CENOTAPH AND ELEGY OF CONTAGION: COLLECTIVE IMMERSION OF NATIONAL-LEVEL COVID-19 FATALITY DATA IN THE UNITED STATES

M. Jerry Huang, Stephanie Loveless, and Jonas Braasch

Rensselaer Polytechnic Institute
110 8th Street, Troy, New York, Unites States
[huangm5, lovels3, braasj]@rpi.edu

ABSTRACT

Effective representation of epidemiological data is essential for informing the general public about the extent of disease contagion, yet hard to achieve. On the one hand, the expressiveness of statistical visualization is often limited. On the other hand, explicit data sonification is lost in translation, failing to capture the meaning of the data at hand. The ongoing COVID-19 pandemic, which has caused excessive fatalities in the world and the United States, has presented us with a timely opportunity to address both challenges at once. In this work, we use the room-oriented immersive system (ROIS) to sonify and visualize the United States' national-level fatality data for the first 1000 days of the pandemic. Reimagining ROIS as a cenotaph that can be experienced collectively from within, a visual timeline in upward motion is spatially distributed to represent individual fatality cases in scattered point densities, resulting in the weighty perception of downward self-motion. The visualization data is spatially subdivided and compressed into multi-channel spike chains with variant embeddings. The sonification, or elegy, uses a self-assembled dataset of heavy breathing sounds and is synthesized as virtual sound sources through a finite convolution-sum technique. The integrated result demonstrates the method's broader ability to effectively convey the relationship between human behaviors and pathogenic evolution through collective immersion while delivering a responsible representation of infectious disease statistics.

1. INTRODUCTION

A cenotaph is a monument of lost lives displaced. Like other monuments of the deceased, the cenotaph installs a sense of permanence in public memory. Yet, the absence of bodies in a cenotaph suggests a different level of grief that goes beyond remembering an individual. The hollowness represents the weight of emotions, as no individual can fully capture the universal feeling of mournfulness behind its story.

The ongoing COVID-19 pandemic has saturated human society with contagions and deaths on a global scale and has had a profound impact on people's psyche. In the United States, more than 1.1 million lives to date have been lost to the SARS-CoV-2 virus since the pandemic started in January 2020 [1]. This deceased collective also tells a story about the Anthropocene: the fragility of being under the unintended consequences of human engagement with non-humans, the political and scientific visions it engenders, and the human footprint, means of collective movement, and migration on the planet. The ability of the virus to take away our senses and even our ability to breathe has, at one point, decelerated the society at large. While humans gradually learn to manage its impact, the pandemic has inevitably left us with a collective trauma filled with numbness, frustrations, and sorrow.

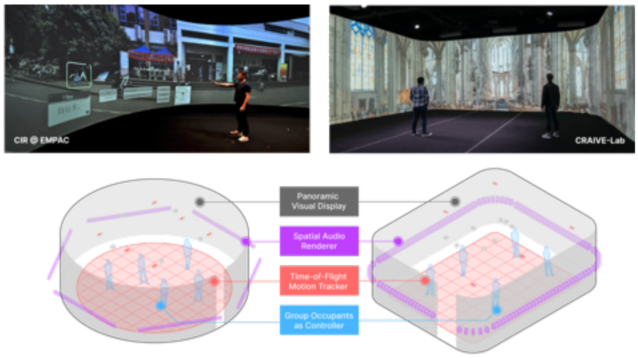

Figure 1: Photograph of the cenotaph being rendered on the panoramic display in one of the ROIS facilities. Here, the visualization moves in upward direction at a constant speed, showing changes of the fatality trends.

The underlying emotional signal of this pandemic motivates a re-imagination of the collective voice of US COVID-19 fatalities from a statistical perspective. Indeed, *while numbers do not lie, they can hardly be felt.* Even though computational modeling and data visualization of disease trends in the pandemic can ground our awareness of the pandemic's evolution, it usually abstracts away the lived experiences perturbed beyond the reduced dimensions of data. Representing this aspect, especially to cultivate collective empathy, requires a new configurative perspective that makes these data more accessible and embodied in space.

In this work, we address this representational issue using a room-oriented immersive system (ROIS). A ROIS is a multi-user virtual reality system with a room-scale panoramic display surrounded by a multi-channel loudspeaker system for spatial audio rendering. Using synchronous data visualization and data sonification, we re-conceptualize the spatial geometry of ROIS as a cenotaph in which multiple people can share a common experience.

In addition to visually experiencing the case and fatality statistics as a linear timeline, the data is augmented and animated aurally with variant distributions as an elegy. The relationship between the pandemic's mass fatality and the underlying pathogenic evolution resonates constantly within the cenotaph, thereby rendering an enveloping audiovisual structure that amplifies the intensity of the collective affect. Through this setup, we provide people with a space to grieve and experience the human cost of the pandemic in a raw and embodied manner.

2. BACKGROUND

Due to the use of ROIS, this work contains an equal and interwoven contribution from both visual and audio realms. In this section, we provide a background review of data visualization and data sonification for epidemiological data, with a specific focus on the existing works related to the COVID-19 pandemic, particularly in the United States. Our goal for this review is to argue that, while visualization and sonification techniques can be mutually informative, they complement each other expressively through motion-driven audiovisual spatialization in ROIS.

