PERCEPTUALLY-MOTIVATED SONIFICATION OF SPATIOTEMPORALLY-DYNAMIC CFD DATA

Lucas Temor
University of Toronto
Biomedical Simulation Lab
Toronto, Canada
lucas.temor@mail.utoronto.ca

Daniel E. MacDonald
University of Toronto
Biomedical Simulation Lab
Toronto, Canada
dmacdo@mie.utoronto.ca

Thangam Natarajan
University of Toronto
Biomedical Simulation Lab
Toronto, Canada
thangam@mie.utoronto.ca

Peter W. Coppin
OCAD University
Perceptual Artifacts Lab
Toronto, Canada
pcoppin@faculty.ocadu.ca

David A. Steinman
University of Toronto
Biomedical Simulation Lab
Toronto, Canada
steinman@mie.utoronto.ca

ABSTRACT

Everyday perception and action are fundamentally multisensory. Despite this, the sole reliance on visualization for the representation of complex 3D spatiotemporal data is still widespread. In the past we have proposed various prototypes for the sonification of dense data from computational fluid dynamics (CFD) simulations of turbulent-like blood flow, but did not robustly consider the perception and associated meaning-making of the resultant sounds. To reduce some of the complexities of these data for sonification, in this work we present a feature-based approach, applying ideas from auditory scene analysis to sonify different data features along perceptually-separable auditory streams. As there are many possible features in these dense data, we followed the analogy of "caricature" to guide our definition and subsequent amplification of unique spectral and fluctuating features, while effectively minimizing the features common between simulations. This approach may allow for better insight into the behavior of flow instabilities when compared to our previous sonifications and/or visualizations, and additionally we observed benefits when some redundancy was maintained between modalities.

1. INTRODUCTION

The visualization of multidimensional spatiotemporal data presents a challenging representational problem that often overlooks the strengths of the multiple sensory channels of the human perceptual system. Various audio or audiovisual techniques have been proposed (e.g., [1, 2, 3, 4, 5]) to offload some of this visual information to the auditory channel in order to reduce problems introduced by visual overload and/or occlusion that are inherent to many spatiotemporal data, and to take advantage of the unique pattern identification abilities of the auditory system.

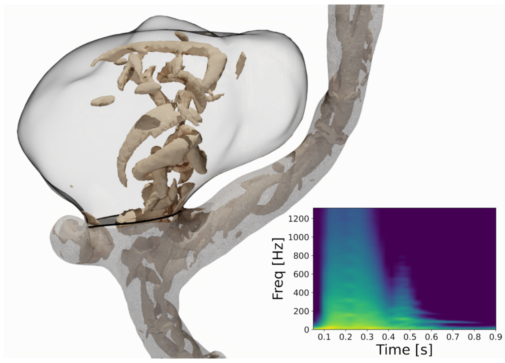

Figure 1: Visualization of Q-criterion (vortex core) isosurfaces in an aneurysm sac (left) with corresponding sac averaged spectrogram (right). Note: still frame from animations/sonifications at http://tinyurl.com/4w38sujx

1.1. Clinical motivation

Our interest in this area stems from our research into the use of computational fluid dynamics (CFD) for predicting cerebral aneurysm rupture. Briefly, a cerebral aneurysm is a region of localized arterial bulging within the brain. It is widely believed that complex blood flow patterns (hemodynamics) may be related to an aneurysm's growth and rupture [6], however these are difficult to measure non-invasively. As a result, image-based CFD techniques have been developed, whereby medical images are used as input for computational models [7] in order to simulate aneurysmal blood flow in a patient-specific manner [8]. Relatively recently, high-fidelity CFD simulations have revealed the presence of "turbulent-like" flow features in aneurysms, which may be an overlooked factor in their growth and rupture [9].

Such high frequency velocity fluctuations can be visualized spatiotemporally over a cardiac cycle using Q-criterion (revealing the location and persistence of vortical structures e.g., [10]) or in the frequency domain using spectrograms [11] (Figure 1). We have also presented various prototypes for sonifying these complex data

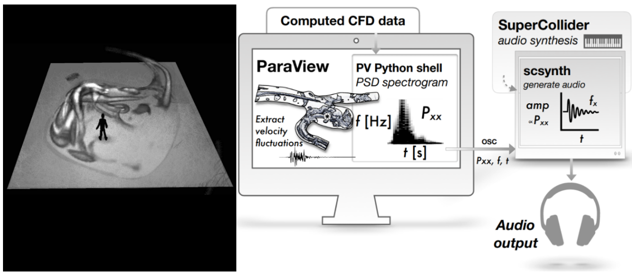

Figure 2: Examples of previous sonification prototypes. Left: sketch of ground plane representation, from [12]. Right: data-driven sonification of velocity spectrograms, from [13].

[12, 13]. Those prototypes focused mostly on developing sonification algorithms and aesthetics that, as expressed by Rabenhorst et al., *"permit the user to visually concentrate on one field while listening to another" [1]*, with less consideration given to how a user may be perceiving and making sense of the presented sounds, especially when presented in combination with a visualization.

1.2. Previous approaches to sonification of turbulent-like flow

One of our first sonification prototypes was presented at ICAD 2017 and virtually situated the user directly inside of the 3D velocity field of an aneurysm sac, sonifying data from velocity-time traces using "flow-inspired" auditory icons which aimed to indicate the presence of high energy flow structures [12]. That approach for sonifying spatiotemporal velocity data was informed by some of the existing (albeit limited) techniques for sonifying similar data (e.g., [2] and [3]) with respect to its algorithmic implementation and interactivity which allowed for local probing of 3D regions of interest. A conceptual sketch that motivated the work is shown in Figure 2, left, where the listener is situated on a "ground plane" and flow structures are moving around them. We will revisit this analogy in the current work (shown in Section 2.2).

