. The core acoustic material was collected with AudioMoth autonomous recording devices deployed across six field sites as part of a PAM survey tracking the impact of anthropogenic noise [19]. These devices captured one-minute recordings every five minutes from October 2023 to March 2024, accounting for roughly 4,000 hours of audio across all sites (see figure 2). Weather data was collected from the Solander's Eye field station [IK Foundation] at fifteen-minute intervals from September 2022 to April 2024, totalling over 53,000 sensor readings measuring temperature, wind, rainfall, humidity, barometric pressure, among other available measurements (see figure 3). This disparity in temporal coverage between approximately six months of acoustic data and nineteen months of continuous sensor data reflects the imbalance mentioned above.

The collected recordings were analyzed and then formatted in a variety of ways to inform our sonification design. We computed FFT analysis to determine the amplitude of individual frequencies across the human range of hearing, aggregated into third-octave level (TOL) spectrograms across 30 frequency bands from 25 Hz to 20,000 Hz, providing a spectral profile of each recording. This profile was used for resynthesis, creating a continuous and dynamic sounding layer in the digital sound print. In addition, the sounds of each field site were clustered based on the spectral density and amplitude deviation of their sound events. This analysis allowed for similar sounds to be quickly grouped and categorized by their acoustic characteristics using DBSCAN clustering algorithm [20], which in our project serves a dual purpose: first to identify and count event frequency across four acoustic categories (biophonic–principally bird sound, anthropogenic, rain, and wind sounds), and second to allow for quick retrieval of individual recordings for the sonification system.

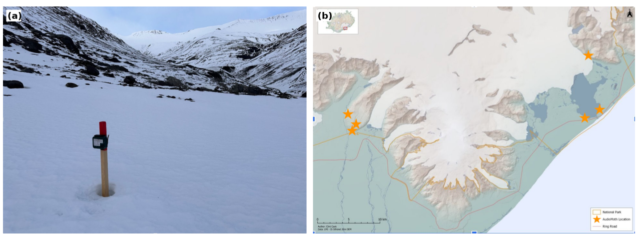

Figure 2: (a) AudioMoth autonomous recording device deployed at one of the field sites in Vatnajökull National Park. (b) Map of AudioMoth deployment locations across the study area.

Figure 3: The Solander's Eye field station with a mounted AudioMoth.

For the sensor data, we selected temperature, wind, rainfall, humidity, and barometric pressure as environmental variables from the thirty available from the field station. The selection was based on identifying environmental measurements that could give us symbolic meaning for mapping them to sound parameters, and variables that have low correlation with each other (see figure 4), disregarding the ones that were derived from or highly correlated with those already selected, such as heat index (correlation = 1.00 with temperature) and dew point, derived from temperature and humidity (correlation = 0.85). An exception was made for wind, which we accessed through two measurements, wind speed and wind run. Even though these are highly correlated, with speed representing instantaneous intensity and run the accumulated distance over a defined time interval, they carry distinct symbolic meanings that we mapped to different sound parameters (detailed in Section 3.4). These values were normalized before entering the harmonization and imputation pipeline.

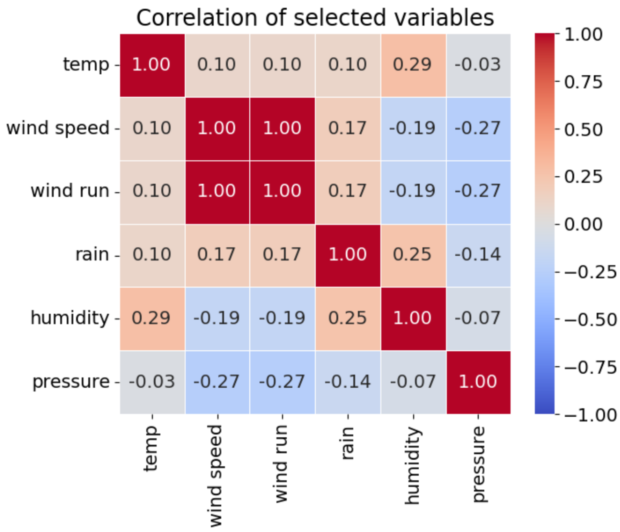

Figure 4: Sensor variable selection. Correlation matrix of the selected variables.