2.1. Epidemiological data visualization

From daily case and fatality reportings, [1, 2, 3] to statistical modeling of disease trends [4], epidemiological data is mainly visualized as spatiotemporal distributions. Representation techniques of case and fatality data often include timelines and maps, which show disease trends over temporal and geographic differences. These visualizations are essential ways to provide direct and intuitive information for the general public to cultivate awareness about the progression of the COVID-19 pandemic. They have made abundant appearances on government websites and news media. In addition, they have been used as a metaphorical device for government and health officials to discuss mitigation strategies[1].

While statistical visualization is reliable and efficient, the accessibility of this data visualization is not sufficiently questioned. However, beyond the most obvious observation that these visualizations do not always consider the visually impaired [6], accessibility is also not always guaranteed, even for healthy-sighted individuals. In particular, due to the low-dimensional nature of the data at hand, where cases and fatalities are of the most interest, their expressiveness is restricted – there are only limited ways to reinterpret these data to accommodate for their augmentation both in terms of vision and in other sensory domains. This limitation presents an accessibility challenge because reinterpreting these data almost certainly means relying on a semantic representation for translation, especially in terms of audio (e.g., reading out the disease trends using a screen reader as if they are text). As one of the goals of greater accessibility is to directly address the affective magnitude of the pandemic's impact on the collective population, such direct representation becomes inadequate.

On the other hand, some fatality visualization strategies try to break away from direct representation. For example, Gamio and Leatherby [7, 8] presented a timeline visualization of COVID-19 fatalities that adopts representation of individual cases as grayscale noise gradients, with simple pixelated representation of individual fatality cases aggregated as densities in a confined sample space. By redistributing fatality statistics as simple individual case pixels, this approach takes advantage of visual motion in space as part of the viewing experience: changes in the pandemic's severity are represented by the variant degrees of noise density.

However, one cannot directly interpret these visualization strategies as auditory information. If we consider the visuals as "spectrogram" of some kind, it would result in the sound of white noise with varying intensity. But in this pandemic, *the lives lost are signals instead of noise*. The visual noise texture lacks features of a collective voice that represents the excessive fatalities that otherwise constitute a supposedly lived experience. As a result, it compromises the bodily visceral impression often associated with the deceased. It is crucial to represent their voices as a collective since the bodily impression gives meaning to the context of sonification here, in which the mortality is especially relevant to the desire for the air to breathe, to speak, and to voice. Considering this impression, the sonic environment of fatality data shall depart from noise so that it can deliver a more authentic and more solemn affective message. This departure has the potential to augment the limited expressiveness of visualization.

2.2. Epidemiological data sonification

Compared to the vast amount of examples in visualization, sonification of COVID-19 disease statistics still finds limited instances. The focus of sonification works covers different facets of the pandemic, from rendering viral genomes [9, 10] to providing auxiliary health supports [11]. For our purpose, where the disease trends are of concern, existing examples mostly take a direct approach. Some combine common forms of musical languages and features for an explicit data representation. For example, Nguyen et al. [12] uses the case statistics of Denmark since the pandemic onset for a structured rhythmic sonification, while Rebelo et al. [13] focused on a pitch-based representation of case and fatality accumulation. Thompson et al. [14] uses a synthetic approach, representing the case and fatality statistics through timbral and intervallic manipulations. Lemmon et al. [15] uses a geospatial mapping of case statistics combined with demographic information, and the sonified results are combinations of simple waveforms of different frequencies. Others, such as *Vox Aeterna* [16], augment fatality statistics by representing their corresponding age group as vocal features. The use of vocal gestures in this approach introduces affective and embodied meanings to the data [17].

The sonification approaches provided above present several issues. For the direct sonification, similar to the limited expressiveness of visualization, these approaches did not adequately address data augmentation. While the explicit mapping of the disease trend parameters to acoustic parameters can preserve the semantic structure of the data, the semiotic representation remains undefined. In the context of product-sound design, the *product identity* of sound does not go beyond the numeric values of the dataset [18]. The consequence is that the sonification resulting from these approaches can be under-specified and be generalized into other unrelated domains that capitalize on time-series data. For example, the same approach that Rebelo et al. used can also translate to representing weather patterns or financial statistics where similar trends exist. In other words, without being verbally introduced to the context behind the direct sonification, the meaning of sound is lost.

Meanwhile, using a rhythmically and melodically driven ap-

[1]The notion of *flattening the curve* refers to the statistical visualization of disease trends and is often related to the communication of pandemic-related public policies [5].

proach to augment the data, with an intention to treat the data as materials to "compose" in a conventional sense, can propose a representational imbalance where the data is overly aestheticized. While this issue is more apparent in genomic sonification, it also applies to timeline sonification. For instance, the HER research group adopted a global-scale dataset of the pandemic onset and used it for an implicit tonal mapping by compressing the geographical locations embedded in the case statistics [19]. While the compression introduces variability to the sonification, it compromises the structure of the data: isolating the listening experience from the data display results in a disassociation between the sound and the timeline it represents.