At ICAD 2018, our paper presented a refinement of our previous approach, adopting a spectral analysis-resynthesis technique that sonified velocity spectrograms [13] (Figure 2, right). In that prototype, the user could again interactively probe various regions of vasculature in 3D, and velocity spectrograms were computed and synthesized in real time in a way that is analogous to a phase vocoder, achieving a fluid-like "bubbly" aesthetic which allowed for better detection of high frequency velocity fluctuations when compared to previous sonifications of single velocity-time traces.

Our 2018 ICAD paper was awarded *Best Student Paper*; however, it was cited in a 2019 editorial [14] where its mapping of multidimensional data to multiple low-level acoustic dimensions (i.e., frequency and loudness, a consequence of our approach to spectral synthesis) was questioned due to their inherent perceptual conflation (see, for example, [15] and [16]). Our decision to use the aforementioned spectral analysis-resynthesis technique was partially due to the fact that we had not yet defined exactly which features in our data we aimed to communicate, and so spectral analysis-resynthesis (which generally aims to synthesize sounds to be as realistic as possible, e.g., in computer music [17] or speech synthesis [18]) provided us with an established way of sonifying our data with little to no abstraction. The end effect of this was that the technique worked well for creating physically realistic sounds, for example, when presented to a patient with pulsatile tinnitus, a correspondence was reported between the sounds they had been hearing as a result of their pathology and the synthesized sounds [19]. As suggested in the editorial, however, the sonifications did not adequately communicate multidimensional data relationships. A prevalent issue was that the listener was only able to attend to a single feature in the data: the outer *spectral envelope* that rises and decays in response to variations in the cardiac cycle. Other relationships in the data (e.g., spectral harmonics, seen as horizontal striations in the spectrogram in Figure 1) were imperceptible.

1.3. "Caricatured" flow features and their sonification

The communication of a flow's spectral envelope by the previous sonification allowed for reasonable distinctions to be made between laminar and turbulent-like flows, however it was difficult to refine this classification further. As mentioned, other features in the data such as spectral harmonics, which we believe to be a marker of periodic vortex shedding, were not properly communicated. Furthermore, when presented in sync with a Q-criterion visualization, it became evident that the sonification was not communicating the inherent quasi-random nature of fluctuations in these turbulent-like flows.

These issues, their associated perceptual challenges, and some of our preliminary findings, which challenge the above Rabenhorst et al. quotation (namely, that our bimodal representations seem to be better served when there is some representational congruency, as opposed to complete orthogonality, between modalities) led us to consider the current feature-based approach. As suggested earlier, an immediate challenge that arises is the *definition* of these features. In our large multidimensional datasets, there are many possible features one could chose, and because direct associations between flow patterns and pathlogy have not yet been established, we do not necessarily know a priori which features are of greatest interest. To guide the definition of these features, we invoke the analogy of "caricature" as presented in [20]. Following this analogy, we aim to identify features for sonification by amplifying those that are "unique" to individual cases, and minimizing those that are common, which should in turn facilitate the classification of flow phenotypes with greater refinement. For example, in our velocity signals, the global carrier that drives all of these pulsatile flows (i.e., a heartbeat) is inherently common to every case, whereas the presence of harmonic structures or flow instabilities can vary, arguing for the latter features to be amplified in the sonification, while minimizing the former.

Once we have defined these features, we can sonify them along different auditory streams. This streaming approach in turn offers a convenient way of conceptualizing the sounds we are presenting within existing perceptual research, mainly that of auditory scene analysis [21], providing an established framework that we use to inform the design of novel sonification mappings. We further aim to design these mappings to bear ecological sound-source similarity to the flow phenomena they aim to describe [22], so as to avoid the reliance on low-level acoustic mappings. As will be discussed in Section 4.2, in the current work we have, for now, removed some of the interactivity from our existing pipeline so as to not be constrained by real-time computational limitations, in order to focus more on the perception and sense-making of the sounds.

In Section 2 we present our techniques for defining and algorithmically extracting features in our data based around spectrograms and the motion of Q-criterion vortex cores, and describe

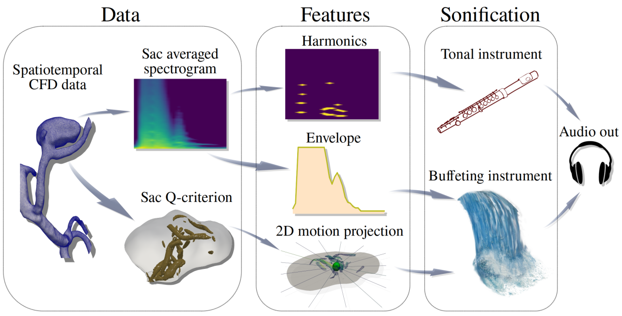

Figure 3: Overview of the presented sonification pipeline

how these drive acoustically distinct instrument models. In Section 3, we show how these new sounds allow for better classification of flow phenotypes than before. Finally, in Section 4 we discuss how a listener may be perceiving/making sense of these sounds, and directions for future end-user validation.