3.2. Data Harmonization

To combine sensor and acoustic analysis into a single shared data structure, we developed a harmonization pipeline aiming to create a common anchor to unify both sources. While these sources capture different dimensions of the environment through their technological affordances, they register the same ecological event, where weather and biophony are entangled [1]. By bringing them together we aimed at extending the environmental sound representation, with our acoustic analysis being our main focus for data-to-sound representation and sensor data expanding the perceptual relation to the environment by using the recorded sounds as its sonic material (sonification design explained in section 3.4). However, our two sources present two temporal challenges. They differ in sampling resolution, with sound being recorded for one minute every five minutes and sensor data registered every 15 minutes, and they further differ in temporal coverage, with sensor data spreading over 19 months and acoustic data covering 6 months, as mentioned in section 3.1. This harmonization pipeline solves the sampling resolution issue, while the coverage issue is addressed by the imputation strategy described in the following section.

To address the sampling resolution difference between our sources, we organized all data around sunrise and sunset times. The decision to use this solar framework draws on Farina et al. [21] observation that dawn and dusk transitions are ecologically significant periods where species choruses concentrate and acoustic activity shifts substantially. Sunrise and sunset times were computed using the Sunrise-Sunset API [1] for the Solander's Eye field station geolocation, and used to define four solar windows per day: dawn (sunrise time ±1.5 hours), dusk (sunset time ±1.5 hours), dawn-to-dusk (daytime), and dusk-to-dawn (nighttime). We set the 1.5-hour transition radius based on the extreme seasonal variation in daylight at this latitude, ranging from approximately 4.3 hours at winter solstice to over 21 hours at summer solstice (see figure 5). These extreme differences in sunrise and sunset times informed our experiment design heavily, as the windows shift dynamically across the year, with sunrise moving from 02:24 in June to 10:55 in December, reflecting the ecological rhythms of the recording sites in both weather patterns and biophony. A wider transition window would have collapsed the daytime period entirely; at ±2 hours, only 17 minutes of daytime would remain at the winter solstice.

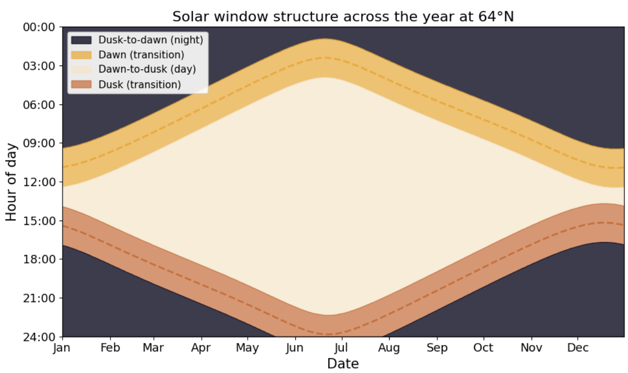

Figure 5: Solar window structure across the year at 64°N (Jökulsárlón).

[1] sunrise-sunset.org

The resulting dataset allocates a row for each window per day, representing an ecological transition throughout the daily cycle. Within each window, sensor data features were aggregated into statistical summaries, event detections were summarized as counts per category, and TOL spectrograms were averaged across each frequency band (see figure 6). While this harmonization confirmed the coverage imbalance mentioned above, it also created rows where both sources coexist, providing data pairs that capture the relationship between weather conditions and acoustic characteristics at the same temporal frame. These complete rows enabled the imputation strategy described in the following section.

TEMPORAL INDEX					SENSOR DATA (AGGREGATED)				…	TOL SPECTROGRAM (AVERAGED)				…	FLAG
Year	Month	Day	Window	Sol	Temp	Wind	Rain	Humid.		B1	B2	B15	B30		is_imputed
2022	9	22	dusk-dawn	03:00	6.84	5.72	0.41	18.43	…	22.14	23.60	25.56	29.59	…	True
2022	9	22	dawn	06:46	5.71	4.67	0.39	18.43	…	22.13	23.58	25.50	29.46	…	True
2022	9	22	dawn-dusk	09:00	10.02	5.83	2.93	18.66	…	22.18	23.61	25.56	29.54	…	True
…															
2024	4	16	dusk	20:46	2.77	0.50	0.83	18.49	…	22.15	23.60	25.57	29.54	…	False
2024	4	17	dawn	02:45	0.19	0.06	0.01	18.67	…	22.38	24.13	26.18	31.12	…	False
2024	4	17	dawn-dusk	05:00	1.44	0.33	0.83	18.46	…	22.16	23.64	25.60	29.65	…	False

2,130 rows × 80 columns. Green: sensor data aggregated per window. Blue: TOL spectrogram bands averaged per window. Yellow flag: imputed rows.