While the indirect approaches give space to the meaning-making process of the data through augmentation, they do not take any advantage of spatialization. Spatialization has been a crucial component in increasing the interactivity of sonification [20]. In the sonification of disease trends, spatialization can accentuate geographical differences commonly observed in the visualized disease trends. In addition, one can use it to encode other information, such as the evolution of variants and demographic differences. Since visualization takes advantage of spatial environments, one can derive spatialization strategies for sonification from visualization. For instance, one can distribute the data into multiple synchronous streams of virtual sound sources in a multichannel setting and render each stream in correspondence to the visualization generated from the same dataset. Such a distributed setup promotes greater accessibility by depersonalizing the experience.

The above arguments do not intend to discount the aesthetic value of the existing sonification works. The precedents have opened up a needed conversation about how representation in sonification affects our perception of the associated data. Our argument, however, is that a balance between information and aesthetics in sonification can be negotiated through its augmentation in spatial dimensions. Because our bodies are spatial objects, and our auditory perception engages spatial features constantly, such augmentation can result in active listening experiences that are easier to embody, making the sonified information more accessible.

2.3. A synthesized approach

For this work, we use cenotaph and elegy as metaphorical devices to augment the limited expressiveness in COVID-19 cases and fatality visualization with spatially distributed sonification. Using ROIS as a spatial carrier, we integrated three general strategies: 1) incorporating spatially distributed listening modalities; 2) applying textures to the sonification using contextual sounds; and 3) synchronizing sonification with an equivalently-textured visual timeline. Through the use of these three strategies, we provide an experience of the fatality timeline where multiple people can be present and embody the same spatial experience.

3. METHODOLOGY

3.1. System Architecture

The overall architecture of the system can be found in Figure 2. In general, this system uses fatality data and variant statistics as inputs, from which both visualization and sonification patterns are generated. For spatiotemporal synchronicity, both visualization and sonification depend on the same set of algorithms for data configuration and processing. The algorithm generates individual fatality cases as a spatial distribution with variant embeddings, drawing from Gamio and Leatherby's visualization approach [7], and does so for the first 1000 days of the pandemic in the US (from January 20, 2020, to October 16, 2022). These embeddings are then transformed into spatially distributed bitmaps and multichannel spike chains that synthesize virtual sound sources from a handcrafted breathing sound dataset. The visualization and sonification algorithm is implemented in Python. In the following, we discuss the computational aspect of each component in greater detail.

3.2. Data Collection and Processing

We have collected three structured datasets for the use of both visualization and sonification:

1. The national-level daily case and fatality dataset from the New York Times (7-day rolling average)[2];
2. The national-level weekly variant proportion dataset from the Center for Disease Control and Prevention (CDC) [21, 22, 23, 24][3]; and
3. A set of self-assembled audio dataset consisting of various sound samples of heavy breathing.

As mentioned in section 2, regional-level data can introduce variations of patterns in both visualization and sonification. Here, however, we intend to augment the national-level dataset. The national-level dataset not only symbolizes a shared experience independent of geographical differences and human behaviors in the virus' contagious spread but also represents more accurately the weight of grief brought by the pandemic as a unified and collective affect shared by the broader society.

The complexity of visualization and sonification is introduced by the combined use of the case and fatality dataset and the variant statistics. However, variant statistics play a secondary role. Specifically, we use the case and fatality dataset for the spatiotemporal structuring of the entire experience, while the variant statistics introduce individual differences within the structure. The use of variant statistics also reflects the dynamics between the human consequence of the pandemic and its underlying pathogenic evolution without losing the integrity of the affective message itself. For sonification, variant statistics also introduce variational behaviors of the use of sound materials in the dataset.

A list of important variables used for data processing is presented in Table 1. We use the visual display dimensions of ROIS, (X, Y), as a starting point for data processing. This dimension determines the number of consecutive days, N_d, to be visualized within ROIS at a time based upon a visual reference size of individual fatality case s, by which a sample space,

[2]The case and fatality dataset can be found in the following GitHub repository: https://github.com/nytimes/covid-19-data/blob/master/rolling-averages/us.csv. From a practical standpoint, this dataset is chosen over the official public-access dataset from the Center for Disease Control and Prevention (CDC) due to its simplicity. The latter contains rich statistics that are aggregated over multiple sources. Disentangling these sources within the dataset is not realistic without deep domain expertise. Still, to the best of our ability, we used the CDC dataset for comparison during the production of initial prototypes, which shows that the New York Times dataset provides an equally effective representation of fatality trends.

[3]The variant dataset can be found in the following repository: https://data.cdc.gov/Laboratory-Surveillance/SARS-CoV-2-Variant-Proportions/jr58-6ysp.

