AI Enhanced Sonification of 3D Datasets for Improved Accessibility and Insight Using Python Based Data Processing and Visualization Libraries

[1]Nakul Kamatkar, [2]Chinmay Kamble, [3]Nikita Patil
[1]Independent Researcher, [2] Independent Researcher, [3] Independent Researcher
[1] Savitribai Phule Pune University, Pune, India
[1]nakulkamatkar@gmail.com, [2]chinmaykamble@gmail.com, [3]nikipatil3004@gmail.com

Abstract— **Data sonification, the translation of data into sound, offers a complementary modality to visual representation, especially for complex multidimensional datasets. This paper introduces Sonific, an AI-enhanced pipeline designed to sonify three dimensional datasets, such as point clouds, by combining machine learning–based feature extraction with spatial audio rendering. Unlike prior methods relying on raw data or bespoke mappings, Sonific automatically identifies salient structures clusters, surfaces, and anomalies and maps them to distinct auditory cues in a 3D soundscape. Our prototype demonstrates improved perceptual differentiation of shapes and spatial features compared to raw sonification approaches. The system fosters accessibility for visually impaired users, supports scientific data analysis, and opens new pathways for interactive multimodal data exploration. We discuss implementation details, prototype evaluation, and potential applications in scientific research, education, assistive technology, and creative domains.**

Keywords: Data sonification, 3D datasets, point cloud, machine learning, spatial audio, accessibility, auditory display, feature extraction,interactive-visualization.

I. Introduction

Data sonification the translation of data into sound offers a powerful complement to visual representation by leveraging the human auditory system's sensitivity to temporal patterns. This approach reveals hidden structures in complex datasets while reducing cognitive overload when visual displays become saturated. From astronomy to neuroscience, sonification serves both scientific inquiry and accessibility needs. NASA's Chandra Centre, for example, has successfully sonified astronomical data to make multiwavelength imagery accessible to blind and visually impaired audiences.

Current sonification techniques, however, primarily target one and two dimensional datasets such as time series or images. General approaches for three dimensional data like volumetric images, hyperspectral cubes, simulations, or point clouds remain limited. While efforts like the vOICe device for vision substitution and augmented brain connectome visualizations show promise, existing methods struggle to convey detailed spatial structures. Early attempts at raw point cloud sonification helped users perceive obstacles but failed to communicate object shapes effectively.

Additionally, most sonification efforts rely on custom, dataset specific mappings requiring significant domain expertise and manual tuning. As noted by experts like Dr. Arcand of NASA, one size fits all sonification schemes have limited utility, posing barriers for widespread adoption, especially for non expert users or rapid exploratory analysis of novel 3D datasets.

We present Sonific, an AI enhanced pipeline that integrates automated feature extraction with spatial audio rendering to generate meaningful auditory representations of complex 3D data. Unlike previous work, our approach combines machine learning driven analysis with modern computational tools to create a reproducible, extensible framework. The system intelligently identifies and highlights salient structural features, transcending the limitations of raw data sonification that conveys only basic properties like distance or intensity.

Our prototype demonstrates this approach through 3D point cloud sonification, showing how AI driven feature analysis is combined with spatial audio and can enhance listeners' ability to perceive structural information including distinct objects and shapes. This methodology aims to make 3D data exploration more inclusive and intuitive for both researchers and the general public.

The following sections review related work, detail the Sonific pipeline implementation, and discuss evaluation results alongside potential applications. Our goal is to establish new pathways for accessible and scalable exploration of three dimensional data through multimodal interaction.

II. Background and Related Work

The use of sound to represent scientific data has a rich history in both research and accessibility contexts. In astronomy, mapping telescope data into audio has not only opened new ways for blind and visually impaired (BVI) individuals to experience cosmic phenomena but has also provided novel insights for sighted audiences. For instance, NASA's Chandra Centre has leveraged sonification by linking image features such as brightness and position to musical parameters, allowing users to aurally identify stellar structures and gas clouds through variations in pitch and volume. Likewise, studies in neuroscience have shown that augmenting 3D visualizations such as connectome network models with sound enhances participants structural understanding and navigation within complex datasets. These efforts underscore the auditory system's remarkable ability to recognize intricate temporal and spatial patterns, often revealing information that may not be easily detected visually.