Figure 6: Excerpt of the harmonized dataset.

3.3. Data Imputation

As addressed in the preceding section, our dataset was constituted by acoustic data accounting for approximately a third of the total time series and sensor data spanning its totality. Based on this limitation, we considered either restricting the sonification to the six-month period where both sources were available, or developing a strategy for extending the acoustic representation across the full sensor time series. Given the experimental nature of this project, we decided to develop an imputation pipeline to generate the missing acoustic analysis from sensor data. This gave us the possibility to design for long-term time series, and to explore a method that could be tested and refined with other long-term PAM datasets in future experiments. Most importantly, the approach we applied preserves the original data, imputing only the missing values in existing rows. Being aware of these constraints, we approached this imputation as a speculative strategy, as the available acoustic data was not sufficient to reliably validate the imputed results against original observations.

To address this issue, we applied Macias et al. [22] autoencoder imputation method using an average code technique. This method trains an autoencoder on the subset of rows where all features are present, in our case the rows where both sensor and acoustic analysis data coexist, and uses the trained model to reconstruct missing values. The method creates multiple copies of each incomplete row and passes each copy through the encoder architecture, producing different compressed representations of the input data, which the authors mention as codes. These codes are averaged within the latent space of the encoder, and the resulting mean code is passed through the decoder to produce the final reconstruction. Once the complete row has been generated, only originally missing values are replaced in the dataset, with the original existing data preserved. The parallel with our harmonized environmental data is loose but addresses a similar problem, where the merging of acoustic analysis with sensor data left us with substantial gaps where acoustic recordings did not exist. In our project, we applied imputation only to TOL spectrogram values, not to event detection counts. This is due to the fact that event detections are identified from the original sound recordings through classification algorithms, making the relationship between sensor data and event counts too indirect for meaningful reconstruction.

The result of this imputation process extended the acoustic representation from the 6 months of collected acoustic data across the full 19-month sensor data span. In the imputed dataset, an is-imputed flag accompanies each row to maintain transparency about which rows contain only original values and which include generated data. To validate the reconstruction quality, we held out 30 percent of complete rows per site and randomly masked 30 percent of their values, with R^2 scores ranging from 0.66 to 0.92 across all harmonized sites (see figure 7). Thus, we acknowledge that this imputed data constitutes a conjectural extension of our sound representation rather than a validated prediction of acoustic analysis. The optimization of imputation methods for ecoacoustic data remains an open area for future work out of the scope of this experiment. The following section describes how both original and imputed data feed the sonification design.

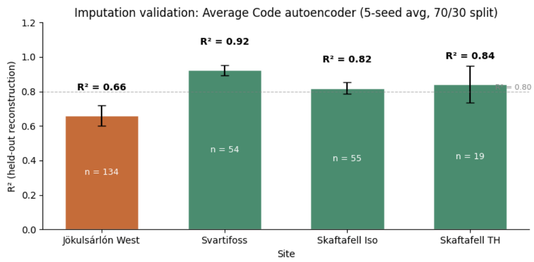

Figure 7: Imputation validation results across four harmonized sites.

3.4. Sonification Design

The sonification design for our digital sound print follows a parameter-based mapping approach, where data features are mapped to sound synthesis parameters through scaling functions [23]. For instance, rain event counts are scaled to control the playback quality of the other sounds in the granular synthesis engine. Our data-to-sound representation was designed in two layers to offer perceptually distinct streams. The first layer is a continuous sound representation driven by spectral analysis, aiming to follow the temporal unfolding of the acoustic environment through a sustained drone-like texture. The second layer uses granular synthesis to trigger sounds from the PAM recordings, thereby forming a more concrete, less abstracted connection to the recorded environment. The result is a two-layered sonification where acoustic analysis produces synthetic sound and event detection reintroduces the original recordings, with weather data mapped to parameters that shape how perceivable these elements are within the overall mix.