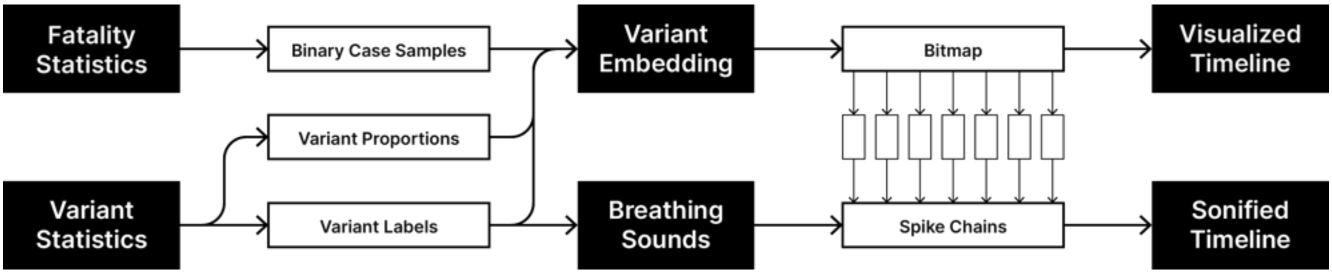

Figure 2: Overall system architecture for this work. The system considers daily fatality reporting as the primary input and variant statistics as the secondary input. The fatality statistics generate distributed and permutation-invariant binary samples that represent the individual fatality cases as spatial density. The variant statistics include variant proportions and variant labels. The variant proportions further partition the binary case samples while the labels encode the case samples and the breathing sound dataset. The variant-embedded case samples generate a grayscale bitmap for visualization while simultaneously generating a series of spike chains through subdivision and compression. The breathing sounds convolve with the encodings on a spike-by-spike basis to generate the spatially distributed sonification.

Table 1: List of basic variables for the system.

Symbols	Category	Descriptions
N_f	Data	Number of fatality cases
p_v	Data	Proportion of variant
N_v	Data	Proportion of variant embedded within the sampled fatality cases
X, Y	Visualization	True width and height of resolution for panoramic display
X_s, Y_s	Visualization	Derived size of sample space for daily fatality cases
N_s	Sonification	Number of virtual sound sources
$N_{f,t}$	Sonification	Number of fatality cases per sound source

$$(X_s, Y_s) = \left(\frac{X}{s}, \frac{Y}{N_d s}\right), \tag{1}$$

is also determined for generating the spatial distribution of fatality cases. This spatial distribution is used both for visualization and for sound synthesis. The geometry of this sample space is constrained in two ways: 1) X_s and Y_s must have integer division to ensure the appropriate data correspondence for the generated virtual sound sources, and 2) the total size of this sample space must be greater than the maximum number of fatalities for the entire duration of the fatality case reporting in the dataset. Within this sample space, daily fatality reporting populates a collection of binary data through random permutation.

The variant dataset [21, 22, 23, 24] consists of weekly and bi-weekly data of individual variant proportions, reported both nationwide and per Health and Human Services (HHS) regions. Consistent with the case and fatality statistics, we choose national-level weekly reporting. In addition to the distribution of proportions in this dataset, we also use individual variant's PANGO Lineages [25] (see Figure 4 *left*) as labels. We classify the variants into prominent and non-prominent types based upon a variable threshold (see Figure 3). This threshold allows us to select dominant variants of interest based upon its peak proportion across the entire duration of the dataset (see also Table 2)[4]. This selection process determines how much sound materials are used within the breathing sound dataset and further influences the degree of variability for the sonification.

[4]For the BQ.1 and BQ.1.1 subvariants, the peak proportion does not appear until after 1000 days since the pandemic's onset. However, they are included as key variants because their dominance started to emerge towards the end of the designed timeline.

Table 2: Key variants and their peak proportions.

WHO Label	Pango Lineage	Peak	Peak Week
Alpha	B.1.1.7	0.6886	15
Delta	B.1.617.2	0.9890	44
Omicron	B.1.1.529	0.4161	48
Omicron	BA.1.1	0.7281	56
Omicron	BA.2	0.7330	63
Omicron	BA.2.12.1	0.6227	69
Omicron	BA.5	0.8621	81
Omicron	BQ.1	0.2521	97
Omicron	BQ.1.1	0.3728	100

For each sample space of fatality cases, the variant proportions, p_v, are embedded using a multinomial sampling process:

$$N_v \sim \text{Multinomial}(N_f, p_v). \tag{2}$$

This sampling process encodes variant labels to the individual fatality cases, to reflect the variants' proportions.

3.3. Visualization

The variant embedding process allows us to augment the visualization concept from Gamio and Leatherby's infographics [7] with greater complexity. From the embedded binary case samples, a timeline is generated as a continuous bitmap to cover the vertical surface of the entire projection area of ROIS. A photograph of this visualization in a ROIS facility is shown in Figure 6. Here, the variant labels are represented as pixelated icons, as visualized in Figure 4 (*left*). The size of these icons corresponds to the visual reference size s, which determines the pixel density of the entire