2. METHODS

The presented sonification pipeline is shown in Figure 3 and will be discussed throughout the remainder of this paper. First, CFD simulations are performed as described in [23] in patient specific models obtained from the Aneurisk dataset (http://ecm2.mathcs.emory.edu/aneuriskweb/index). With the resulting spatiotemporal velocity data, spectrograms are generated for the aneurysm sac using the techniques described in [11], and the scalar Q-criterion field is also computed in the sac. To prepare data features for sonification, harmonic content and the outer envelope are extracted from spectrograms (Section 2.1) and Q-criterion vortex cores are projected onto a 2D plane allowing for their motion and strength to be approximated (Section 2.2). This information (spectrogram envelope, spectrogram harmonics, average vortex core strength, approximate location of projected cores) is then sent to two Supercollider [24] instrument models and sonified along timbrally distinct streams: an acoustically complex (i.e., more than a simple sine wave) tonal instrument to sonify the presence and persistence of harmonic structures and a buffeting instrument modulated by fluctuations in the spectral envelope and Q-criterion motion (Section 2.3).

It is important to note that the decision to sonify a flow's the spectral envelope and Q-criterion motion together along one auditory stream does not contradict the physical relationships that are present in the data. The spectral envelope feature can be thought of as a global indicator of instability (analogous to a sort of carrier signal), to which local spatiotemporal fluctuations are added. As introduced in Section 1.3 the sonification of the envelope alone was not able to sufficiently communicate the quasi-random nature of these fluctuations, which resulted in monotony between different cases and/or limited classification of flow patterns within cases. Generally, these local spatiotemporal fluctuations are instead best captured using Q-criterion, which motivated its sonification. As vortical structures of Q-criterion will naturally stabilize during the diastolic portion of the cardiac cycle when flows decelerate, we cannot rely on them alone for communicating flow behavior over an entire cycle, necessitating the combined sonification of Q-criterion and the spectral envelope, further demonstrating the coupled relationship between these two features that we aim to capture with the presented mapping.

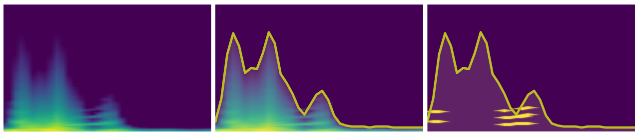

Figure 4: Feature extraction from spectrograms. Left: the original spectrogram. Center: extracted spectral envelope. Right: extracted harmonic features, and final caricatured representation

2.1. Extracting spectral features

Following our caricature analogy, we identified major spectrogram features that we believe facilitate their classification. The features identified are: (i) a spectrogram's outer envelope (as was already communicated in previous sonifications); (ii) the relationship between its frequency bands (i.e., the presence/persistence or absence of harmonics); (iii) and their relative strength compared to the other spectral content of the flow. This approach to extracting and preprocessing spectral features for sonification is conceptually reminiscent of some existing work on the sonification of EEG data [25], though most of the techniques do not translate directly due to differences in the definition and interpretation of features in these data between domains.

To extract the outer frequency envelope of the spectrogram, a binary threshold operation is applied to its associated matrix of power spectral density (PSD, in dB), such that ones are assigned to entries where the matrix is above a heuristic threshold power level, and zeros elsewhere. The binarized matrix is then traversed column-wise (i.e., along points in time) and the value of each column's largest non-zero frequency row is appended to a one dimensional array which corresponds to the temporal evolution of the spectrogram's envelope. We can see the results of this operation in Figure 4 - the thick line in the center image traces the outer edge of the spectrogram on the left.

Harmonic features in the data can broadly be understood as periodic peaks in column-wise traces of PSD vs. frequency that persist in time. To extract harmonic structures from spectrogram data, we perform the following steps along each column in its matrix. First, in order to differentiate smaller fluctuations in power from harmonic structures of interest, we approximate the global falloff of power by fitting a linear spline to successively decreasing minima along each power trace. This is then subtracted from the original signal which allows us to extract the height of all peaks relative to the baseline falloff of power (i.e., the strength of the peaks relative to their local power levels). To locate the presence and strength of harmonic striations, peaks with height above a heuristic threshold in this floored signal are identified. Finally, the original spectrogram is masked such that only points corresponding to these detected peaks are retained.

Once harmonic peaks are identified, they are temporally grouped and labeled in the masked spectrogram in order to drive the tonal instrument model described in Section 2.3. To do this, we aim to assign a unique identifier to each horizontal striation and then sonify each of these labelled vertically stacked bands as coherent temporal units (i.e., features), rather than scattered peaks. To label horizontal bands we first dilate the detected peak matrix in order to fill in any small discontinuities. We then identify regions

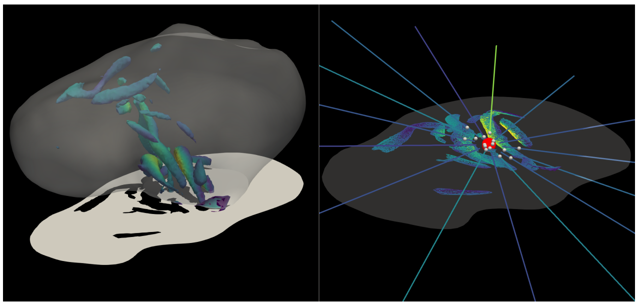

Figure 5: Visualization of Q-criterion projection. Left: aneurysm sac containing thresholded Q-criterion and its 2D projection. Right: radial lines show the directions along which the scalar field is averaged and are colored on its magnitude. The larger sphere indicates the geometric centre of the flattened sac and the smaller spheres show the centres of the intersecting geometry along each line. We emphasize the similarity to Figure 2, left.

of connected pixels (i.e., matrix entries), which we choose to be harmonic features in the spectrogram. Figure 4, right, shows the results of this operation. Harmonic bands that are present in the spectrogram on the left are extracted and "amplified " on the right. Overall, this shows a representative or caricatured visualization of the spectrogram features to be sonified.