Direct sonification of three dimensional spatial data remains relatively unexplored. A significant effort in this domain is Commère et al. (2018), who investigated the sonification of 3D point clouds for sensory substitution in blind users. Adapting concepts from vision substitution (such as the 2D vOICe system), their prototype rendered each point as an independent sound source, encoding spatial coordinates into auditory cues like pitch and panning. This minimal, "raw" approach allowed users to localize objects and gauge relative sizes but was limited in its ability to communicate object shape or detailed structure, highlighting the perceptual bottleneck of low level sonification. Without some extraction of structural features, more complex attributes of 3D data can be difficult to convey through sound alone.

Some domain specific methods have attempted to bridge this gap by introducing feature extraction prior to sonification. For example, earlier work referenced by Commère et al. included contour extraction from 3D shapes, where sonifying the extracted outlines helped users recognize object boundaries more intuitively. In medical imaging, Mellor et al. (2017) introduced "TripleTone Sonification," segmenting brain PET/CT scans into interpretations tied to different audio streams. Mapping metabolic activity in distinct brain regions to independent tones allowed interactions between those tones (frequency beating and modulation) to highlight clinical features, supporting more accurate diagnoses of diseases like Alzheimer's. Such targeted mapping enhances the interpretability of sonified data, but typically these methods are crafted for specific applications or tasks rather than offering broad generalizability.

The literature thus identifies two main gaps motivating our approach: (1) the lack of a general purpose, extensible framework for sonifying arbitrary 3D datasets, and (2) the limited ability of raw parameter mapping to represent complex features like geometric shapes, clusters, or anomalies. To address these, our proposed work employs a hybrid strategy: leveraging automated machine learning and signal processing to extract salient features, and systematically mapping those to perceptible auditory cues. Importantly, we aim to balance feature augmentation with the "directness" advocated in raw data methods providing both raw and enhanced information streams, and avoiding overinterpretation or misleading simplifications through excessive automation. This dual approach strives to make 3D sonification both informative and trustworthy, supporting broader accessibility and deeper insight for varied audiences.

III. Proposed Method

Sonific is designed as a modular pipeline with three core stages: Feature Extraction, Sound Mapping, and Auditory Rendering, complemented by an optional Interactive Agent module for user feedback or customization. Figure 1 conceptually illustrates this workflow. The process begins with machine learning driven feature extraction to identify salient structures in 3D datasets. These features then guide how distinct aspects of the data such as clusters, surfaces, or anomalies are mapped to specific audio parameters. The final auditory rendering produces a soundscape that users can experience interactively or as an offline sequence for data analysis and insight.

