Exploring sonar data using clustered visual features and coordinated geovisualization

Asikur Rahman
Memorial University of Newfoundland
St. John's, NL, Canada
asikur.rahman@mun.ca

Orland Hoeber
Memorial University of Newfoundland
St. John's, NL, Canada
hoeber@mun.ca

Abstract

Exploring and analyzing oceanographic sonar data is a difficult task due to the extreme ratios in the dimensions of the data. While sonar data may consist of many hundreds of thousands of sonar pings coving hundreds of kilometres, the ocean depth of the data is at a much lower resolution. As a result, visual representations of the sonar data (echograms) are normally shown as long and narrow ribbons of data. As an analyst zooms in to show the echogram in sufficient detail, much of the contextual information is lost and the analyst must perform horizontal scrolling to explore the data. In this research, we outline an approach that couples a technique for visually clustering slices of the echogram based on visual similarity, with a geovisualization method that shows the spatial location of echogram slices on a virtual globe. Panning and zooming within each of these views of the data results in coordinated filtering, such that data outside of the viewport in one view is dimmed and de-emphasized in the other view. This approach provides data analysts with a powerful geovisual analytics tool for exploring sonar data. In particular, analysts may filter the data based on spatial regions of interest, visually identify important features within the data, and observe the spatial relationships among the locations of the echogram slices.
Keywords: geovisual analytics, sonar data, visual clustering, geovisualization, multiple coordinated views

1 Introduction

Sonar is commonly used to measure sub-sea phenomena in disciplines such as fisheries research and physical oceanography [5]. Often, such sonar datasets are collected over large geographic regions, and visualized using an echogram, where the length represents the number of sonar pings in the data set, and the height represents the depth of the data. Colour is used to show the strength of the sonar pings, representing the size or density of objects that are below the water surface.

A common use of such acoustic methods is to monitor and analyze fish stocks [6]. Vessels equipped with acoustic gear travel over some region of interest, collecting sonar datasets that may contain hundreds of thousands of sonar pings. Software systems such as Echoview [8] can then be used to process this sonar data and generate echograms. Fisheries scientists and environmental managers analyze and explore these echograms in order to understand the sub-sea environment [9].

The main challenge with analyzing sonar data using echograms is that the ratio of the length to the height can be very high (see Figure 1). The alternatives for viewing the data are to either view the entire echogram and not be able to see any detail, or zoom in so that detail can be seen, but then lose the contextual information provided by the full echogram. Furthermore, an echogram does not include any facilities for showing the geographic locations related to the data, resulting in additional cognitive load as the analysts attempt to keep track of the spatial locations of the features while they analyze the echogram.

To address this problem, we have developed an approach that combines the visual clustering of slices of the echogram with a geovisual representation of the spatial locations of each slice. Coordination between these views allows the analyst to dynamically filter the data based on the spatial extent and visual features within the echogram slices. The goal of this geovisual analytics approach is to support knowledge discovery activities through exploration and analysis of the data [2].

2 Method

2.1 Echogram Slice Extraction and Clustering

It is common for sonar dataset to be large, both in the number of sonar pings as well as the geographic distance covered. As a result, the corresponding echograms may be hundreds of thousands of pixels wide. By partitioning a high resolution echogram, we can produce a large number of low resolution echogram slices. For example, a 300000×1000 pixel echogram may be partitioned into a set of 300 1000×1000 pixel echogram slices. The problem then is how to visually organize and cluster these slices such that an analyst can identify features of interest. To further complicate this problem, the screen space even on a very high resolution monitor will not be sufficient to simultaneously show all the slices in detail when the source echogram is large.

Our solution is to cluster the echogram slices in a hierarchical manner based on their visual similarities, using a multi-resolution Self Organizing Map (SOM) [11]. While others have explored the use of SOMs within the context of geographic information systems [1], our approach is fundamentally different than these. Rather than clustering the raw data, we focus on clustering geographically continuous subsets of the sonar data, as represented by the echogram slices. This approach has been successfully employed in the context of image retrieval [12]. Since the task of visually identifying interesting features within the echogram is similar to image retrieval tasks (i.e., visually iden-

Figure 1: A sample echogram consisting of 30,000 sonar pings and a depth resolution of 1000 pixels. Note that this is 1/5 of the dataset used in the other examples in this paper.

tifying interesting images), we expect it to be well-suited to this problem domain.

