

Web-based Visualization of Long-term Ocean Acoustic Observations and Modeled Soundscapes

Thomas Butkiewicz, Colin Ware, Ilya Atkin, Jennifer Miksis-Olds, Anthony Lyons
Center for Coastal and Ocean Mapping
University of New Hampshire
Durham, New Hampshire, USA

Abstract— This paper describes the design and implementation of an interactive web-based geospatial and acoustic visualization interface for the Atlantic Deepwater Ecosystem Observatory Network (ADEON) project. An array of hydrophones deployed along the U.S. Mid- and South Atlantic Outer Continental Shelf collected four years of acoustic measurements. The overwhelming size of the resulting datasets (over 73 terabytes) poses an accessibility problem for researchers and the public. A web-based visualization interface was developed to allow anyone to easily explore the massive datasets generated by the project in order to gain insight about the ecology and soundscape of the region.

Keywords—web-based visualization, ocean noise, acoustics, spectrograms, marine mammal detection, visual analysis

I. INTRODUCTION

The Atlantic Deepwater Ecosystem Observatory Network (ADEON) is an array of hydrophones deployed along the U.S. Mid- and South Atlantic Outer Continental Shelf, that collected four years of acoustic measurements of natural and human sounds in the region from 2017 to 2020. There were seven physical landers deployed in the network, each with multiple directional and omnidirectional hydrophones that recorded various frequency ranges. In total, these hydrophones generated 73 terabytes of audio files.

This ADEON data needed to be made available to the public, but its massive size poses many accessibility and visualization issues. It is difficult for researchers to download all of this data, due to both bandwidth and physical data storage limitations, and finding what one is looking for in the files can be just as difficult, due to the overwhelming number and length of the recordings.

The recordings were further processed to produce additional datasets, such as event detections based on different filters (e.g. dolphin clicks or seismic surveys). The project also had a soundscape modelling component that generated 5D (lat/long/depth/frequency/time) datasets of predicted sound energy levels for sound sources such as ships and surface winds. This visualization project sought to build a single, cohesive visual interface for researchers both within and outside of the ADEON project to easily explore the acoustic data and the derived datasets, without having to download and store the data or install any particular software.

Study concept, oversight, and funding were provided by the U.S. Department of the Interior, Bureau of Ocean Energy Management, Environmental Studies Program, Washington, DC under contract number M16PC00003, in partnership with other NOPP funding agencies.

II. MAPPING INTERFACE

A. Main Map

Visitors to the ADEON website (<https://adeon.unh.edu/map>) can access an interactive map of the project region, through which the project's dataset can be explored using a number of different visualization interfaces. The map initially displays the locations of the hydrophone deployments over a hill-shaded bathymetric map, as shown in Fig. 1. The bathymetric map helps explain the intentional placement of the hydrophones in relation to the continental shelf.

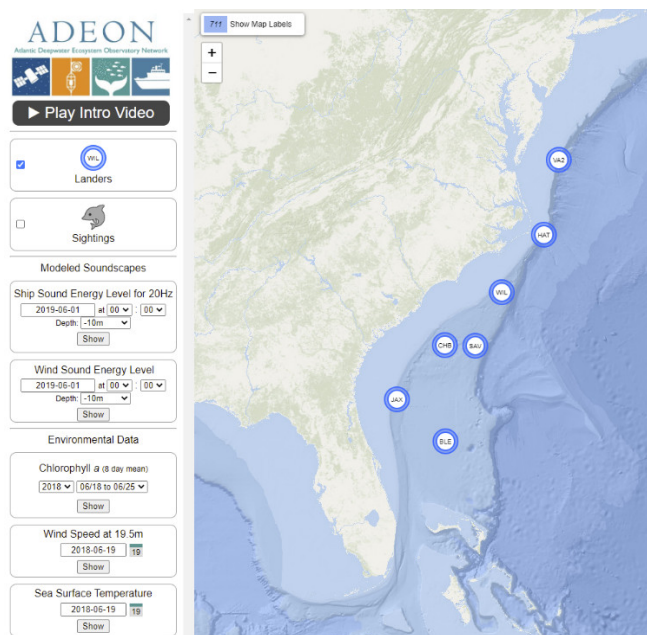


Fig. 1. The initial view of the ADEON interactive map, showing hydrophone locations over a hill-shaded bathymetry map that reveals the continental shelf.