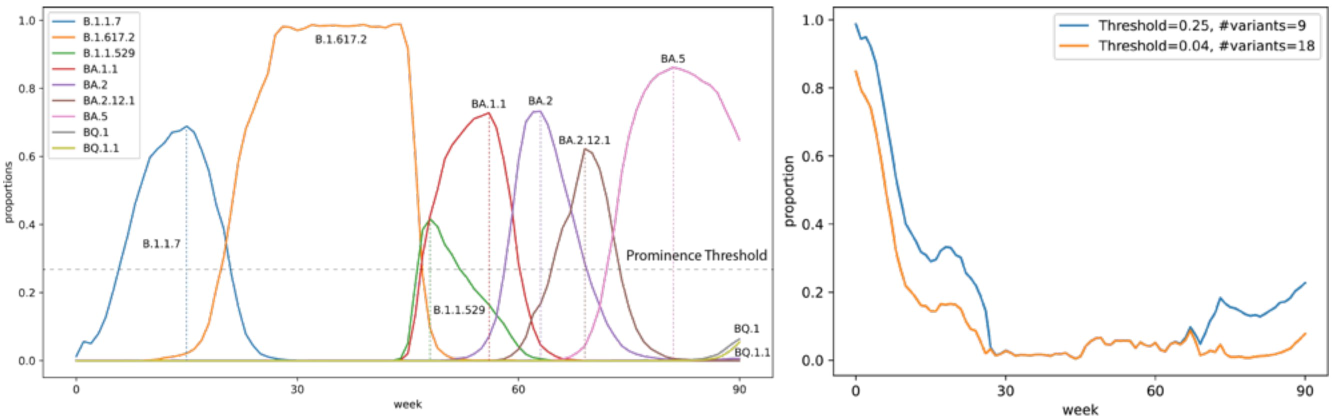

Figure 3: Timelines showing changes of variant proportions. *Left*: weekly proportions of dominant and emergent variants and sub-variants over the timeline (90 weeks since first observation). The prominence thresholds are set to select these key variants based on their peak proportions and the week of their peaking. The BQ.1 and BQ.1.1 sub-variants are included to indicate their emergence as key variants post BA.5 dominance. *Right*: total proportions of non-prominent variants over the timeline of interest as the left figure, based upon two prominence thresholds (0.25 and 0.04). The dark (blue) line indicates a higher threshold, resulting in 9 key variants; the light (orange) line shows a lower threshold that results in 18 key variants. The variability of thresholds suggests different degrees of variability for the encoded sound materials.

visualization relative to the room. This reference size can vary depending on the physical size of ROIS's infrastructure (see Figure 5).

Once triggered, the visualization moves steadily upward at one day of case samples per second. This upward movement is a crucial component behind the concept of the "cenotaph": people experience the hollow monument from within, and the upward movement of the visualization creates the illusion that the cenotaph is moving downward, similar to the experience of a sinking elevator. This movement serves as a metaphor for weight: the accumulation of fatality cases introduces a sense of heaviness.

3.4. Sonification

3.4.1. Sound Synthesis

The variant-embedded case samples used for visualization is also used to synthesize sound. To this end, an array of virtual sound sources is distributed across the entire display area, and located at ear height. The number of virtual sound sources, N_s is simply determined by:

$$N_s = \frac{X_s}{Y_s}. \tag{3}$$

With this, the sample space is vertically sliced into N_s pieces, each representing a timeline for a single virtual sound source. These timelines are compressed from top to bottom into a "spike chain" for sonification. This compression process does not affect the sequence of variant embeddings previously attributed to the visualization. When cascaded throughout the timeline, these spike chains form a running signal akin to the irregular firing (discharge) pattern of neurons.

For one second of the visualization, a total of Y_s^2 positions in the sample space is traversed. This means that, for a synchronous representation of the variant, one sample of audio can be further partitioned into Y_s^2 subsamples, with the subsample size τ for an individual spike is simply:

$$\tau = \frac{f_s}{Y_s^2}. \tag{4}$$

Within this subsample, the location of the spike is selected at random, to avoid undesired rhythmic pattern between the sound sources.

The spike chain is the backbone of synthesis for the sonification. Using the sound materials from a self-assembled dataset of heavy breathing sounds, with each sound material corresponds to one PANGO Lineage of the key variants, the synthesis takes advantage of the multiplicative identity in linear discrete-time convolution, namely

$$m[t] * \delta[t] = m[t], \tag{5}$$

where m is any sound materials, and δ is the Delta function that represents individual spike. Using this identity, the various breathing sounds overlay each other with their corresponding spikes as the time of onset, and this overlay is cascaded for the entire duration of display. The resulting sonification retains a noise-like texture as a result of the spike chain, while rendering them reverberant because of the trailing convolution-sum technique.

3.4.2. Spatialization

The spatialization is managed using the Ircam Spat library in Max/MSP [26]. The positions of individual virtual sound sources is determined by the corresponding vertical slices of the visualization. For synchronicity, the playback of sonification uses the same trigger as the one for the visualization.

4. DISCUSSION

The sonification and visualization presented in this project have both conceptual and perceptual significance. To discuss the significance systematically, we analyze this work in two parts. First, we draw from an existing analytical framework for distributed user