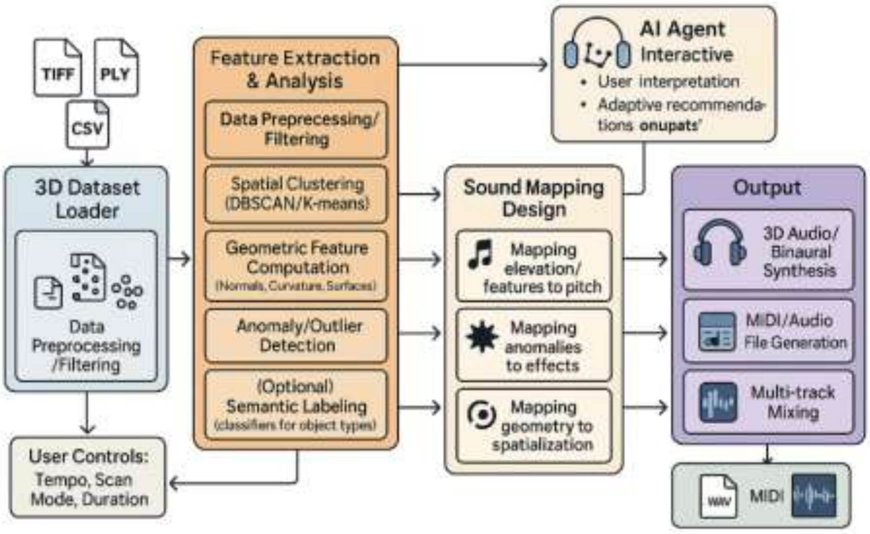

- *Figure 1: Detailed pipeline diagram of the Sonific system, showing data flow from 3D dataset input through multi-stage feature extraction, mapping, auditory rendering, and output, including the planned AI Agent module.*

Feature Extraction and Analysis

The Sonific pipeline accepts a broad range of 3D datasets, including point clouds, voxel grids, or volumetric image cubes. Automated analysis at this stage focuses on extracting higher level features that are most informative for sonification:

- **Spatial-Clustering:**
  Unsupervised algorithms cluster the data (e.g., by applying DBSCAN or K means), grouping points or voxels into distinct

objects or regions of interest. This step provides a coarse semantic segmentation, addressing the limitations of raw sonification where every data point is treated equally. By separating a point cloud of an environment into, for example, “furniture”, “walls”, and “floor”, the system can later give each object an identifiable sound palette, conveying structure beyond mere location.

- **Geometric Feature Detection:**
  The pipeline computes geometric descriptors such as surface normals, curvature, and edges for each object or region. In point clouds, this may involve detecting flat planes, corners, or other surface features; for volumetric data, edge detection and boundary analysis identify meaningful structural divides. These descriptors inform the auditory rendering about the form and complexity of each segment, allowing the sonification to communicate not only position but also shape and texture.

- **Anomaly and Salience Detection:**
  Machine learning models or statistical methods (such as outlier detection) identify anomalies or high salience areas within the data. For instance, unusual points, standout clusters, or regions with strong gradients are flagged and later represented with distinctive sound events to draw immediate user attention.

- **Semantic Labeling (Optional):**
  Where available, domain trained models can classify data regions into semantic categories (such as “tree,” “building,” or “vehicle” in LiDAR scans). This semantic layer enables intuitive auditory icons like assigning water like sounds to regions labeled “water” further simplifying the audio scene for non expert listeners.

The result is a set of feature descriptors (e.g., cluster IDs, surface labels, anomaly flags) attached to the original data, primed for the next processing stage.

Sound Mapping Design

This critical module translates data and extracted features into a coherent, informative audio format. Careful parameter mapping ensures that distinct aspects of the dataset generate perceivable, non overlapping cues:

- **Spatial-Mapping:**
  Sonific employs 3D audio synthesis, assigning each point, cluster, or feature to a virtual sound source in a rendered audio environment. Horizontal angles are mapped to stereo panning and binaural cues, vertical positioning to pitch or filtering, and depth to volume or reverberation. This spatialization gives listeners a natural sense of direction and distance, allowing immersive exploration of the 3D data landscape.

- **Pitch_and_Timbre_Mapping:**
  Scalar point attributes (intensity, reflectance, temperature, etc.) control musical parameters such as pitch or timbre. Perceptually appropriate scaling is often log based for pitch which ensures subtle differences remain audible. Timbre or instrument identity can differentiate material or semantic types (for instance, “metal” as a bell tone, “wood” as a mellow synth). These mappings are adaptable to the available semantic labels.

- **Feature_Driven_Cues:**
  Extracted clusters are assigned unique sonic identities (motifs, timbres, or musical instruments). Events such as entering a new cluster or encountering a detected surface or anomaly are marked by specific musical gestures, motifs, or sound effects. For example, cluster transitions might be announced by a chord or motif, while edges generate brief percussive tones. Anomalies could trigger attention grabbing sounds (like chimes or alarms), ensuring critical patterns are not missed.