The first layer is split into two areas based on frequency, with the lower register covering approximately 25 Hz to 631 Hz (TOL bands 1 through 15) and the upper register covering 794 Hz to 20,000 Hz (bands 16 through 30). In the lower register, the 15 frequency bands are grouped into four overlapping bins, each assigned to a sine wave voice. Within each bin, the system selects the frequency with the highest amplitude and plays a sine tone at that pitch and level, producing four voices that move independently through their respective frequency ranges. For instance, if the first bin contains bands at 25, 32, 40, 50 and 63 Hz, and the 40 Hz band has the highest amplitude in a given solar window, the sine voice assigned to that bin plays at 40 Hz at the corresponding amplitude level. The bins slightly overlap at their borders, so when a dominant sound spans across two bins, voices can converge into fewer pitches, occasionally producing three or two tones instead of four. In the upper register, white noise is passed through a bank of filters corresponding to each of the remaining TOL bands, with each filter's gain set by the amplitude of its corresponding frequency band. This produces a textured layer that reflects the spectral density of the recorded environment in the higher frequencies (see figure 8). The sound output from both registers is further shaped by weather data that affects the clarity, brightness, and presence of this layer. These mappings draw on both acoustic phenomena observed in the physical environment and symbolic associations between an environmental variable and a sonic quality. Wind speed inversely affects the amplitude of the four sine voices, reflecting how wind masks sound transmission in the real environment, so that stronger wind reduces the presence of the tonal voices. Humidity controls the range of pitch detuning across the four voices, drawing on how moisture affects sound propagation through the air. Temperature was mapped to the brightness of the noise layer, with higher temperatures producing a brighter filter, a symbolic association treating warmth as a brightening of the spectral texture. Finally, wind run was mapped to the particle density of a dust-like texture in the upper register, translating the accumulation of wind over time into a granular sonic accumulation. The complete mapping between data features and sound parameters for this layer can be seen in figure 9. The perceptual result is a continuously evolving sound texture where a choral movement in the lower register is enriched by spectral density in the upper register.

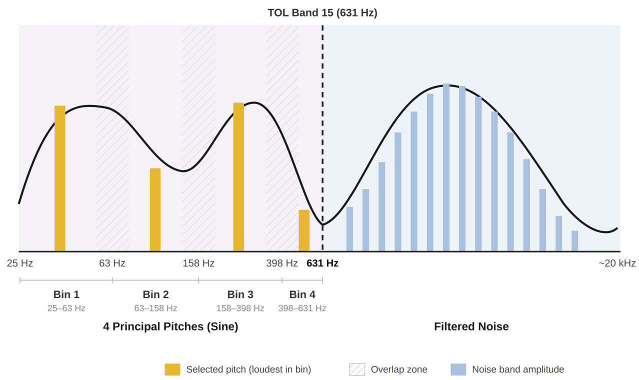

Figure 8: Layer 1: spectral synthesis mapped to sine voices (lower register) and filtered noise (upper register)

The second layer uses granular synthesis with categorized PAM recordings as source material, organized into three voices for birds/biophonic, motor/anthropogenic, and rain/geophonic, where each voice uses sound grains from its respective category. The amplitude of each voice is controlled directly by its event detection count, so periods with more bird activity produce louder bird grains, and similarly for motor and rain. The grain size for bird and motor voice, however, follows an inverse relation, meaning that as motor event counts increase, bird grain size decreases, and vice versa. This design decision is informed by the ecological observation that the presence of anthropogenic activity correlates with a reduction in biological sound activity [24]. Conversely, rain grain

LAYER 1: BACKGROUND SOUNDPRINT

Acoustic analysis (TOL) and sensor data mapped to synthesis parameters

	DATA FEATURES				
	TOL Spectral Analysis (acoustic)	Wind Speed (sensor)	Humidity (sensor)	Temperature (sensor)	Wind Run (sensor)
LOWER REGISTER (TOL BANDS 1–15, 25–631 HZ): 4 SINE WAVE VOICES ACROSS 4 FREQUENCY BINS					
Pitch (×4 voices)	Loudest frequency in each bin sets voice pitch	—	—	—	—
Amplitude (×4 voices)	Amplitude of selected frequency sets voice level	More wind = quieter	—	—	—
Detuning (×4 voices)	—	—	Higher humidity = wider detuning range	—	—
UPPER REGISTER (TOL BANDS 16–30, 794–20,000 HZ): FILTERED WHITE NOISE					
Noise dynamics	Each band's amplitude filters noise at that frequency	—	—	—	—
Noise brightness	—	—	—	Higher temp = brighter filter	—
Dust wind texture	—	—	—	—	More wind run = more dust particle density

Direct mapping Inverse mapping

Figure 9: Layer 1 mapping matrix.

size is controlled by the combined count of rain and wind events, as both contribute to the environmental noise texture captured in the recordings. Moreover, rain count works as a distortion effect where more rain maps onto a more atomised texture. This decision was informed by listening to the recordings, where rain creates a denser and noisier texture; slowing playback when rain is absent produces a more spacious sound output. The complete mapping between data features and sound parameters for this layer can be seen in figure 10. The perceptual result of this layer is a direct reference to the recorded environment, connecting the listener back to the acoustic environment the data is representing.