The underlying data is processed on a Linux-based server running CentOS 8. Mapping data is served using GeoServer via Web Mapping Service for raster data (e.g. sea surface temperature), and Web Feature Service for point data (e.g. animal sighting locations). The Python-based Django web framework is used both as an interface to a PostgreSQL database (which stores detection and supplemental data) and to serve the mapping site itself. The web map is built using the Leaflet JavaScript mapping library, enhanced with new (developed for this project) components for animated time sliders and environmental data legends.

A sidebar interface allows users to toggle on and off the display of additional spatial datasets. This includes modeled soundscapes generated as part of the ADEON project, as well as environmental data layers, which provide important contextual information about the region during the times-of-interest being explored, such as the passing of a hurricane, or where the Gulf Stream was bringing warmer water.

Users can choose to display marine animal sightings on the map. These are point-based records of animal sightings that occurred during the various ADEON cruises. This feature displays small icons of various animals on the map where they were spotted. As shown in Fig. 2, these appear mostly around the lander locations. Clicking any of these animal icons will pop up an info panel with detailed information about that sighting event (e.g. species and time sighted).

B. Contextual Environmental Data Layers

Much of the sound detected at landers can be traced back to environmental factors in the surrounding sea and the atmosphere above. For example, surface winds can create high frequency sound, and the presence of high levels of chlorophyll can invigorate the food chain, attracting animals that make sound. As such, a number of helpful layers of environmental data are provided that can provide context and help users make sense of the datasets.

A chlorophyll layer, shown in Fig. 3, was extracted from NASA satellite data, and shows the weekly chlorophyll concentration in the region, the presence of which could significantly affect marine mammal activity and movements. Due to limitations of cloud cover, this data layer often contains many large “holes” and regions with no data values. For display, these holes are preserved as transparent, but behind-the-scenes values are interpolated to completely fill the layer. This is done so complete time-series values for particular regions (e.g. at a lander location) can be extracted for purposes of displaying contextual information in other interfaces (e.g. heat maps).

A wind speed layer, shown in Fig. 4, presents the near-sea surface (19.5m height) wind speed, extracted from a larger weather model. As surface winds are a significant source of high-frequency sound, this can provide helpful contextual information as to why some regions of time and space may be noisier. For example, Fig. 4 shows the passage of Hurricane Florence over ADEON hydrophones in September 2018.

Finally, a sea surface temperature layer is available, as shown in Fig. 5. These sea surface temperatures were extracted from NOAA’s RTOFS model. These temperatures influence marine wildlife activity levels and migration, and so are helpful contextual data when exploring the other datasets such as event detections of marine mammal calls. This layer also clearly reveals the location of the Gulf Stream current.

C. Modeled Soundscapes

The other main category of datasets that can be displayed on the main map is modeled soundscapes. These are 4D (lat/long/depth/time) model-derived soundscapes that show the sound energy levels at particular frequencies across the ADEON project region as modeled from sound sources such as wind and ships.

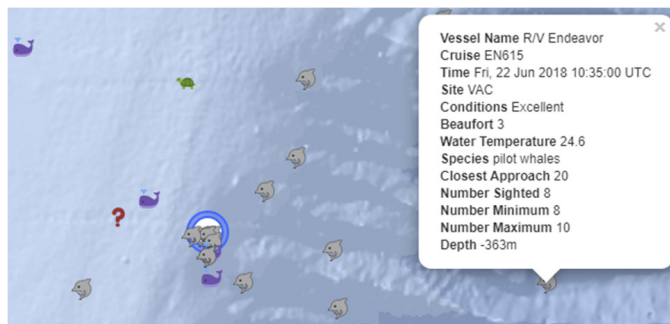


Fig. 2. Marine animal sightings around a lander site. One sighting has been selected, and an information panel shows the species and other details.




Fig. 3. ADEON map displaying chlorophyll data. Notice the incomplete coverage (due to cloud cover) is preserved in this presentation.




Fig. 4. ADEON map displaying a wind speed data from Sept 2018, in which Hurricane Florence passes over the ADEON network.