In order to cluster the echogram slices based on similarity, we extract visual information from each slice in the form of a colour-gradient correlation feature vector [11]. Although there are other approaches to converting images into high-dimensional vectors, this approach has been found to be both efficient to calculate and provides good organizational performance [10]. As a result, rather than comparing the echogram slices to one another, such comparisons are made based on the Euclidean distance between feature vectors.

Fundamentally, a SOM is a type of artificial neural network that can be trained to organize and cluster data [7]. We use the SOM to map the high-dimensional visual feature vectors extracted from the echogram slices into a 2D grid. The topology-preserving property of the SOM ensures that similar feature vectors are mapped to cells in the SOM that are near one another. As a result, placing the corresponding echogram slices at their assigned locations within the 2D space produces a clustering of the slices based on their visual features.

However, when the collection contains a large number of echogram slices, it is impractical to display all of the slices at the same time. The solution to this problem is to use a multi-resolution SOM. At the highest resolution, all of the echogram slices are organized and visually clustered. This space is progressively divided in half over both the x and y dimensions, producing lower resolution spaces. For a given cell in the lower resolution space, the average feature vector is calculated from the four cells in the higher resolution space, and the nearest corresponding echogram slice is used as the representative for this region. The end result is a hierarchical clustering of the echogram slices.

This interface supports pan and zoom operations, whereby the analyst can zoom into a region of space that appears to contain some visual feature of interest in the echogram slices. Once sufficient space has been created as a result of this zoom operation, a higher resolution SOM is chosen and additional echogram slices are shown. Zooming further and further pushes those echogram slices that are distant from the focal point of the zoom operation out of the viewport, and loads additional echogram slices that are similar to those near the focal point.

2.2 Geovisualization

While the visual clustering described in the previous section can allow an analyst to easily identify interesting features within the echogram slices, what is lost is the continuity of the echogram. To address this, and to further enhance the analysts' understanding of the spatial aspects of the data, a geovisual representation is provided to show the locations of each of the echogram slices on a map.

With the geographic space, directional glyphs are used to represent the location of each echogram slice as well as the direction in which the source echogram was measured. In order to disambiguate these locations within congested areas of the map, paths are drawn between successive echogram slice locations using cubic Hermite splines.

Like the visual space, this geographic space also supports pan and zoom operations, making it easy for analysts to focus on specific geographic locations by zooming into these regions.

2.3 Coordinated Interaction

To further support the exploration of the data, the visual space and geographic space are linked to support coordinated interaction [3]. In the visual space, the echogram slices are organized such that those with visual similarities are grouped together; simultaneously, the geographic space shows their locations. Panning and zooming within each of these views of the data results in coordinated filtering, such that data outside of the viewport in one view is dimmed and de-emphasized in the other view. Figures 2 - 5 provide an example of this process, whereby an analyst first zooms into a geographic region of interest, then zooms into a visual region of interest, and finally zooms in to increase the sizes of the echogram slices such that specific features can be examined in detail.

This coordinated interaction helps analysts to explore the echogram data based on both the features of the echogram slices and their geographic locations, removing visual complexity from both views along the way. This allows the analysts to more readily identify and explore the underlying patterns and features as a result of the de-emphasis of the data that was pushed out of the viewport of the opposite view.

Rather than providing persistent visual links between each echogram slice in the visual space and each glyph in the geographic space, an interactive highlighting mechanism is provided. By selecting a specific echogram slice (or glyph), the system automatically highlights the corresponding object in the other view. Furthermore, when the echogram slice is selected, it is temporarily zoomed to fill the visual space, allowing the analyst to examine it in greater detail.

Two features are provided within this system to mitigate the risk of slicing an echogram through specific features of interest. The first of these is the ability for analysts to merge multiple slices into a larger subset of the echogram. This is performed by holding down the shift key and selecting the beginning and end point within the geographic space. The merged slices are displayed within the visual space.

The second of these features is the ability for the analysts to control the width of the echogram slices. By making the echogram slices narrower, more will be generated and the quality of visual clustering will be increased; however the likelihood of dividing an important feature between multiple slices will be increased. By making the echogram slices wider, the opposite tradeoff occurs. Whether wider or narrower echogram slices are appropriate depends on the features of the phenomenon that is being investigated within the sonar data.

Figure 2: Initially, the echogram slices are clustered in the visual space (left), and their locations are shown in the geographic space (right).