LAYER 2: GRANULATION LAYER

Event detection counts mapped to granular synthesis of field recordings

	DATA FEATURES (ACOUSTIC EVENT DETECTION)			
	Bird Count	Motor Count	Rain Count	Wind Count
GRAIN SIZE PER VOICE (SOURCE: CORRESPONDING FIELD RECORDINGS)				
Bird grain size	—	More motors = smaller grains	—	—
Motor grain size	More birds = smaller grains	—	—	—
Rain grain size	—	—	Rain + Wind combined	Rain + Wind combined
VOLUME PER VOICE				
Bird voice volume	Bird count = bird volume	—	—	—
Motor voice volume	—	Motor count = motor volume	—	—
Rain voice volume	—	—	Rain count = rain volume	—
SHARED GRANULAR PARAMETERS (ALL VOICES)				
Distortion (all voices)	—	—	More rain = more distortion	—

Direct mapping Inverse mapping Combined mapping

Figure 10: Layer 2 mapping matrix.

4. INSIGHTS AND LIMITATIONS

Using the digital sound print pipeline described in the previous section, we created multiple display realizations of our sonification. For each realization, the entire pipeline was applied independently to each of the six recording sites, producing different sonifications for the same time periods. In one sound installation, we mapped each field recording to one of six speakers distributed around the venue, allowing listeners to traverse the soundscape and hear the differences between sites (see figure 11). We also developed an interactive system that allows a listener to move between the six recording locations and move through the temporal dimension of the dataset through a map of the study area and a time slider. In this way, the listener can move to any point and hear the evolution of the sound print at their own pace, allowing for comparison across wide time spans and locations (see figure 12). We also produced fixed renderings that layer the site recordings into a path through time and space, trading the responsiveness of the live approaches for an easily transmissible form.[2]

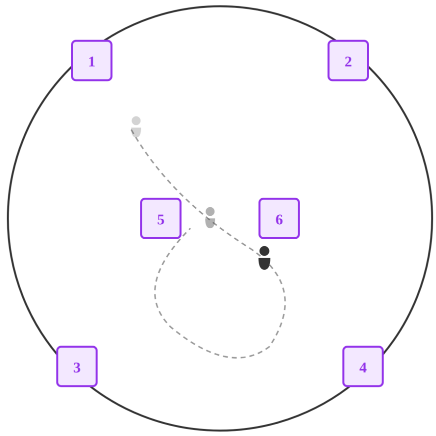

Figure 11: Sound installation layout.

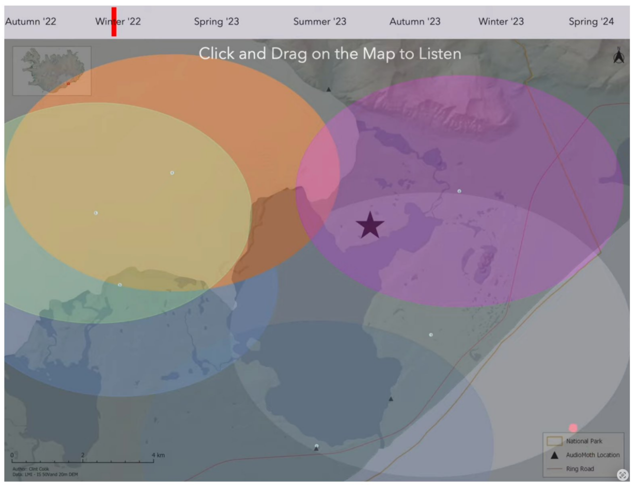

Figure 12: Interactive display.