Fig. 5. ADEON map showing sea surface temperature data, revealing the path of the Gulf Stream over lander sites, and colder waters around the VA2 lander.

The wind soundscape, shown in Fig. 6, displays the contribution of surface winds to ocean sound at a variety of depths and at the sea floor. The ship soundscape, shown in Fig. 7, shows the modeled sound from surface ships (tracked using recorded AIS data) at different depths and at the sea floor.




Fig. 6. Modeled soundscape of wind sound at a depth of 10 meters.




Fig. 7. Modeled soundscape of ship sound (20Hz) at the seafloor in the ADEON region.

D. Additional Components

Two custom Leaflet library components were developed and released back into the open source software community, so that others can reuse them in their web-based visualization projects. The first component, shown in Fig. 8, generates and displays legends for the various custom color maps (scales) used in the environmental and soundscape layers. The second component, shown in Fig. 9, is a custom time-bar that is used to manipulate the time dimension of the mapping interface. It allows users to change the currently displayed time, step forward and backward through individual time steps, set time ranges for animations, and adjust the playback speed of animations.




Fig. 8. Leaflet component developed to provide custom color map legends.




Fig. 9. The time bar Leaflet component, developed to control which times are displayed, as well as the range and speed of animations.

Selecting individual lander locations on the map will pop up visualization interface windows that are visually linked to those landers and provide multiple ways to visualize the data associated with them. As shown in Fig. 10, multiple such interfaces can be opened simultaneously for multiple landers, and can be rearranged to permit side-by-side comparisons between landers. These panels initially display basic information about the selected lander, while a series of tabs along the top provide access to the different visualizations available, including *Heat Maps*, *Spectrograms*, and *Deviations*.

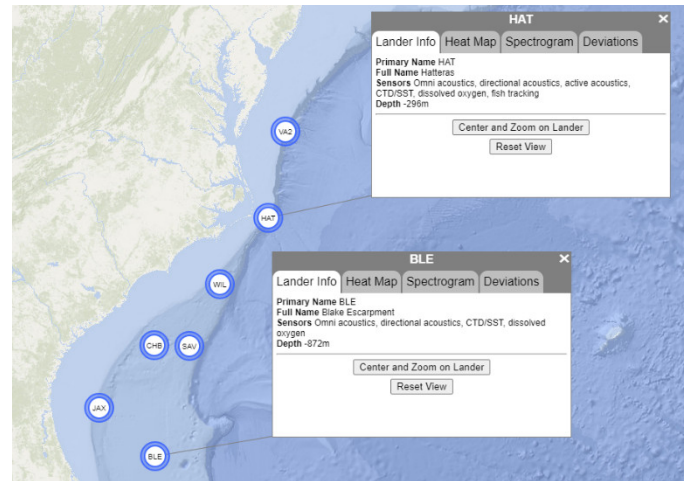


Fig. 10. : Example of two lander interface windows opened simultaneously. These are visually linked to the associated lander locations on the map.

III. HEAT MAPS FOR EVENT DETECTIONS

The heat map interface, shown in Fig. 11, provides an interactive visualization of the event detection datasets available for a lander location. The central heat map plot uses color intensity to indicate the number of detections that occurred within each one-hour cell. The plot is 24 cells tall, i.e. each column of cells represents a single day. While there was sufficient screen real estate to have more temporal granularity (e.g. 5 minutes/cell), this size better matched the temporal granularity and event density of the datasets in this project.




Fig. 11. Example view of the heat map interface, showing a continuous plot of dolphin clicks detected around a lander over a two year period. Notice that dolphin activity appears to correlate with the night time hours (line-shaded band across the plot) and cooler water temperatures (green vs orange/yellow in the contextual data bands above and below the plot).

An optional day/night indicator can be toggled on/off, and can be helpful for revealing daylight-dependent animal behavior patterns, as can be seen with dolphin clicks in Fig. 11. This option draws a wavy band across the plot, which indicates the hours of darkness, as calculated based on the lander’s latitude and longitude. After some experimentation, it was decided that a visual style of using subtle, gray horizontal lines, behind the filled cells, was sufficient enough to provide this indication, without being distracting or imposing too much relative visual weight relative to the actual data.