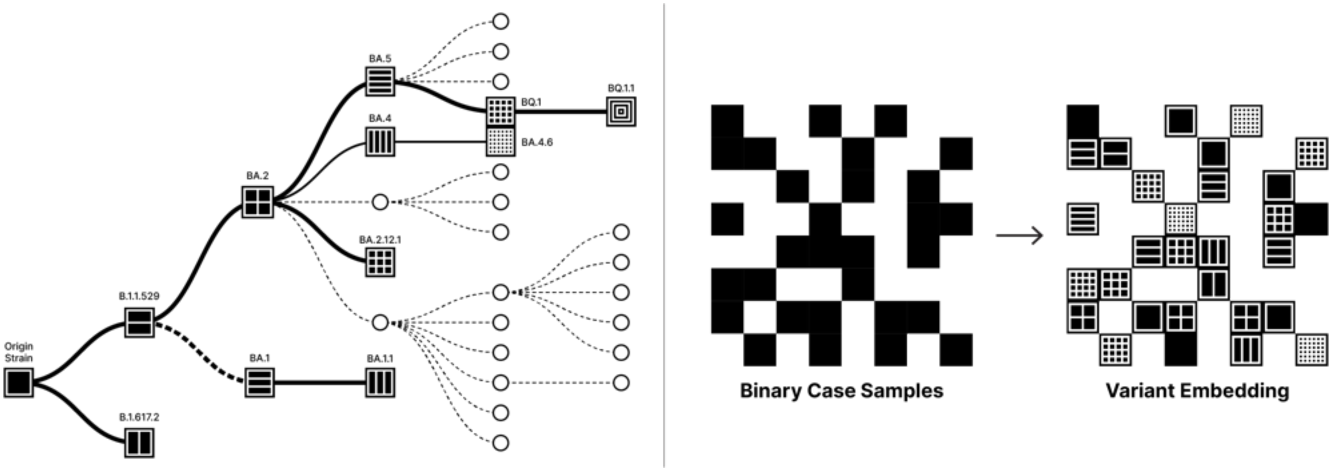

Figure 4: Diagrams showing the process of variant embedding for the system. *Left*: a snapshot of PANGO Lineage [25] showing the virus' mutation during the first 1000 days of the pandemic. Prominent variant types during this period are assigned pixelated symbols for visualization and a label for sonification. Other non-prominent variants (shown in small circles) are classified as "other variants." *Right*: The sampled bitmap for the binary representation of individual cases is used as the basis for embedding variant symbols and labels. The embedding is generated through random spatial sampling of individual cases within the bitmap using the proportion of individual prominent variants.

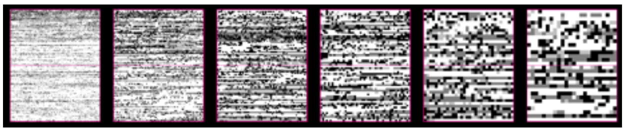

Figure 5: Study of various pixel densities for visualizing the fatality timeline.

interaction design in ROIS [27] and dissect the sonification system in terms of the spatial relationship it engages, its hardware-related design modalities, and its group interaction properties. Second, we further explore this work from a human-centered perspective and analyze it from the perspective of human perception and mental representation of space.

4.1. Design

The group interaction design use consists of three observable qualitative factors: 1) topological properties, which specify the spatial relationship between ROIS and the augmented virtual environment, both from the perspective of the system itself and from a global perspective; 2) hardware-specific design modalities, which describes how each system components of ROIS are used for the experience of viewing, listening, and navigation; and 3) group interaction properties, which defines the general awareness of ROIS users to the system infrastructure due to the virtual environment itself. For the cenotaph and elegy, these factors can be outlined below:

1. **Topological properties**: the upward motion of the visual timeline introduces a reciprocal downward self-motion of the physical space. The impact of this self-motion suggests that the interactivity goes *inward*, namely, the people experiencing the cenotaph are simultaneously experiencing the spatial enclosure of the system it is embedded in. This is further supported by the nature of the corresponding sonification: the information-driven sound field is generated within the same spatial enclosure. From an allocentric perspective, the interaction design of the cenotaph uses both the spatial *volume* and the *surface* of the spatial topology, capitalizing on the interaction between these two components;

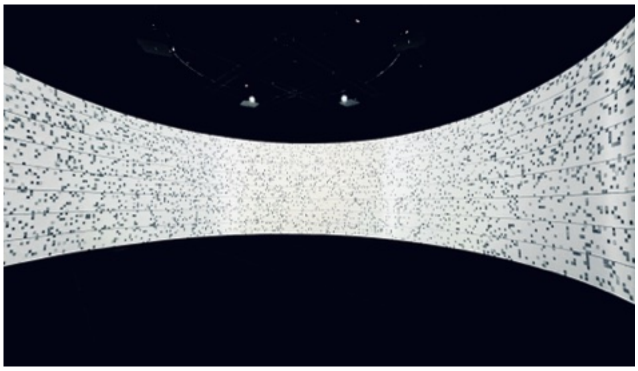

Figure 6: Photograph of the cenotaph being rendered on the panoramic display in one of the ROIS facilities. Here, the visualization moves in upward direction at a constant speed, showing changes of the fatality trends.

2. **Hardware-specific design modalities**: The evenly distributed spatial density changes of both visualization and sonification govern how the cenotaph uses each hardware component of ROIS. On the one hand, from the visual perspective, the entire panoramic display is engaged, but the representation does not extend beyond ROIS's spatial footprint. This configuration suggests a *2.5D* design modality. On the other hand, the spatialized sound field contains synchronous elements with the visual field but does not encode

the acoustic signature of an extended environment. Its even spatial distribution renders a uniformed *ambient* sound field within the physical space with no location-specific interaction parameters of interest. In addition, the motion vector of the entire experience is highly *axial*. and

3. **Group interaction properties**: The experience of the cenotaph is passive but induces an active perception of spatial motion. This spatial awareness makes *explicit* the passivity of user interaction. In addition, the even spatial distribution creates no focal point for the experience. Regardless of where the user is located in the cenotaph, the experience remains the same. This *decentralized* interaction design resonates with the conceptual language that the affect conveyed through the cenotaph and elegy is universal.