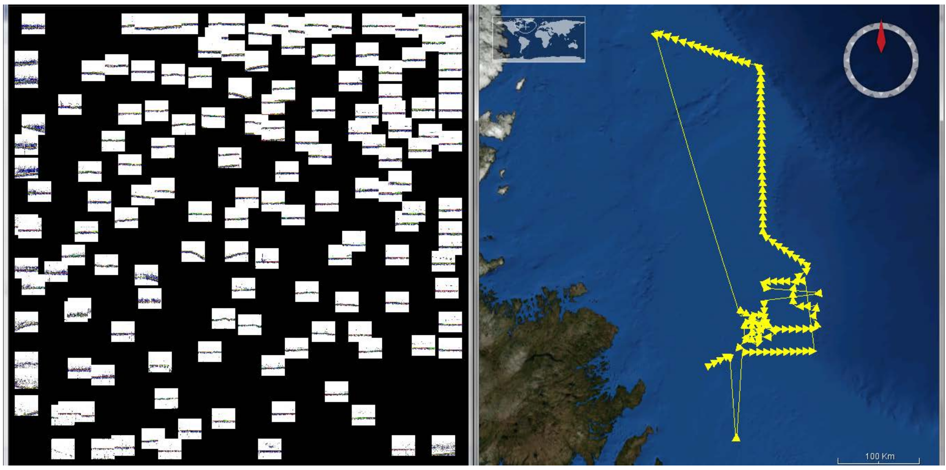

Figure 3: As the analyst zooms in within the geographic space (right), the corresponding echogram slices that are located outside of the viewport are dimmed in the visual space (left).

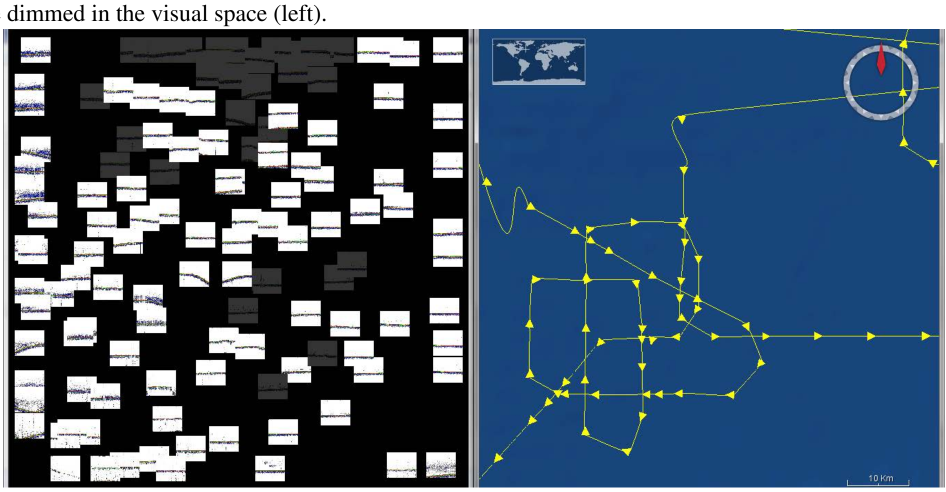

Figure 4: Zooming into a region of interest in the visual space (left) results in the locations of the echogram slices that are outside of the viewport to be dimmed in the geographic space (right).

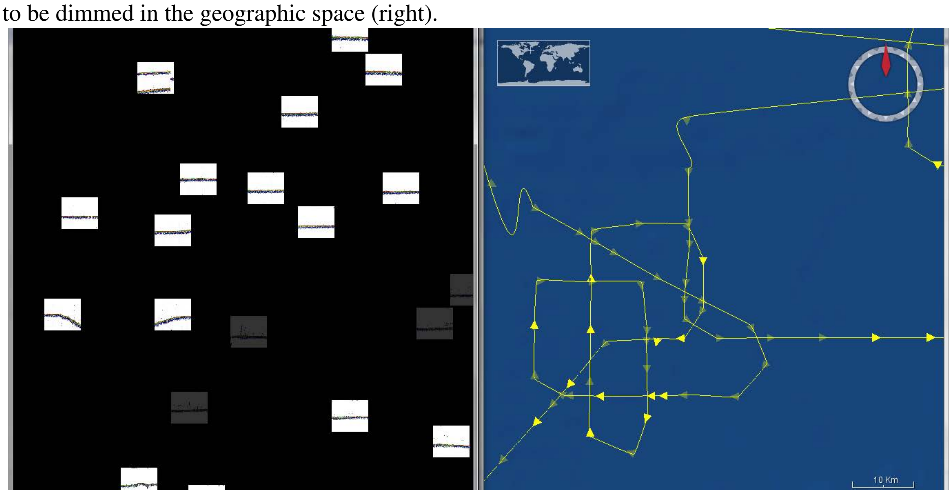

Figure 5: Continuing to zoom in once all of the hidden echogram slices are shown increases the sizes of the echogram slices, allowing the analyst to examine the features in detail.