The plot can be toggled between either a ‘continuous’ mode, in which the plot expands horizontally to fill the available screen space and displays all records in temporal order, or a ‘cyclic’ mode, which stacks each year of data on top of the other in a single plot 365 cells wide.

The cyclic mode helps reveal patterns which repeat each year, and an ‘emphasis’ feature allows users to mouse over the year labels to indicate the contribution of each individual year to the overall stacked plot. As shown in Fig. 12, this is done by desaturating the content of cells with content from other years, thus highlighting the contribution of the selected year, while preserving the context of the other years.

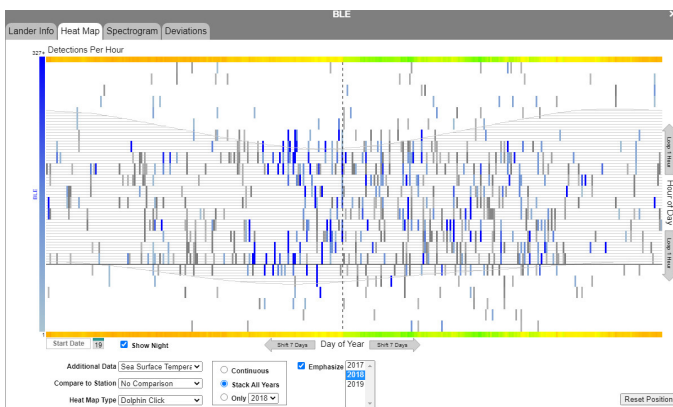


Fig. 12. View of the heap map interface in stacked mode, showing dolphin clicks from all years (gray cells), but emphasizing those from 2018 (blue cells). Notice the plot has been shifted to center the winter/night pattern, instead of wrapping it around the default Dec 31st/Jan 1st & midnight boundaries.

Arrow buttons below and to the side of the plot can be used to adjust the plot vertically and horizontally to change when wrapping occurs in terms of day of the year and time of day. This is useful because patterns can often straddle arbitrary midnight and December 31st/January 1st wrapping boundaries, which can cause them to appear disconnected. As illustrated in Fig. 12, shifting the wrapping boundaries to noon and July 1st can make a pattern centered around winter nights appear whole.

Contextual data bars can be added redundantly to both the top and bottom of the plot to indicate environmental values at the lander location during the times shown in the pot. This can reveal relationships between the patterns in the plot and surrounding conditions such as temperature, wind speeds, and chlorophyll. Fig. 11 and Fig. 12 show examples of contextual data bars indicating sea surface temperature. Values could also be extracted from modeled soundscapes and used this way, e.g. to evaluate model performance.

Comparisons between landers can be made by selecting a second lander in a drop down box. This draws the second lander’s data using another color. A question that arises is how to best display two values in a single cell. After some experimentation, it was decided that cells with data from multiple landers would automatically subdivide, presenting the two colors individually, as shown in Fig. 13. This allows users to visually search for, and see the overall distribution of, each color. While subdivided cells add visual noise, alternative solutions that mixed colors were difficult to interpret.




Fig. 13. Example heat map comparing event detections at two landers.

Users can mouse-over cells in the heat map plot to get additional information about them. As shown in Fig. 14, a small popup window follows the mouse cursor and displays the time that cell represents, and the number of detections in that hour period for the selected event type. If the plot is showing a comparison with another lander, mousing over subdivided cells shows the numbers of detections for both landers. For stacked plots with multiple years per cell, these popups show small bar graphs, which break down detections by year. If multiple event types are displayed, these graphs further break detections by event type.




Fig. 14. Small popup windows follow the mouse over the heat map, showing the time associated with each cell, as well as the number of events detected during that time. If a comparison is being made between two landers (center) the popup shows the counts for each lander. If the plot is stacking multiple years, bar graphs (right) show the distribution of events by year.

Clicking on a heat map cell switches the visualization interface to the Spectrogram Viewer (described in the following section), and automatically centers the spectrogram views onto the time associated with the selected heat map cell. This allows users to quickly find and examine the raw audio that triggered an event detection filter.

IV. TRI-LEVEL SPECTROGRAM VIEWER

The second visualization tab contains a multi-level spectrogram viewer that presents linked spectrograms at three different time-scales. This interface enables rapid exploration of multiple years of raw recordings. An example view of this interface is shown in Fig. 15.