4.2. Perception

In ROIS, the shared interactive experience between its group users is facilitated by a shared mental model of space [28]constructed from the direct perception of spatial layout induced by visual and auditory motion. From the visual perspective, people can experience the physical space of ROIS as a representation of the cenotaph because of a strong ground dominance effect [29] given by the distributed and near homogenous audiovisual representation. When the ground is the only stationary object of reference, one perceives the space itself. This notion is further supported by the fact that gradual textural change of the distributed virtual sound sources that sonify the fatality statistics introduces interlocking cross-modal attentional interaction shared by the people immersed in ROIS [30, 31].

From a conceptual standpoint, the globally distributed sonification introduces a sensibility of compression that reinforces the grazing visual movement. Since this time-as-motion approach is more effective when the focus is on change over time[5], the motion allows the experiencers to embody the surroundings. The use of gasping sound as base material introduces an affective message that the contagion is strongly related to the action of breathing while the variability of these base materials preserves the information needed to demonstrate the evolving relationship between our collective mentality and the pathogen throughout the pandemic.

5. CONCLUSION

In this work, we have recontextualized the notion of cenotaph and elegy as objects of memorability and designed and implemented a system for visualizing and sonifying US national COVID-19 fatality data using room-oriented immersive systems (ROIS). Through this system, we have demonstrated the efficacy of large-scale audiovisual data display as one of ROIS's core capabilities, which has also opened up questions about perceptually-driven data representation.

We believe that the methodology involved in this work can generalize to broader scenarios where epidemiological statistics can become experiential and its accessibility more pronounced to the public. In the future, we intend to explore this generality by solidifying the system for other disease trends.

[5]As opposed to the relationship between individual data points, see [17].

6. RESOURCE AVAILABILITY

The Python code and organized data for generating the sonification and visualization for this work are currently available as a public repository on GitHub under the MIT License. A sampled sonification excerpt is also available within the repository[6]. An compressed binaural prototype of the sonification excerpt is also available as a supplemental material.

The cenotaph as an installation is also available for display during this conference, using the Cognitive Immersive Room (CIR) at the Experimental Media and Performing Arts Center at Rensselaer Polytechnic Institute. The total duration of this installation is about 17 minutes.

7. REFERENCES

[1] Center for Disease Control and Prevention, "COVID Data Tracker," U.S. Department of Health and Human Services, CDC, 2024, https://covid.cdc.gov/covid-data-tracker.

[2] E. Dong, H. Du, and L. Gardner, "An interactive web-based dashboard to track COVID-19 in real time," *The Lancet Infectious Diseases*, vol. 20, no. 5, pp. 533–534, May 2020, publisher: Elsevier. [Online]. Available: https://doi.org/10.1016/S1473-3099(20)30120-1

[3] B. D. Wissel, P. J. Van Camp, M. Kouril, C. Weis, T. A. Glauser, P. S. White, I. S. Kohane, and J. W. Dexheimer, "An interactive online dashboard for tracking COVID-19 in U.S. counties, cities, and states in real time," *Journal of the American Medical Informatics Association*, vol. 27, no. 7, pp. 1121–1125, June 2020, _eprint: https://academic.oup.com/jamia/article-pdf/27/7/1121/34153390/ocaa071.pdf. [Online]. Available: https://doi.org/10.1093/jamia/ocaa071

[4] C. Tsay, F. Lejarza, M. A. Stadtherr, and M. Baldea, "Modeling, state estimation, and optimal control for the US COVID-19 outbreak," *Scientific Reports*, vol. 10, no. 1, p. 10711, July 2020. [Online]. Available: https://doi.org/10.1038/s41598-020-67459-8

[5] J. Santos, "Reflections on the impact of "flatten the curve" on interdependent workforce sectors," *Environment Systems and Decisions*, vol. 40, no. 2, pp. 185–188, June 2020. [Online]. Available: https://doi.org/10.1007/s10669-020-09774-z

[6] D. Fan, A. Fay Siu, H. Rao, G. S.-H. Kim, X. Vazquez, L. Greco, S. O'Modhrain, and S. Follmer, "The accessibility of data visualizations on the web for screen reader users: Practices and experiences during covid-19," *ACM Trans. Access. Comput.*, vol. 16, no. 1, mar 2023. [Online]. Available: https://doi.org/10.1145/3557899

[7] L. Gamio and L. Leatherby, "How 450,000 Coronavirus Deaths Added Up," The New York Times, 2021, https://www.nytimes.com/interactive/2021/01/27/us/us-coronavirus-deaths-rate.html.

[8] N. Coleman, "On the Front Page, a Wall of Grief," The New York Times, 2021, https://www.nytimes.com/2021/02/21/insider/covid-500k-front-page.html.

[6]https://github.com/jerrymhuang/Cenotaph-Elegy..

[9] M. D. Temple, "Real-time audio and visual display of the Coronavirus genome," *BMC Bioinformatics*, vol. 21, no. 1, p. 431, Oct. 2020. [Online]. Available: https://doi.org/10.1186/s12859-020-03760-7

[10] E. J. Martin, "Sars coronavirus replicase sonification," in *Proceedings of the International Conference on Auditory Display*, 2021.