2.2. Q-criterion as the primary driver of fluctuating sounds

To inform the present sonification of temporally-dynamic Q-criterion, we turned to an approach we had previously employed to sonify temporally-static Q-criterion data generated from a computational simulation of flow around wind turbine blades [26]. In that case, point-wise Q-criterion values were probed in 3D space at a single time step and used to control synthesis parameters for a buffeting noisy instrument model in a parameter mapping sonification designed to be ecologically evocative of moving wind. The "buffeting" aesthetic worked very well in that application as it allowed for a sonically-familiar way of making sense of the intensity flow instabilities through the interpretation of its quasi-random fluctuations. As we were aiming to obtain much the same insight from our current sonifications, we adapted that instrument model to meet our needs, modifying it to change with our temporal data as well as spatial locations of Q-criterion structures and their magnitude. More details of this parameter mapping are presented in Section 2.3.

To modulate the buffeting sound in a data-driven way, our initial approach considered sonification techniques based on motion tracking in videos. Mainly, we began by sonifying visualization videos of Q-criterion in a way almost identical to the motiongram technique presented in [27]. This work was reminiscent of some of our very early sonification prototypes which used a similar method of video analysis to drive instrument models (e.g., http://tinyurl.com/nujdub6y, http://tinyurl.com/ac9vjcyk). While this provided promising initial results advocating for the continued exploration of Q-criterion motion sonification, we realized that for our purposes, sonification of video has many inherent limitations as it is largely dependent on various subjective rendering decisions such as the choice of camera position. Furthermore the one-dimensionality of the motiongram technique often resulted in sonifications that sounded too similar between cases.

To overcome these limitations we decided to work with 2D projections of 3D Q-criterion structures. The use of 2D projections allows us to drive our sonification in much the same way as we were with video (which is also two dimensional), however is able to preserve features in the data such as spatial relationships or scalar magnitudes that are not retained in video, and does not require any animation rendering as a preprocessing step. This decision was further motivated by some of our early sonification prototypes such as Figure 2, left, where the listener is standing directly on a plane surrounded by flow structures, as well as existing work in our domain on the two-dimensionalization of 3D aneurysm geometries (e.g., [28]). The remainder of this section will briefly describe our approach for the two-dimensionalization and spatial discretization of Q-criterion structures as input for the buffeting instrument described in the next section.

To begin, we compute pointwise Q-criterion in the aneurysm sac over the entire cardiac cycle. The volumetric Q-criterion scalar field is then thresholded to only retain stronger vortical structures, which are written to unstructured mesh files at each timestep. The choice of scalar threshold for Q-criterion does involve some subjectivity, however, as long as a fairly inclusive threshold is chosen the resulting sonification should not be considerably changed, as lower-magnitude structures tend to be relatively stable, thus not contributing to huge fluctuations in motion.

Next, the plane of the aneurysm's ostium is defined and the entire sac geometry is flattened into 2D by projecting each of its points downward onto this plane along its normal direction. The points contained in the previously computed vortical structures (each one with an associated Q-criterion value) are also projected along this direction onto the same plane, preserving their magnitudes and connectivity (Figure 5, left). On this plane, we now want to approximate the location and strength of projected Q-criterion structures. At the plane's geometric centre (indicated by the larger sphere positioned where the radial lines converge in Figure 5, right) we cast out 16 rays spaced with equal angles along the flattened surface. Along each ray, we compute the average magnitude of Q-criterion as well as the location of the centres of the intersecting structures (smaller spheres in Figure 5, right). This results in a series of temporal arrays, each containing corresponding values of average Q-criterion and the distance of the intersecting structures to the centre of the plane for each of the angular positions.

2.3. Sonification mappings and sound synthesis

Once the data have been preprocessed, all relevant arrays (spectrogram envelope, harmonic features, average Q-criterion, core location) are sent to a Supercollider audio server which handles all sound synthesis using a pipeline similar to what we presented in our 2018 sonifications [13]. The sonification consists of two major components which may be broadly understood as the previously introduced buffeting sound and an acoustically complex tonal sound, each of which is implemented as a server side instrument model in Supercollider. These two streams are synthesized independently and mixed together to create the final sonification. The interpretation and perceptual implications of these sonification mappings will be discussed in Section 3 and Section 4, respectively.

To drive the tonal instrument, the extracted and labeled harmonic bands are mapped at each timestep to subsequent notes in a major chord, ascending in octave as the bands ascend in frequency.

To further facilitate the detection of the presence/absence of these tonal sounds amongst the noisy sounds described below, pitches are synthesized with acoustically complex timbres. For example, in the sonifications presented in Section 3, spectrogram harmonics are mapped to a sawtooth-triangle oscillator with rich overtones that are modulated and filtered to achieve a sort of glassy whistling sound. There are a variety of aesthetic choices that could be made for this instrument model to suit user preferences, however, regardless of this decision, it is essential that a clear timbral/acoustic distinction is maintained between the buffeting and tonal sounds.