The topmost level of spectrograms shows a few weeks on screen at a time, the middle level shows a few hours, and the bottom level shows a few minutes. (The exact amounts of time shown in each level varies based on the available screen resolution of the user's display.) All levels are linked; users can quickly scroll through months or years of data in the top level until something catches their eye. Then, they can click on a region of interest, which centers the other two levels on that time. Likewise, clicking on a region of interest in the second level will further refine the time shown in the other levels.

This multi-view, multi-resolution design facilitates exploratory analysis by letting users browse high-level overviews of the data, while being able to instantly drill-down into regions of interest to get details on demand. Instead of having to download terabytes of data, only small (~500KB - 1MB) compressed image files need to be transferred from the server as needed. By pre-caching a few image files on either side of the currently active view in each level, on most broadband internet connections, scrolling is generally a seamless experience, as users do not need to wait for data to load.

As can be seen in Fig. 15, there are highlighted regions in both the top and middle levels which indicate the portion of that level which is currently being shown in the level below it. A black line also follows in the upper regions to indicate the time displayed at the lowest level.




Fig. 15. Example view of the spectrogram interface running on a high resolution (2560x1600) monitor. The big orange spikes seen in the top and middle levels are ships transiting the lander region, while the small repeating orange blips near the bottoms of the middle and lower levels are fin whale calls.

Because the recordings do not continuously/completely cover all the time (e.g. there were ~5 minute gaps between many files), it was important to indicate to users where gaps were, while not disconnecting the data visually on either side of the gap, or introducing unnecessary clutter. As shown in Fig. 16, a small black bar was included between files with short gaps, with the bar's width equal to the missing time. However, for longer gaps, instead of a large black region that would overly-disconnect data visually (and waste screen real-estate), gap widths were limited to 100 pixels and a label was inserted to indicate the total amount of missing time.




Fig. 16. Spectrogram viewer showing three individual recording files and the gaps between them. The gap between the center file and the right file is very short, and is simply blacked out. However, the five-minute gap between the center file and the left file is condensed and a label indicates its actual length.

Mousing over the bottom level shows the exact time and frequency at that point in the spectrogram. Selecting a 1D region-of-interest within the bottom spectrogram enables users to play or download an audio clip of that particular time. An “audio filters” option allows users to select a 2D rectangular region-of-interest of a particular time and frequency range. This feature makes it easier to listen to a specific pattern seen in the spectrogram by removing distracting noise at other frequencies (e.g. to hear an animal vocalization drowned out by ship engine noise).

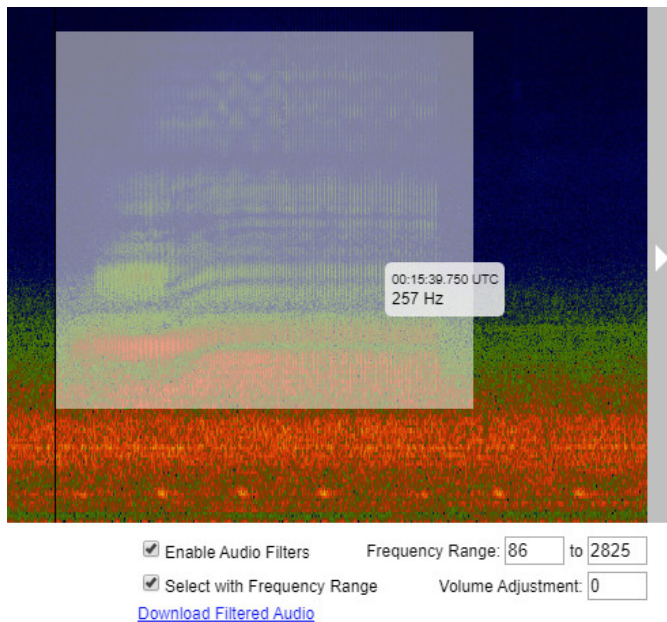


Fig. 17. Selecting a region of interest in both time and frequency. After a user makes such a selection, the frequency-filtered clip can be downloaded.