Initial explorations suggest that the Digital Sound Print sound representation carries information about both dominant sounds and noise level from the acoustic environment, represented by dominant pitches and volume in the background sound layer. By moving through the sound representation across time, it is possible to hear long-term changes in soundscapes, and by selecting different dates to contextually compare them, the system allows for a study of the transformation of ecosystems across a variety of time scales. Moreover, even though the data imputation remains speculative, the combination of harmonized weather and acoustic data

[2]Examples of interactive display and fixed realization can be accessed at `http://tiny.cc/w2a4101`.

shows an interesting potential. If the imputation method can be refined and validated with more diverse training data, this approach could open the possibility of studying the soundscape of ecosystems wherever weather data is collected alongside a sufficiently long set of PAM recordings. In its current iteration, this is a limited instrument that does not provide a high degree of fidelity, and more work is needed to refine the approach, but the possibilities are promising for the ability of this system to represent the impact of climate change on soundscapes. Furthermore, the solar window structure proved particularly useful for observing how dawn and dusk choruses and their spectral structure change, exposing a temporal pattern through the day and across seasons of ecological significance. The application of our digital sound print pipeline across multiple sites further informed this perception, as listening to the same time period across different locations revealed differences based on site-specific acoustic characteristics.

Our experiment also faced multiple limitations. The acoustic data available for training the imputation model was concentrated in the winter period, providing scarce and in some locations highly variable sound material. Jökulsárlón West, an exposed glacial lagoon site with significant wind variability, produced the lowest reconstruction accuracy ($R^2 = 0.65$) despite having the most training rows, while sites with more homogeneous conditions achieved stronger results with fewer examples. Future deployments should consider recorder placement in relation to environmental variability and data collection across seasons. Our first impression for future implementation of this method is that it is better suited for filling scattered gaps in a dataset than for reconstructing entirely absent seasonal cycles. Moreover, the solar window design uses a fixed ±1.5 hour transition radius that does not account for the ecological difference between summer and winter dawn transitions; a mobile window length proportional to day length could better represent the significance of each transition. In the sonification design, weather data was not used to feed the granular layer, unlike the first layer where it shapes timbral parameters such as amplitude, detuning, and brightness. Applying a similar strategy to the granular layer could deliver a more cohesive experience across both layers.

We would further like to acknowledge that the data imputation process introduces uncertainty into the consolidated dataset, and thus into the sonic representation. How to perceptually map uncertainty so that it is rendered audible has been addressed by Ballatore et al. [25], who identify loudness, order, and clarity as the most intuitive sound dimensions for this purpose, and Vriend et al. [26], who examine how audio and visual channels can be paired to enhance the perception of uncertainty, among others. Even though this dimension was not central to our experiment, our imputation process produced an "is imputed" flag indicating whether a row's values were generated, supporting transparency in the data structure. This flag could drive sound parameters that render audible the difference between collected and imputed data in future iterations of this experiment.

5. CONCLUSION

In this paper we presented the Digital Sound Print, a sonification experiment that explores the temporal dimension of environmental recordings captured at Vatnajökull National Park, Iceland, by sonifying acoustic analysis derived from them. At the same area as our recording sites, a field station collects publicly accessible weather data, which motivated us to extend the sound representation with sensor data as it captures a different dimension of the same environment. However, these sources differ in sampling resolution and temporal coverage, with sensor data extending over nineteen months and acoustic recordings covering only six. We developed a harmonization pipeline that consolidates both sources into a unified dataset organized around solar windows, resolving the resolution difference. To address the coverage imbalance, we applied an autoencoder-based imputation to speculatively extend the acoustic representation into periods where recordings were unavailable.

The consolidated dataset was sonified in two perceptual layers: a continuous synthetic texture that follows the temporal unfolding of the data organized around ecological transitions from dawn to dusk, and a granular layer that uses the source recordings as sonic material, creating a perceptual connection between the listener and the environment they are exploring through the sonification.

Our initial experience from developing and exploring this sonification suggests that the digital sound print allows for perceiving the temporal development of acoustic analysis organized around ecological cycles. This project navigated multiple tensions: between the limitations of recording technology and the goal to sonify large acoustic time series, between scientific accuracy and aesthetic decisions in sound design, and between the speculative data imputation process and the meaning of our continuous data-to-sound representation. Finally, as PAM surveys grow in scale and more environmental monitoring data becomes publicly available, we see a need for approaches that display the temporal unfolding of this data around ecologically meaningful cycles and that combine diverse sources to represent the complexities of dynamic environments through listening experiences. The Digital Sound Print aims to contribute to this vision.