The buffeting instrument is based on a frequency modulation (FM) synthesis technique which consists of a carrier brown noise source sent to a resonant low-pass filter (RLPF) which is modulated by an additional brown noise source. As shown in Figure 6, Q-criterion magnitudes are scaled and are added to the modulator signal which controls the cutoff frequency of the RLPF such that high values of Q-criterion will result in the synthesis of higher frequency noise which modulates more rapidly, giving the effect of more chaotic buffeting. The resonance of the filter and its total gain are also varied in time by the distance to centre of the Q-criterion cores along each cast ray, such that closer cores will sound louder and create more resonant peaks in the synthesized sounds.

Along the two dimensional Q-criterion plane, each angular position is mapped to a different spatial location within the plane of the user's listening environment. The listening perspective is situated at the centre of the flattened model (i.e., the sphere in the centre of Figure 5, right). Along each line, average Q-criterion values and their distance to the centre are used to control the buffeting instrument model. The use of spatialization here is mainly to create the impression of motion, for example different instrument models may periodically turn on or off as a core swirls around on the plane. As will be discussed in the next section, the goal of such a mapping is not to permit the user to attend individually to each of the 16 spatialized locations or to access information (e.g., make loudness judgments) along a certain direction; rather, it is to create an auditory stream that provides insight into the nature of quasi-random fluctuations in the flow, as well as help to anchor to the 3D Q-criterion visualization, giving an overall impression of the flow behavior.

To sonify the envelope feature in a way that is coupled to the Q-criterion motion, an additional buffeting instrument model is added to the motion projection sonification that is controlled by the extracted spectral envelope. In a similar fashion to Q-criterion magnitude, the outer spectral frequency envelope (i.e., Figure 4, center) is used to modulate the brown noise signal that modulates the RLPF cutoff along each angular direction.

It is important to note that the musical vocabulary used to describe the tonal sonification mapping is not essential for the interpretation of the sounds, i.e., the user is not expected to be able to attend to variations in the sounds of the chords generated. What we are interested in is for the listener to be able to temporally detect the *presence* or *absence* of harmonic structures. A musician or trained user could learn to attend to various differences in chord voicing, but this would only, at the most, provide minimal supplementary information about the relationships between harmonics, or at the least, aesthetic support.

3. RESULTS

A video playlist containing bimodal examples of the presented sonification pipeline synchronized with Q-criterion visualizations

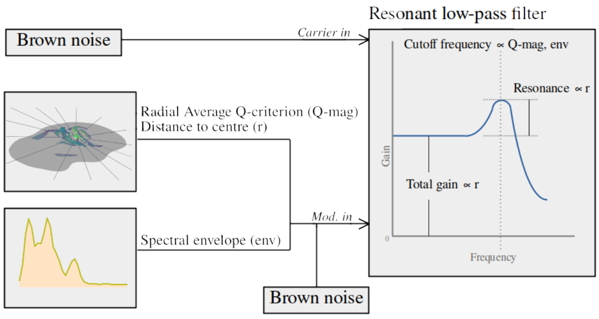

Figure 6: Buffeting instrument model for sonifying Q-criterion motion with spectral envelope

is found at http://tinyurl.com/4w38sujx. In general, the current sonification allows for classification of flow phenotypes with respect to the detection of broadband and harmonic spectral features, as well as improved identification of finer-scale spatiotemporal fluctuations in the data, offering a noticeable improvement when compared to previous prototypes.

Similar to our previous sonifications, the present sonifications allow for classification of flows into laminar and turbulent-like, allowing for the listener to attend to variations in its spectral envelope. For example, contrast the sonifications of case 0053 to case 0032. It should be quite straightforward to gather that the former exhibits much more chaotic flow patterns as is indicated by its “noisier” synthesized sound, both spectrally and with respect to Q-criterion motion, facilitating a quick classification of the former as turbulent-like and the latter as laminar.

The first way in which the present sonification improves upon its predecessors is in its ability to allow for the identification of spectral harmonic structures which are associated with periodic vortex shedding phenomena. When we compare variations in the spectral envelope of case 0053 to case 0060, we note that both rise and decay in a similar fashion, which may lead us to identify both as turbulent-like. We can now further refine this classification as we detect the presence of tonal sounds in the former that are not present in the latter. When viewed in sync with the Q-criterion visualization, the presence of this tonal sound helps to reinforce that the nature of the vortical structures in the former are the result of a more pure, periodic vortex shedding, rather than those of the latter which suggest more chaotic turbulent-like flow, a task which is difficult to do through Q-criterion visualization alone, but facilitated through multiple modalities.

Secondly, changes in the buffeting sounds produced by Q-criterion motion may also provide reinforced information on the fluctuating nature of instabilities and/or their associated vortex core behavior. Consider the sonifications and associated visualizations of case 0053 and case 0004. We can see that spectrally, the envelopes of these two cases are quite similar and both show harmonic structures, however the associated nature of their flow patterns is quite different. In case 0053 we visually observe coherent vortical structures being shed periodically just past the aneurysm ostium. In case 0004, we observe one larger core that forms in the sac that appears to periodically fluctuate and rotate in place. When we attend to the associated sonification of Q-criterion in sync with the visualization, the nature of these fluctuations becomes highlighted in a way that is difficult through visualization alone. In case 0004 the periodicity of the buffeting sound seems to sync with the rotations in the visualized core movement, and in case 0053 the buffeting seems to sync with the frequency of the periodic shedding. Here, the auditory cues seem to allow for the perception of motion in the visualization to be reinforced in a way that is not possible through a single modality, instead suggesting that there are benefits to audiovisual redundancy which allow the user to integrate information from either modality to form a more complete mental representation of the flow patterns.