The mapping of frequency on the y-axis of the spectrograms is linear up to 80 Hz, at which point it begins to smoothly transition to a logarithmic scale. This significantly compresses the higher frequency ranges (> 1 kHz) while providing more emphasis on the content in lower frequency ranges (< 300 Hz).

The spectrogram images are all pre-processed. Individual wav files are read in and an FFT is performed. The FFT results are then turned into a spectrogram image using a custom color map. The color map used here was designed based on research into human visual perception issues, e.g. [1]. It was perceptually optimized to best reveal salient features (e.g. marine mammal sounds) and deemphasize low-level sound. As shown in Fig. 18, the color scale is also compatible with several types of color-blindness. It is mostly compatible with: red-weak protanomaly, green-weak deuteranomaly, blue-weak tritanomaly, blue-blind tritanopia, and blue cone monochromacy. However, it is not compatible with red-blind protanopia, green-blind deuteranopia, and monochromacy achromatopsia.




Fig. 18. Coblis [2] simulation of spectrogram color map appearance under (from left to right): normal color vision, red-weak protanomaly, green-weak deuteranomaly, blue-weak tritanomaly, blue-blind tritanopia, and blue cone monochromacy.

Image compression was an important consideration. The spectrogram images needed to be compressed as much as possible to save server disk space and ensure that users could download spectrogram images faster than they could scroll through them (to provide a seamless experience). The high-resolution details in spectrogram images can make efficient compression challenging, and preserving as much of this detail as possible necessitated a lossless compression algorithm.

Ordinary PNG image compression initially proved to be acceptable. However, with further experimentation it was discovered that the images could be compressed substantially smaller by taking advantage of various PNG compression settings. Ultimately, it was found that the open source software pngcrush [3] was the best solution, as it tries many different combinations of settings until it finds what works best for each individual image (as the optimal settings vary based on image content). Pngcrush also strips out all unnecessary metadata (such as color-correction data). Fig. 19 shows the experimental findings, with pngcrush providing the smallest file sizes.




Fig. 19. Average file sizes for spectrogram images compressed using a variety of common PNG compression algorithms/settings versus pngcrush, which tries many different combinations of settings, and produced the smallest file sizes.

To create the second-level, zoomed-out images, as .wav files are processed into full resolution (0.5 second) “lv11” spectrogram images, they are also processed into lower resolution (15 second) spectrograms using the correct FFT calculations (versus simply squishing the full resolution images). These very narrow (~20 pixel wide) spectrogram images are appended horizontally into a collage image, until the collage reaches an arbitrary width of 1800 pixels, at which point the collage is saved as a single “lv12” image. A “lv12” text file is also saved with the same filename as the image. This text file contains a list of the “lv11” source filenames and timestamps for each of the clips in the collage, the length of each clip in pixels, and the start and end index of the pixels. These allow for rapid lookup of timestamps and alignment of images in linked views when the images are clicked on.

To create the top-level, most zoomed-out images, the contents of each individual recording is collapsed (averaged) into a single column of pixels and added to a “lv13” collage. Similar to the “lv12” collages, these “lv13” images are saved out once they reach an arbitrary size, along with a “lv13” text file that contains, for each pixel column, the timestamp and source image in both the original level 1 data, and the level 2 collage. This approach was designed to be mostly agnostic to file contents and timestamps (beyond those contained in the filenames), while making integration with the web server simple, and minimizing interaction times on the website.

V. DEVIATIONS VIEWER

The final visualization tab contains a similar tri-level viewer. This interface presents the calculated deviations from normal sound levels around a lander for different frequency bins over time. Instead of showing sound pressure levels (as in the previous spectrograms), these plots use a diverging blue-white-red color scale to reveal times when the ocean around a lander

was unusually louder or quieter than normal at various frequencies. An example view of this interface is shown in Fig. 21.

The recordings were processed to calculate the mean sound pressure levels across 40 decade frequency bins every 60 seconds. A variable-length (week/month/year) moving window was then used to calculate the standard deviations of these mean sound pressure levels around each minute of each frequency bin. The mean sound pressure value for each bin was then compared to the moving window mean value, and the number of moving window standard deviations away from the moving window mean was calculated for each frequency and time.