- **Temporal_Encoding_(Sonification_Scanning):**
  The data is traversed along a defined temporary path such as a “depth based sweep” from nearest to farthest layer or a scan along a spatial axis. Each sweep creates a progressive auditory scene: near objects sound first, distant ones later, mimicking physical space traversal. Scanning speed is configurable, balancing dataset coverage with listener comfort. Users can interactively alter scan direction or focus on particular subregions for flexible exploration.

All mappings are modular and easily user tunable in the interface, with the AI assistant providing personalized interpretation and parameter recommendations; future work will enhance interactive, AI driven mapping customization. The mapping of 3D data features to auditory parameters is depicted in **[see Fig. 2]**.

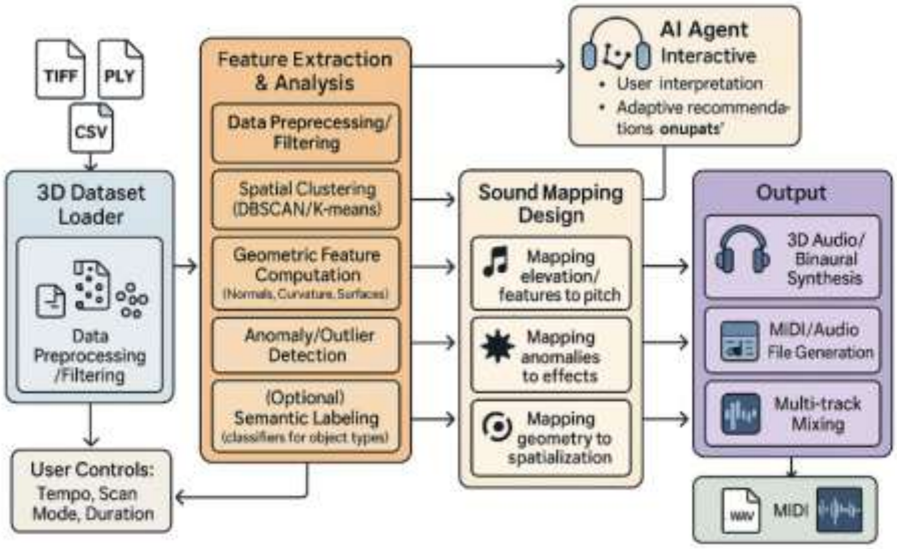

Figure 2: Schematic showing how 3D data features including clusters, elevation, anomalies, and geometric structure are mapped to auditory parameters such as instrument/timbre, pitch, spatial location, and sound effects in the Sonific pipeline.

Auditory Rendering Engine

Once parameter mappings are defined, Sonific synthesizes the corresponding audio to create a meaningful and accessible auditory representation of the input data. Our Python based prototype leverages reliable libraries such as MIDIUtil for MIDI file creation and modern audio synthesis tools for digital signal processing (NumPy and SciPy for waveform synthesis). The rendering engine processes the sequence of events derived from the mapping stage, producing either MIDI files or direct waveform outputs (such as WAV files). We employ an offline rendering approach, computing the entire sonification in advance for consistent, high quality playback and ease of distribution.
The auditory rendering pipeline incorporates several important components:

- **Spatial_Audio_Simulation:**
  Sound events are spatialized using stereo panning or full binaural synthesis. Specifically, Head Related Transfer Function (HRTF) based panning models are applied; each sound is filtered to reflect its simulated direction relative to the listener. Directions are precomputed across a sphere for computational efficiency, with each data point's orientation matched to the nearest available filter. This technique yields a convincing 3D soundstage in standard stereo headphone output, enhancing the user's spatial perception of dataset features.

- **Sound_Synthesizers:**
  Each data derived event (such as a point or feature) is mapped to a short tone, using simple waveforms (e.g., sine, square) or instrument samples for varied timbres. Tone duration may be constant or data driven (for instance, brighter points result in longer notes). For continuous structures, textures or drones represent ongoing features, while discrete "blips" highlight individual points. Multiple sonic elements are mixed into a single audio track for an immersive and information rich soundscape.