4. DISCUSSION

We have presented a pipeline for the extraction and sonification of features from dense spatioteporal data. This feature-based approach provided us with a sort of "dimensionality reduction", which eliminated our need to work with/sonify large portions of these multidimensional datasets, more easily facilitating their classification through the identification and communication of unique features. This brings us back to the introduced analogy of caricature [20], again highlighting that the aforementioned classification is facilitated by the amplification of defining features in the data and minimizing those features which are closer to the sort of metaphysical "mean" of its feature space. Conceptually, this is how the current sonification has moved forward. Our previous spectral synthesis technique had close physical/acoustic similarity to the actual flow phenomena, but failed to communicate relationships between multidimensional data features. In the current work, we reduced the total amount of data we simultaneously translate into sound, but on the whole have achieved a gain in the information we are able to access from these sounds (e.g., we can derive the same insight into flow behavior from sonifying an entire 2D spectrogram matrix as we can from sonifying its 1D envelope feature).

As presented in Section 3, the distinctions that can be made between case 0004 and case 0053 caused us to revisit the Rabenhorst et al. quotation presented in the Introduction, which motivated some of our sonification work in the past. With respect to the use of bimodality, the idea that we can allow the user to *"visually concentrate on one field while listening to another"* did not seem to hold for our previous sonifications. For example, if we were to simply add a spectral sonification (generated using our previous methods) to a Q-criterion visualization, it was difficult to localize and/or make sense of finer spatiotemporal fluctuations within models. When we instead (as with the current pipeline) maintain some causality between the sounds presented and the visualization (in our case, this is effected by Q-criterion) these localized spatiotemporal fluctuations are highlighted in a way that seems to be superior to the presentation of different information to different sensory modalities, which is in line with how we make sense of spatiotemporally-dynamic stimuli in everyday scenarios. This advocates not necessarily for a completely redundant or completely orthogonal presentation of information, but rather one that at the least facilitates some sort of unity between the senses.

4.1. Perceptual aspects of the sonifications

In line with some existing sonification design suggestions (e.g., [5, 29, 30]), guiding most of the current work are ideas taken from perceptual research in he domains of auditory scene analysis [21] as well as ecological psychoacoustics [31]: the former informing the auditory streaming approach that we applied to better understand how to sonify unique features along different streams in a way that facilitates their segregation; and the latter guiding the way in which we consider the sense-making of the streams once they have been perceptually segregated. We took this approach in order to try and minimize the multiplicity of low-level perceptual conflations that arise when using mappings that rely on simple acoustic dimensions. This section will contextualize the sonification mappings presented in Section 2.3 within models of auditory perception in order to advocate for their perceptual efficacy. As we will discuss in Section 4.3, the intention of this is not to replace design validation that can only be achieved through formal user testing, but rather to demonstrate how we have applied ideas from auditory perceptual research to refine the design of our sonifications prior to formal user testing, providing us with a theoretical foundation upon which we can design future validation studies.

We chose to break down the synthesis of our sounds into separate tonal and noisy streams in order to be able to communicate both the global spectral evolution of the flow and its more local spatial fluctuations, as well as their nature (e.g., periodic vortex shedding vs. "turbulence"). As suggested by auditory scene analysis, these apparent timbral/acoustic differences may help to facilitate their segregation and the user's ability to attend to information along them, especially when we consider the schema-based and/or attentional implications of this segregation. That is, because of their timbral differences, we believe the primitive scene analysis system is able to partition these unique acoustic groups (i.e., noisy and tonal) within the presented sonifications, and the schema/attentional based system allows for the user to tune into (i.e., segregate) either stream independently to access information with minimal interference. It is to this extent, in much the same way that a user can make attentional-based visual queries with a visualization, one can auditorily query streams of information in this sonification.

It has been shown that under certain conditions musicians outperform non-musicians at timbral stream segregation tasks involving the identification of a mistuned harmonic in concurrent sounds [32]; however, we believe that the acoustic differences in the two streams presented in our sonifications are large enough that these differences in performance are minimized. For example, in the above referenced work, as harmonic mistuning between concurrent sounds increased, the likelihood of identifying two separate sounds increased within both groups, and differences between groups further decreased as harmonic mistuning approached 16%. Because our sonification does not rely on segregation of such subtle harmonic differences (rather, vast distinctions of "tonal" or "noisy"), these differences in segregation ability should not greatly impact the segregation performance of non-musicians.

Once a stream has been perceptually segregated, the user may infer properties of the flow behavior by listening to changes in its sound. We again make the important note that the goal of our sonifiations is not to judge absolute magnitudes, but rather to facilitate a "caricatured" perception of the flow features, allowing for different auditory streams to be queried to gain insight into their independent behavior, and their relationships. For example, we do not expect the listener to be able to report the magnitude of the average Q-criterion values in the sonification; instead we aim to facilitate a relative identification of flow behavior between models (e.g, the differences in fluctuation behavior between case 0053 and case 0004, as previously discussed). We believe that our sonification mappings allow for this because they were chosen to bear ecological (sound-source) similarity to the phenomena they are trying to describe [22]. As introduced in Section 2.2, the buffeting sound provides a familiar ecological anchor for anyone who has been in the presence of heavy winds or fast moving water, both of which can be associated with turbulent flow phenomena. Based on these ecological experiences of *everyday listening* [33] it should be a familiar task for a user to map perceived changes in the buffeting sound to perceived changes in the flow behavior, functioning on the assumption that users have an inherent intuition for the ways in which acoustic changes in a fluid-like sound relate to changes in its dynamics. A similar rationale can be applied to understand the ecological-perceptual implications of the tonal sounds, which in everyday listening environments are often associated with vibrational phenomena (e.g., strings on a musical instrument). Such tonal bruits have been observed for decades in invivo vibrating aneurysms, within the audible range of human hearing [34], which we currently hypothesize may be associated with the harmonic structures that are present in our spectrograms [35]. The choice of this tonal sonification mapping therefore offers a cyclic representational schema wherein the associated ecological sounds and hypothesized clinical implications are tied through a common understanding of this single auditory stream.