6. ACKNOWLEDGMENT

The authors would like to express their gratitude to the Hornafjörður Research Centre and the Intelligent Instrument Lab at the University of Iceland for their support and resources. Further thanks to the Fulbright Commission Iceland and Vatnajökulsþjóðgarður (Vatnajökull National Park) for their support. The Intelligent Instruments Lab project (INTENT), funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 101001848).

7. REFERENCES

[1] B. C. Pijanowski, L. J. Villanueva-Rivera, S. L. Dumyahn, A. Farina, B. L. Krause, B. M. Napoletano, S. H. Gage, and N. Pieretti, "Soundscape ecology: The science of sound in the landscape," *BioScience*, vol. 61, no. 3, p. 203–216, Mar. 2011. [Online]. Available: http://dx.doi.org/10.1525/bio.2011.61.3.6

[2] T. Bradfer-Lawrence, N. Gardner, L. Bunnefeld, N. Bunnefeld, S. G. Willis, and D. H. Dent, "Guidelines for the use of acoustic indices in environmental research," *Methods in Ecology and Evolution*, vol. 10, no. 10, p. 1796–1807, July 2019. [Online]. Available: http://dx.doi.org/10.1111/2041-210X.13254

[3] N. D. Merchant, K. M. Fristrup, M. P. Johnson, P. L. Tyack, M. J. Witt, P. Blondel, and S. E. Parks, "Measuring acoustic habitats," *Methods in Ecology and Evolution*,

vol. 6, no. 3, p. 257–265, Jan. 2015. [Online]. Available: http://dx.doi.org/10.1111/2041-210X.12330

[4] D. Stowell, "Computational bioacoustics with deep learning: a review and roadmap," *PeerJ*, vol. 10, p. e13152, Mar. 2022. [Online]. Available: http://dx.doi.org/10.7717/peerj.13152

[5] E. C. Knight, S. Poo Hernandez, E. M. Bayne, V. Bulitko, and B. V. Tucker, "Pre-processing spectrogram parameters improve the accuracy of bioacoustic classification using convolutional neural networks," *Bioacoustics*, vol. 29, no. 3, p. 337–355, Apr. 2019. [Online]. Available: http://dx.doi.org/10.1080/09524622.2019.1606734

[6] T. M. Aide, C. Corrada-Bravo, M. Campos-Cerqueira, C. Milan, G. Vega, and R. Alvarez, "Real-time bioacoustics monitoring and automated species identification," *PeerJ*, vol. 1, p. e103, July 2013. [Online]. Available: http://dx.doi.org/10.7717/peerj.103

[7] https://wilding.radio.

[8] S. R. P. Ross, D. P. O'Connell, J. L. Deichmann, C. Desjonquères, A. Gasc, J. N. Phillips, S. S. Sethi, C. M. Wood, and Z. Burivalova, "Passive acoustic monitoring provides a fresh perspective on fundamental ecological questions," *Functional Ecology*, vol. 37, no. 4, p. 959–975, Feb. 2023. [Online]. Available: http://dx.doi.org/10.1111/1365-2435.14275

[9] D. Betchkal, "Acoustic monitoring report, Denali National Park and Preserve – 2012," Natural Resource Data Series NPS/DENA/NRDS—2013/589, National Park Service, Fort Collins, Colorado, 2013.

[10] I. Alcocer, H. Lima, L. S. M. Sugai, and D. Llusia, "Acoustic indices as proxies for biodiversity: a meta-analysis," *Biological Reviews*, vol. 97, no. 6, p. 2209–2236, Aug. 2022. [Online]. Available: http://dx.doi.org/10.1111/brv.12890

[11] E. P. Kasten, S. H. Gage, J. Fox, and W. Joo, "The remote environmental assessment laboratory's acoustic library: An archive for studying soundscape ecology," *Ecological Informatics*, vol. 12, p. 50–67, Nov. 2012. [Online]. Available: http://dx.doi.org/10.1016/j.ecoinf.2012.08.001

[12] T. Bradfer-Lawrence, C. Desjonqueres, A. Eldridge, A. Johnston, and O. Metcalf, "Using acoustic indices in ecology: Guidance on study design, analyses and interpretation," *Methods in Ecology and Evolution*, vol. 14, no. 9, p. 2192–2204, Aug. 2023. [Online]. Available: http://dx.doi.org/10.1111/2041-210X.14194