- **Multitrack_Composition:**
  To ensure clarity, the system can allocate different sound types to separate tracks (channels) during mixing. For example, background tracks carry the raw data ambience, while main feature tracks communicate clusters or salient structures with distinct motifs. Isolating and later recombining channels facilitates control over volume balance and special effects, resulting in a layered, non congested auditory experience.

The rendered output delivers an audio interpretation of the 3D dataset. As with other scientific sonifications, Users can familiarize themselves with how data features are mapped to sound. With some training, listeners can accurately identify patterns such as spatial separations or gradients purely through the sonified output.

Optional Interactive AI Agent

Sonific optionally includes an AI assistant module to support users in understanding and interacting with the sonifications. This conversational agent helps explain complex audio features, clarifies the mapping between data and sound, and offers recommendations to improve sonification clarity and effectiveness.

Users can ask natural language questions such as "What does the high pitched beep represent?" or "How can I better distinguish clusters?" The AI analyzes the underlying data features and audio characteristics, using feature extraction and audio analysis tools like librosa, to provide clear, human readable explanations. This aids users in interpreting how terrain features, anomalies, and spatial information are rendered sonically. Currently, the AI offers guidance and feedback but does not autonomously or dynamically modify the sonification parameters or audio output; parameter adjustments are made manually via the user interface. The assistant provides preset modes (e.g., scientific, and artistic) to address different user preferences and contexts.

Future work will aim to incorporate more advanced AI driven capabilities, such as automated parameter tuning, real time response to user interactions, and deeper integration within the sonification pipeline. This design maintains user control while enhancing interpretability and accessibility for diverse audiences, from experts to novices.

IV. Implementation and Prototype Demonstration

To validate our approach, we developed a prototype centered on point cloud sonification, using a synthetic dataset with simple 3D shapes specifically, clusters of points forming a cube and a sphere positioned in virtual space. This scenario was chosen to parallel the experimental setup of Commère et al., facilitating comparison with earlier work that explored auditory shape recognition in 3D point clouds. Our goal was to assess whether enhanced feature extraction and mapping could improve the perceptibility of object structure through sound.

Feature Extraction in the Prototype

With the ground truth of object identity available, we evaluated the efficacy of our automated clustering and feature analysis. Clustering via DBSCAN distinctly separated the cube and sphere, accurately labeling them as separate objects due to their spatial separation. Additionally, we computed normal vectors for each point: the cube's normals aligned with a few principal directions (corresponding to its faces), whereas the sphere's normals varied smoothly over its surface. While normal vectors were not directly used in this round of sonification, they confirmed that geometric differentiation is feasible and could be leveraged for future enhancements such as distinguishing flat versus curved surfaces.

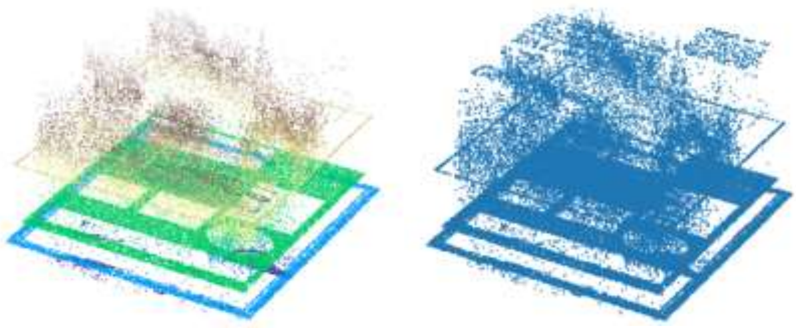

Figure 3: (a) DEM represented as a 3D point cloud. (b) Clusters detected by DBSCAN, showing spatial groupings of terrain features in color.