4.2. Limitations

The presented approach inherently does not allow for interactive probing of the data as previous prototypes did. We chose to focus only on sonifying the data in the aneurysm sac, and do not allow for the user to "listen" to the flows at different locations. We instead approximate the function of probing through our technique for the flat mapping and spatial discretization of the entire aneurysm sac and Q-criterion structures. While it would be fairly straightforward to implement the presented pipeline and sonification mappings in a real time environment (e.g., allowing for the interactive probing and flattening of localized 3D spatiotemporal Q-criterion structures), this can be a computationally-intensive undertaking. We chose not to take this approach in order to limit our design scope to focus on the perception and sense-making of the sounds as well as to reconsider how we present our sonifications to clinical collaborators. From an implementation perspective, the use of pre-rendered audio and video is much easier for clinical users to access as opposed to more intensive interactive 3D environments of which they are not familiar with nor have the time to learn to use. This is not to say that we are abandoning the use of interactivity at all, we are indeed aware of the possible benefits of allowing for interaction in sonification interfaces (e.g., [36]), rather, we aim to reevaluate what this sort of interactivity may look like in a clinical setting, and then begin fit our existing sonification models into that.

4.3. Future directions

We have previously expressed the intention to run more formal user studies [13] to better understand the utility of sonification in CFD/clinical workflows. As we began to think more about how the sounds of our previous sonifications were being perceived and interpreted, as well as noticing that there were certain data features that were not being communicated, it became apparent that conducting a formal user study would have been premature, and that more attention needed to be directed to the actual perception and meaning-making of these sounds. Now that we have begun to address this problem, the possibility of future user studies appears to offer much more meaningful insight into how these sonifications may prove useful in a clinical environment.

Recent correspondence with our clinical collaborators shows an openness to integrating sound into their regular workflows (especially when links are drawn between this and Doppler ultrasound auscultation, an existing audification technique that has been used in clinical practice to identify turbulent-like blood flows). We further note that we have observed evidence of existing cross-sensory intuitions between Q-criterion visualizations and sound in interactions with one of our clinical collaborators who noted that they were able to *"see the sounds"* that were being generated in a Q-criterion visualization. When asked to elaborate on this, they explained that looking at the visualization is reminiscent of looking at a vibrating subwoofer, suggesting an existing intuition for associations between the physical vibrational phenomena, visualizations, and sounds. Future user studies should further elicit these existing associations that clinical users may have for sound-source acoustic phenomena and integrate these findings into the design and presentation of bimodal representations, ultimately towards a user study which will test the possible utility of bimodal workflows in clinical practice, and perhaps other scientific or industrial domains.

5. REFERENCES

[1] D. A. Rabenhorst, E. J. Farrell, D. H. Jameson, T. D. Linton Jr, and J. A. Mandelman, "Complementary visualization and sonification of multidimensional data," in *Extracting Meaning from Complex Data: Processing, Display, Interaction*, vol. 1259. Int. Soc. for Optics and Photonics, 1990.

[2] M. Kazakevich, P. Boulanger, W. F. Bischof, and M. Garcia, "Augmentation of Visualisation Using Sonification: A Case Study in Computational Fluid Dynamics," in *Eurographics Symposium on Virtual Environments*, B. Froehlich, R. Blach, and R. van Liere, Eds., 2007.

[3] E. Klein and O. G. Staadt, "Sonification of three-dimensional vector fields," *UC Davis: Institute for Data Analysis and Visualization*, 2004.

[4] R. Minghim and A. R. Forrest, "An illustrated analysis of sonification for scientific visualisation," in *Proc. Visualization'95*. IEEE, 1995, pp. 110–117.

[5] K. Vogt, "Sonification in computational physics-qcd-audio," in *Proc. of SysMus Int. Conf. of Students of Systematic Musicology*, 2008.

[6] B. Staarmann, M. Smith, and C. J. Prestigiacomo, "Shear stress and aneurysms: a review," *Neurosurg. focus*, vol. 47, no. 1, 2019.

[7] C. A. Taylor and D. A. Steinman, "Image-based modeling of blood flow and vessel wall dynamics: applications, methods and future directions," *Ann. Biomed. Eng.*, vol. 38, no. 3, 2010.

[8] B. Chung and J. R. Cebral, "Cfd for evaluation and treatment planning of aneurysms: review of proposed clinical uses and their challenges," *Ann. Biomed. Eng.*, vol. 43, no. 1, 2015.

[9] K. Valen-Sendstad and D. A. Steinman, "Mind the gap: impact of computational fluid dynamics solution strategy on prediction of intracranial aneurysm hemodynamics and rupture status indicators," *Am. J. Neuroradiol.*, vol. 35, no. 3, 2014.

[10] T. Natarajan, D. E. MacDonald, M. Najafi, P. W. Coppin, and D. A. Steinman, "Spectral decomposition and illustration-inspired visualisation of highly disturbed cerebrovascular blood flow dynamics," *Comp. Methods in Biomech. and Biomed. Eng.: Imaging & Vis.*, vol. 8, no. 2, 2020.