This value (standard deviations away from mean) was then used with a diverging blue-white-red color map to determine the color value for each pixel in the plots. White indicates sound levels near the mean, while blues of increasing intensity indicate sound levels at multiple standard deviations below the mean (i.e. increasingly quieter than average periods), and reds of increasing intensity indicate sound levels which are multiple standard deviations above the mean (i.e. periods that are unusually loud).

These calculations were run three times, each using different moving window sizes: weekly, monthly, and quarterly. This generates three complete sets of images, each containing three levels of images and text based index files (using the same process as described in the previous section). Within the visualization interface, users can adjust the moving window size from weekly, monthly, or quarterly in order to suppress or highlight the influence of various factors (e.g. temperature or sensor drift). Fig. 22 shows the differences in plot appearance for the same data when viewed using different window sizes.




Fig. 20. The Deviations Viewer interface, showing time and frequency ranges when the ocean around the lander was unusually quieter than normal (blue) or unusually louder than normal (red) based on a moving window size of one week (with monthly and quarterly options selectable via drop-down box).

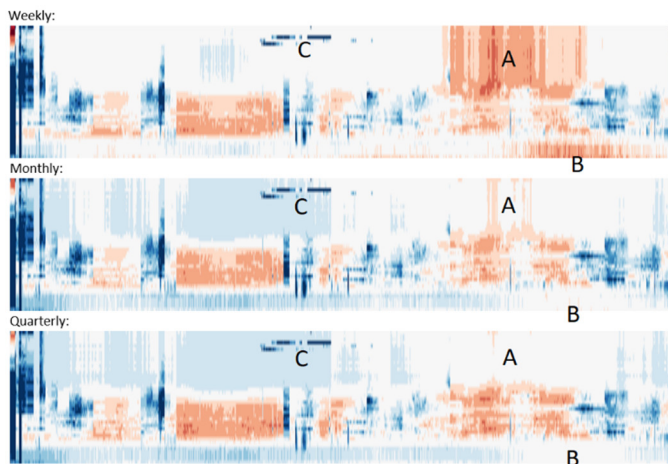


Fig. 21. Deviations from mean sound levels calculated using different moving window sizes. Blue is quieter, red is louder. There are regions of high frequency (A) and low frequency (B) sound that appear unusually loud when looking at weekly variation, but when considered over a longer time period, appear average. (C) shows an anomaly constrained to a narrow high frequency range, which remains unusually quiet across all moving window sizes.

VI. CONCLUSION

Providing public access to ADEON’s massive datasets was never going to be easy. The simplest way to satisfy the need to make the data public would be to offer all the files for download. However, this would mean that few would ever see or benefit from it: At over 70 terabytes, the data’s size poses a significant barrier for most researchers. Furthermore, specialized (and expensive) software is often needed to work with much of the data; and the overwhelming length of the recordings makes it challenging to find features of interest.

The suite of web-based visualization tools presented in this paper directly addresses these accessibility problems. Anyone with a web browser can now easily explore all of the various ADEON datasets without downloading enormous files or installing software. They can quickly skim through years of spectrograms, discover and listen to interesting features, visually analyze the patterns of marine mammals and other detected events, and better understand the sources of ocean noise.

Beyond enabling the public to explore the data, these tools make it easier for researchers, managers, and regulators to gain insight from the massive ADEON dataset, and can help inform future, more targeted studies into the impacts of marine noise.

ACKNOWLEDGMENT

Study concept, oversight, and funding were provided by the U.S. Department of the Interior, Bureau of Ocean Energy Management, Environmental Studies Program, Washington, DC under contract number M16PC00003, in partnership with other NOPP funding agencies.

The Center for Coastal and Ocean Mapping is supported by NOAA Grants #NA15NOS4000200 and #NA20NOS4000196.

REFERENCES

- [1] C. Ware, T. L. Turton, F. Samsel, R. Bujack, D. H. Rogers, “Evaluating the Perceptual Uniformity of Color Sequences for Feature Discrimination”, EuroVis Workshop on Reproducibility, Verification, and Validation in Visualization, 2017.
- [2] M. Wickline, Human-Computer Interaction Resource Network, “Coblis Color Blindness Simulator”, <https://www.color-blindness.com/coblis-color-blindness-simulator/>
- [3] G. Randers-Pehrson, pngcrush, <https://pmt.sourceforge.io/pngcrush/>