Sound Mapping and Audio Generation

Each cluster was mapped to a unique instrument: the cube employed a marimba like percussive tone, while the sphere was represented by a soft continuous synth pad. The dataset was sonified by scanning points in order of increasing distance from the listener's viewpoint. Points in the cube cluster generated short percussive "ding" sounds, their frequency proportional to height (Y coordinate). The sphere cluster produced a continuous hum, also pitch modulated by height. For spatial realism, the cube's sounds panned according to angle, originating from one side if the cube was positioned laterally in the scene, and similarly for the sphere. An anomaly a single outlier point placed above the others was detected as a small cluster and mapped to a distinctive, brief high pitched chime to draw immediate auditory attention.

Audio Output and Evaluation

The resulting sonification, depicted in Figure 4 as a spectrogram, lasted about 15 seconds and corresponded to a spatial sweep from the nearest to the farthest points. Listeners, including one of the authors and several colleagues (some blindfolded), reliably identified two distinct groups of sounds: first, the marimba dings for the cube (predominant on one side), then the pad hum for the sphere (from the other side). The cube's percussive notes clustered at a few discrete pitches, reflecting the limited number of height levels on its flat faces; this formed a distinctive repeated chord pattern, echoing its geometric structure. In contrast, the sphere's sounds presented a continuous sweep in pitch, producing a smooth ascending or descending tone, indicative of its continuously varying surface.

Such textural and melodic differences in the audio allowed participants to intuit structural distinctions between the objects. Even without prior information, listeners described the first object as having "layers or levels" and the second as "continuously changing," mirroring the cube's planar nature and the sphere's smooth curvature. Notably, a participant in our informal trial accurately counted the number of objects and remarked, "the second object seems to go higher than the first" an observation that corresponded with the actual geometry. This feedback suggests that feature enhanced sonification can support spatial comprehension and object counting.

Despite these positive outcomes, distinguishing the precise shape cube versus sphere from sound alone remains challenging and may require more training or additional audio cues. This limitation was anticipated and aligns with earlier findings; enabling precise shape recognition is a complex perceptual task. We did not conduct a formal user study at this stage, but anecdotal evidence is encouraging regarding the system's potential to facilitate spatial insight. Formal experiments are planned to rigorously measure task accuracy, learning curve, and usability with larger participant samples.

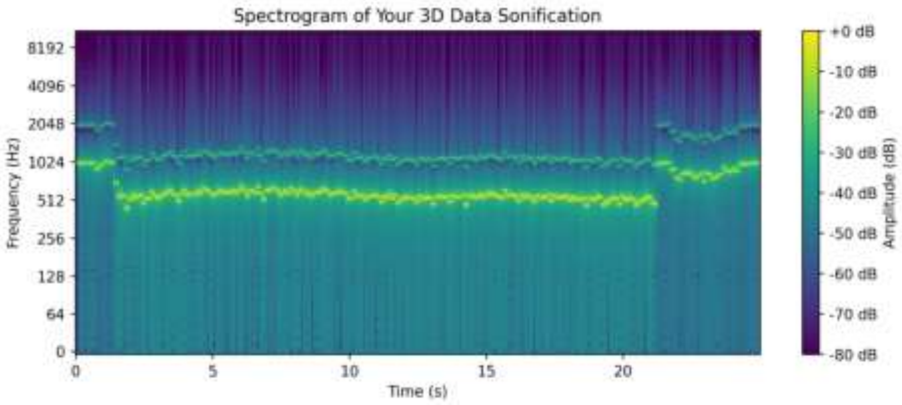

Figure 4: Spectrogram of the prototype sonification. Distinct bands correspond to discrete height levels of the cube, while continuous frequency sweeps represent the sphere's smooth elevation variation.

V. Applications and Impact

The proposed sonification approach offers promising applications across diverse fields, harnessing its ability to transform complex 3D datasets into meaningful auditory experiences.