[11] T. Natarajan, D. E. MacDonald, M. Najafi, M. O. Khan, and D. A. Steinman, "On the spectrographic representation of cardiovascular flow instabilities," *J. Biomech.*, vol. 110, 2020.

[12] P. W. Coppin, R. C. Windeyer, D. E. MacDonald, and D. A. Steinman, "Progress toward sonifying napoleon's march and fluid flow simulations through binaural horizons," in *Proc. ICAD*, 2017.

[13] D. E. MacDonald, T. Natarajan, R. C. Windeyer, P. Coppin, and D. A. Steinman, "Data-driven sonification of cfd aneurysm models," in *Proc. ICAD*, 2018.

[14] J. G. Neuhoff, "Is sonification doomed to fail?" in *Proc. ICAD*, 2019.

[15] J. W. Grau and D. G. Kemler Nelson, "The distinction between integral and separable dimensions: evidence for the integrality of pitch and loudness," *J. Exp. Psychol. Gen.*, vol. 117, no. 4, pp. 347–370, Dec 1988.

[16] J. G. Neuhoff, M. K. McBeath, and W. C. Wanzie, "Dynamic frequency change influences loudness perception: a central, analytic process," *J. Exp. Psychol. Hum. Percept. Perform.*, vol. 25, no. 4, pp. 1050–1059, Aug 1999.

[17] P. Masri, A. Bateman, and N. Canagarajah, "A review of time–frequency representations, with application to sound/music analysis–resynthesis," *Organised Sound*, vol. 2, no. 3, pp. 193–205, 1997.

[18] D. H. Klatt, "Review of text-to-speech conversion for english," *J. Acoust. Soc. Am.*, vol. 82, no. 3, pp. 737–793, 1987.

[19] V. M. Pereira, N. M. Cancelliere, M. Najafi, D. MacDonald, T. Natarajan, I. Radovanovic, T. Krings, J. Rutka, P. Nicholson, and D. A. Steinman, "Torrents of torment: turbulence as a mechanism of pulsatile tinnitus secondary to venous stenosis revealed by high-fidelity computational fluid dynamics," *J. Neurointerv. Surg.*, Nov 2020.

[20] V. S. Ramachandran and W. Hirstein, "The science of art: A neurological theory of aesthetic experience," *J. consciousness Studies*, vol. 6, no. 6-7, pp. 15–51, 1999.

[21] A. S. Bregman, *Auditory scene analysis: The perceptual organization of sound*. MIT press, 1994.

[22] B. N. Walker and G. Kramer, "Ecological psychoacoustics and auditory displays: Hearing, grouping, and meaning making," *Ecological psychoacoustics*, pp. 150–175, 2004.

[23] C. Chnafa, P. Bouillot, O. Brina, M. Najafi, B. M. A. Delattre, M. I. Vargas, V. M. Pereira, and D. A. Steinman, "Errors in power-law estimations of inflow rates for intracranial aneurysm cfd," *J. Biomech.*, vol. 80, 2018.

[24] S. Wilson, D. Cottle, and N. Collins, *The SuperCollider Book*. The MIT Press, 2011.

[25] A. Väljamäe, T. Steffert, S. Holland, X. Marimon, R. Benitez, S. Mealla, A. Oliveira, and S. Jordà, "A review of realtime eeg sonification research," in *Proc. ICAD*, 2013.

[26] D. E. MacDonald, "Data-driven sonification of cfd aneurysm models," Master's thesis, U. of Toronto, 2018.

[27] A. R. Jensenius, "Motion-sound interaction using sonification based on motiongrams," in *Proc. Int. Conf. on Advances in Computer-Human Interactions*, 2012.

[28] M. Meuschke, S. Voss, O. Beuing, B. Preim, and K. Lawonn, "Combined visualization of vessel deformation and hemodynamics in cerebral aneurysms," *IEEE Trans. Vis. Comput. Graph.*, vol. 23, no. 1, pp. 761–770, Jan 2017.

[29] G. Kramer, B. Walker, T. Bonebright, P. Cook, J. H. Flowers, N. Miner, and J. Neuhoff, "Sonification report: Status of the field and research agenda," 1999.

[30] T. Hildebrandt, F. Amerbauer, and S. Rinderle-Ma, "Combining sonification and visualization for the analysis of process execution data," in *IEEE Conf. on Bus. Inform.*, 2016.

[31] J. Neuhoff, *Ecological psychoacoustics*. Brill, 2004.

[32] B. R. Zendel and C. Alain, "The influence of lifelong musicianship on neurophysiological measures of concurrent sound segregation," *J. Cogn. Neurosci.*, vol. 25, no. 4, 2013.

[33] W. W. Gaver, "What in the world do we hear?: An ecological approach to auditory event perception," *Ecological psychology*, vol. 5, no. 1, pp. 1–29, 1993.

[34] R. Aaslid and H. Nornes, "Musical murmurs in human cerebral arteries after subarachnoid hemorrhage," *J. Neurosurg.*, vol. 60, no. 1, pp. 32–36, Jan 1984.

[35] D. E. MacDonald, M. Najafi, L. Temor, and D. A. Steinman, "Spectral bandedness in high fiedlity cfd predicts rupture status in intracranial bifurication aneurysms," in *Summer Biomech., Bioeng. and Biotransp. Conf.*, 2021.

[36] T. Hermann and A. Hunt, "The importance of interaction in sonification," in *Proc. ICAD*, 2004.