[11] B. Van Kerrebroeck and P.-J. Maes, "A breathing sonification system to reduce stress during the covid-19 pandemic," *Frontiers in Psychology*, vol. 12, 2021. [Online]. Available: https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2021.623110

[12] T.-L. Nguyen and M. Kamper-Jørgensen, "The sound of waves: tracking the pandemic," *The Lancet Infectious Diseases*, vol. 23, no. 1, p. 43, Jan. 2023, publisher: Elsevier. [Online]. Available: https://doi.org/10.1016/S1473-3099(22)00791-5

[13] P. Rebelo, "Covid-19 data sonification #65," in *Queen's University Belfast*, 2020, https://www.qub.ac.uk/coronavirus/research/related-research-activity/coronavirus-data-sonification/.

[14] I. Thompson and F. Hollerweger, "Pure covidata: An automated sonification of the covid-19 pandemic's recent state," in *Proceedings of the 7th Interactive Sonification Workshop (ISon)*. CITEC, Bielefeld University, 2022, p. 75.

[15] E. Lemmon, M. Schedel, I. Bilkhu, H. Zhu, L. Escobar, and G. Aumoithe, "Mapping in the emergency: Designing a hyperlocal and socially conscious sonified map of covid-19 in suffolk county, new york," in *Proceedings of the 7th Interactive Sonification Workshop (ISon)*. CITEC, Bielefeld University, 2022, p. 93.

[16] R. Bresciani, "Vox aeterna – the covid variations," in *Proceedings of the International Conference on Auditory Display*, 2021.

[17] S. Roddy, "Embodied sonification," Ph.D. dissertation, University of Dublin, Trinity College, 2016.

[18] U. Jekosch, "Assigning meanings to sound: Semiotics in the context of product-sound design," in *Communication Acoustics*, J. Blauert, Ed. Berlin and Heidelberg: Springer-Verlag, 2005, ch. 8, pp. 193–221.

[19] HER: She Loves Data, "A Song for Corona," 2020. [Online]. Available: https://he-r.it/Corona/

[20] M. Geronazzo, A. Bedin, L. Brayda, C. Campus, and F. Avanzini, "Interactive spatial sonification for non-visual exploration of virtual maps," *International Journal of Human-Computer Studies*, vol. 85, pp. 4–15, 2016, data Sonification and Sound Design in Interactive Systems. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1071581915001287

[21] P. Paul, A. France, Y. Aoki, and et al., "Genomic surveillance for sars-cov-2 variants circulating in the united states, december 2020–may 2021," *MMWR Morb. Mortal. Wkly. Rep.*, vol. 70, no. 846–850, 2021.

[22] A. Lambrou, P. Shirk, M. Steele, and et al., "Genomic Surveillance for SARS-CoV-2 Variants: Predominance of the Delta (B.1.617.2) and Omicron (B.1.1.529) Variants — United States, June 2021–January 2022," *MMWR Morb. Mortal. Wkly. Rep.*, vol. 71, no. 206-211, 2022.

[23] K. Ma, P. Shirk, A. Lambrou, and et al., "Genomic Surveillance for SARS-CoV-2 Variants: Circulation of Omicron Lineages — United States, January 2022–May 2023," *MMWR Morb. Mortal. Wkly. Rep.*, vol. 72, no. 651-656, 2023.

[24] C. Paden, "SARS-CoV-2 Variant Proportions | Data," Center for Disease Control and Prevention, 2021, https://data.cdc.gov/Laboratory-Surveillance/SARS-CoV-2-Variant-Proportions/jr58-6ysp/about_data.

[25] A. Rambaut, E. C. Holmes, O'Toole, V. Hill, J. T. McCrone, C. Ruis, L. du Plessis, and O. G. Pybus, "A dynamic nomenclature proposal for SARS-CoV-2 lineages to assist genomic epidemiology," *Nature Microbiology*, vol. 5, no. 11, pp. 1403–1407, Nov. 2020. [Online]. Available: https://doi.org/10.1038/s41564-020-0770-5

[26] T. Carpentier, "A new implementation of Spat in Max," in *15th Sound and Music Computing Conference (SMC2018)*, Limassol, Cyprus, July 2018, pp. 184 – 191. [Online]. Available: https://hal.archives-ouvertes.fr/hal-02094499

[27] M. Huang, S. Chabot, C. Leitão, T. Krueger, and J. Braasch, "Spatially-aware group interaction design framework for collaborative room-oriented immersive systems," *Applied Ergonomics*, vol. 113, p. 104076, 2023.

[28] D. Kirsh, "Do architects and designers think about interactivity differently?" *ACM Trans. Comput.-Hum. Interact.*, vol. 26, no. 2, 2019.

[29] Z. Bian, M. L. Braunstein, and G. J. Andersen, "The ground dominance effect in the perception of 3-D layout," *Perception & Psychophysics*, vol. 67, no. 5, pp. 802–815, July 2005. [Online]. Available: https://doi.org/10.3758/BF03193534

[30] C. Spence and J. Driver, "Audiovisual links in exogenous covert spatial orienting," *Perception & Psychophysics*, vol. 59, no. 1, pp. 1–22, Jan. 1997.

[31] B. Cohen-Lhyver, S. Argentieri, and B. Gas, *Audition as a Trigger of Head Movements*. Cham: Springer International Publishing, 2020, pp. 697–731.