- **Scientific_Data_Analysis:**
  Sonific enables researchers to audit and explore intricate 3D simulations that are challenging to visualize conventionally. For instance, a chemist could sonify a 3D electron density map of a protein, using distinct instrument sounds to differentiate secondary structures such as helices and sheets. This auditory rendering might reveal repeating motifs or symmetries as recurring musical themes, facilitating new insights or identifying data anomalies. Similarly, astrophysicists could sonify large scale cosmological simulations, where clusters of galaxies emerge as harmonious chords, allowing listeners to detect spatial structures through sound. By incorporating temporal scanning, even static 3D datasets can be presented dynamically, unveiling implicit changes and deepening scientific understanding.

- **Accessibility_in_STEM_Education:**
  Sonific holds strong potential to enhance inclusivity by making STEM concepts and data accessible to blind or low vision learners. In educational settings such as classrooms and museums, sonified 3D models ranging from geometric shapes and terrains to anatomical scans can be combined with tactile exhibits to provide rich multisensory learning experiences. For example, a geology exhibit might pair a tactile topographic map with sonified elevation cues, where high altitudes correspond to higher pitches and valleys to low rumbles, offering an intuitive audio tactile impression of the landscape. This multimodal approach promotes deeper understanding and engagement, extending well beyond traditional visual displays. Prior outreach efforts like NASA's Audio Universe project underscore the impact of sonification in astronomy, and Sonific's generalized framework can scale this success to broader educational content.

- **Assistive_Navigation_and_Object_Perception:**
  With further development toward real time processing, Sonific could serve as the auditory engine in assistive technologies for visually impaired users. Using live depth sensor input from smartphones or wearables, a device could generate spatialized auditory cues indicating obstacles and their relative positions. By leveraging clustering and feature aware sonification, such systems would offer richer and more informative soundscapes than simple echolocation beeps. For instance, a user might hear a "high, sharp click" to their left indicating a small object at head height, alongside a "low, soft hum" on the right signifying a flat ground plane. While analogous to existing vision to audio devices, Sonific's integration of 3D spatial and semantic cues could enhance object recognition and situational awareness. However, deploying such safety critical applications requires rigorous validation to ensure reliability and minimize false alarms.

- **Art_and_Creative_Technology:**
  Beyond utilitarian roles, Sonific opens new creative avenues by enabling artists and composers to transform scientific 3D data into novel sonic compositions. The pipeline's configurable mappings support artistic exploration, emphasizing musical scales, rhythms, and timbres that balance aesthetic appeal with data integrity. Composers might "play" architectural scans as urban symphonies, encoding building heights and spatial layouts into sound textures that engage public audiences emotionally and intellectually. This blending of scientific data and creative expression can deepen societal appreciation of complex datasets, making abstract patterns tangible through immersive sound art. Similar efforts in climate data sonification have demonstrated art's power to communicate pressing global issues, a potential Sonific can extend across domains.

These multifaceted applications illustrate Sonific's capacity to bridge scientific insight, educational inclusivity, assistive technology, and artistic innovation. By providing an extensible, feature aware auditory representation of 3D data, Sonific fosters new modalities for understanding and engaging with spatial information in ways that complement and enhance traditional visualization methods.

VI. Conclusion and Future Work

This paper has presented a novel framework for sonifying three dimensional datasets by integrating machine learning driven feature extraction with carefully designed auditory mappings. Our approach overcomes key limitations in existing methods that either rely on raw data sonification, which can fail to convey complex spatial structures, or require extensive manual tuning bespoke to each dataset. Preliminary results demonstrate the potential of Sonific to distinguish shape characteristics and highlight salient features in point cloud data, providing enhanced perceptual access compared to prior approaches.

The impact of Sonific is twofold. First, it offers scientists an additional sensory modality to explore and interpret complex 3D datasets, enabling auditory pattern recognition which may reveal insights that are difficult to visualize. Second, it contributes to accessibility by making spatial data perceptible to individuals with visual impairments and broadening public engagement through multisensory data experiences.

Looking forward, several key directions merit further investigation. Technically, we plan to refine feature extraction by employing advanced deep learning architectures such as PointNet to learn optimal representations tuned for sonification clarity and distinction. We also intend to develop a real time implementation of the pipeline, enabling live sonification of streaming depth data, potentially complemented by haptic feedback for a multisensory interface.