[13] P. Lindborg, S. Lenzi, and M. Chen, "Climate data sonification and visualization: An analysis of topics, aesthetics, and characteristics in 32 recent projects," *Frontiers in Psychology*, vol. 13, Jan. 2023. [Online]. Available: http://dx.doi.org/10.3389/fpsyg.2022.1020102

[14] E. Zelada and A. undefinedamcı, "Unnatural nature: Spatial sonification of climate change data," in *Proceedings of the 28th International Conference on Auditory Display (ICAD2023)*, ser. ICAD 2023. International Community for Auditory Display, June 2023, p. 121–128. [Online]. Available: http://dx.doi.org/10.21785/icad2023.9190

[15] P. Lindborg, "Stairway to helheim," in *Proc. of DACA (Data Art for Climate Action)*, Hong Kong, 2022, available at: https://scholars.cityu.edu.hk/en/publications/stairway-to-helheim/.

[16] E. Elmquist, M. Ejdbo, A. Bock, D. S. Thaler, A. Ynnerman, and N. Rönnberg, "Birdsongification: Contextual and complementary sonification for biology visualization," in *Proceedings of the 29th International Conference on Auditory Display (ICAD2024)*, ser. ICAD 2024. International Community for Auditory Display, June 2024, p. 34–41. [Online]. Available: http://dx.doi.org/10.21785/icad2024.006

[17] F. Ekeberg, "Ingenmannsland (no man's land)," Sound installation, Museum Angewandte Kunst, Frankfurt, Germany, 2019.

[18] H. R. Spence and M. Ballora, "Layers of meaning: The ocean's natural acoustics and the music of its datasets," in *Proceedings of the 26th International Conference on Auditory Display (ICAD 2021)*, ser. ICAD 2021. International Community for Auditory Display, June 2021, p. 60–64. [Online]. Available: http://dx.doi.org/10.21785/icad2021.017

[19] A. P. Hill, P. Prince, E. Piña Covarrubias, C. P. Doncaster, J. L. Snaddon, and A. Rogers, "Audiomoth: Evaluation of a smart open acoustic device for monitoring biodiversity and the environment," *Methods in Ecology and Evolution*, vol. 9, no. 5, p. 1199–1211, Jan. 2018. [Online]. Available: http://dx.doi.org/10.1111/2041-210X.12955

[20] "What is audio event detection clustering?" Arbimon Help Center, https://help.arbimon.org/article/222-what-is-audio-event-detection-clustering.

[21] A. Farina, M. Ceraulo, C. Bobryk, N. Pieretti, E. Quinci, and E. Lattanzi, "Spatial and temporal variation of bird dawn chorus and successive acoustic morning activity in a mediterranean landscape," *Bioacoustics*, vol. 24, no. 3, p. 269–288, Sept. 2015. [Online]. Available: http://dx.doi.org/10.1080/09524622.2015.1070282

[22] E. Macias, J. Serrano, J. L. Vicario, and A. Morell, "Novel imputation method using average code from autoencoders in clinical data," in *2020 28th European Signal Processing Conference (EUSIPCO)*. IEEE, Jan. 2021, p. 1576–1579. [Online]. Available: http://dx.doi.org/10.23919/Eusipco47968.2020.9287343

[23] F. Grond and J. Berger, "Parameter mapping sonification," in *The Sonification Handbook*, T. Hermann, A. Hunt, and J. Neuhoff, Eds. Berlin: Logos Verlag, 2011, pp. 363–397.

[24] J. R. Barber, K. R. Crooks, and K. M. Fristrup, "The costs of chronic noise exposure for terrestrial organisms," *Trends in Ecology amp; Evolution*, vol. 25, no. 3, p. 180–189, Mar. 2010. [Online]. Available: http://dx.doi.org/10.1016/j.tree.2009.08.002

[25] A. Ballatore, D. Gordon, and A. P. Boone, "Sonifying data uncertainty with sound dimensions," *Cartography and Geographic Information Science*, vol. 46, no. 5, p. 385–400, Aug. 2018. [Online]. Available: http://dx.doi.org/10.1080/15230406.2018.1495103

[26] S. A. Vriend, D. Hägele, and D. Weiskopf, "Two empirical studies on audiovisual semiotics of uncertainty," in *Proceedings of the 20th International Audio Mostly Conference*, ser. AM '25. ACM, 2025, p. 107–123. [Online]. Available: http://dx.doi.org/10.1145/3771594.3771604