Crucially, systematic user evaluations are needed to rigorously quantify the effectiveness of Sonific for different user groups, including both sighted and visually impaired participants. Such studies will assess interpretability, learning curves, and utility across diverse data types and contexts.

In sum, Sonific inaugurates a new paradigm for employing artificial intelligence to enrich sonification of complex spatial data. By enabling users to "hear" shape, structure, and anomalies in three dimensional datasets, our work advances both scientific analysis and accessibility. We envision ongoing enhancements and evaluations will pave the way toward mature, real world tools that integrate seamlessly into researchers' workflows and education environments, fostering inclusive discovery and understanding.

VII. References

[1] [Arcand et al. 2021] Arcand, K., Watzke, M., De Pree, C., Edmonds, P., Arcand, K., Shishkovsky, L., and Lestition, K.: "Why Make Sonifications of Astronomical Data?", Chandra X ray Center Blog (2021).

[2] [Audio Universe 2022] Audio Universe Project Team: "Audio Universe: Tour of the Solar System – an accessible sonification based astronomy show", 2022. Accessible at: https://www.audiouniverse.org

[3] [Commère et al. 2018] Commère, L., Truillet, P., Oriola, B., and Jouffrais, C.: "Sonification of 3D Point Clouds for Substitution of Vision by Audition for Blind Users", Proc. 24th Int. Conf. on Auditory Display (ICAD), 2018.

[4] [Gao et al. 2021] Gao, X., Choi, B., Plumbley, M.D., & Wang, W.: "Quantitative Evaluation of Sonification Strategies for Elevation Perception", Proceedings of the 27th Int. Conf. on Auditory Display (ICAD), 2021.

[5] [Lluís et al. 2022] Lluís, F., Benacchio, F., Munteanu, C., and Giró i Nieto, X.: "Points2Sound: from mono to binaural audio using 3D point cloud scenes", EURASIP Journal on Audio, Speech, and Music Processing, 2022:37, https://doi.org/10.1186/s13636 022 00247 x

[6] [Markoff 2023] Markoff, S.: "Understanding data through other senses: Data Sonification", Blog Post, November 13, 2023.

[7] [Mellor et al. 2017] Mellor, S., Crutch, S.J., and Bogaardt, H.: "A Novel Sonification Approach to Support the Diagnosis of Alzheimer's Dementia", Frontiers in Neurology, 8, Article 647 (2017). https://doi.org/10.3389/fneur.2017.00647

[8] [Papachristodoulou et al. 2014] Papachristodoulou, P., Papaharilaou, Y., and Badar, Z.: "Sonification of Large Datasets in a 3D Immersive Environment: A Neuroscience Case Study", Proc. ACHI 2014, pp. 183–189.

[9] [Schwartz et al. 2020] Schwartz, A., Liu, S., and John, K.: "Mobile Application for Real Time Sonification of 3D Scenes for Visually Impaired Users", IEEE Access, 8, pp. 178354–178367, 2020.

[10] [Sun et al. 2023] Sun, C., Ju, R., Liu, Z., Wang, J., & Wang, Y.: "Research on point cloud hole filling and 3D reconstruction for robust data sonification", Scientific Reports, 13, 18166, 2023. https://doi.org/10.1038/s41598 023 45648 5

[11] [Wu et al. 2024] Wu, L., Jin, C., Mushtary Uttsha, M., & Vidal Calleja, T.: "A Scene Representation for Online Spatial Sonification", arXiv preprint arXiv:2412.05486, 2024. https://arxiv.org/abs/2412.05486

[12] [Zhang 2023] Zhang, H.: "Deep Learning based 3D Point Cloud Classification", arXiv preprint arXiv:2311.02608, 2023. https://arxiv.org/abs/2311.02608

[13] [ICAD 2018] ICAD: Proceedings of the 24th International Conference on Auditory Display. 2018.