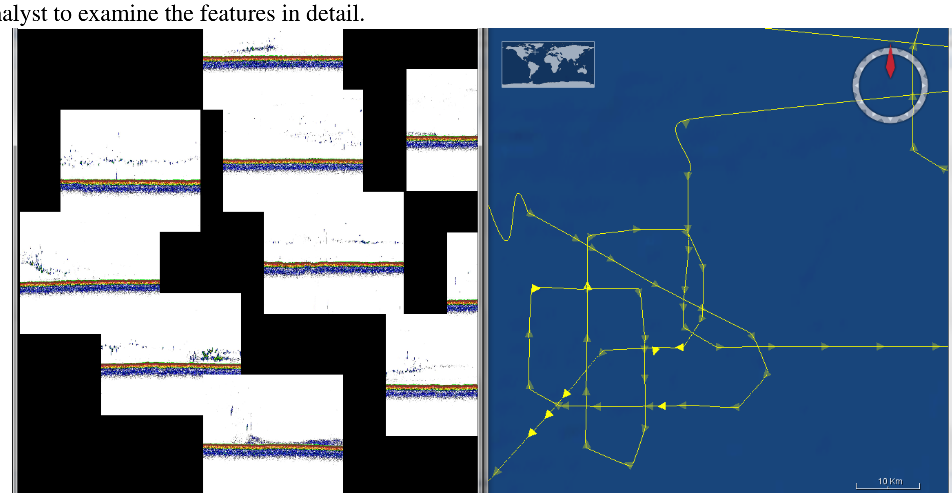

3 Conclusion

Analyzing sonar data using echograms is difficult due to the extreme ratio of the length to the height, and the lack of corresponding geographic representations. To address these problems, we have developed a novel geovisual analytics approach that slices the echogram into smaller pieces and uses coordinated interaction between two views of the data: a hierarchical visual clustering of the echogram slices and a geovisualization of their locations. This approach supports both geographic-based exploration while providing visual feature information, and visual feature-based exploration while providing geographic information.

Field trials that will focus on studying the value and benefits of this approach to analyzing sonar data are currently in the planning stages. Other future work includes analyzing the differences between different feature vector creation methods within the context of visually clustering echogram slices, and studying the benefits and drawbacks of different alternatives to using self-organizing maps, such as multidimensional scaling [4].

References

[1] P. Agarwal and A. Skupin. *Self-Organizing Maps: Applications in Geographic Information Science*. Wiley-Blackwell, 2008.

[2] G. Andrienko, N. Andrienko, P. Jankowski, D. Keim, M. J. Kraak, A. MacEachren, and S. Wrobel. Geovisual analytics for spatial decision support: Setting the research agenda. *International Journal of Geographical Information Science*, 21(8):839–857, 2007.

[3] W. Baldonado, Q. Michelle, A. Woodru, and A. Kuchinsky. Guidelines for using multiple viewsin information visualization. In *Proceedings of the Working Conference on Advanced Visual Interfaces*, pages 110–119, 2000.

[4] T. F. Cox and M. A. A. Cox. *Multidimensional Scaling*. Chapman and Hall/CRC, 2nd edition, 2000.

[5] J. P. Fish and H. A. Carr. *A Guide to the Generation and Interpretation of Side Scan Sonar Data*. Lower Cape Publishing Company, London, 1990.

[6] P. Freon and O. A. Misund. *Dynamics of Pelagic Fish Distribution and Behaviour: Effects on Fisheries and Stock Assessment*. Fishing News Books, Oxford, 1999.

[7] T. Kohonen. *Self-Organization Maps*. Springer-Verlag, Berlin, 1995.

[8] Myriax Software. Powerful and Flexible Software for Echosounder and Sonar Data Processing. `http://www.echoview.com`, May 2011.

[9] G. A. Rose and W. C. Leggett. Use of oceanographic forecasts and echosounders to enhance an inshore gillnet fishery for Atlantic cod (Gadus morhua). *Canadian Journal of Fisheries and Aquatic Sciences*, 50(10):2129–2136, 1993.

[10] G. Strong and M. Gong. Organizing and browsing photos using different feature vectors and their evaluations. In *Proceedings of the International Conference on Image and Video Retrieval*, pages 1–8, 2009.

[11] G. Strong and M. Gong. Similarity-based image organization and browsing using multi-resolution self-organizing map. *Journal of Image and Vision Computing*, 29(11):774–786, 2011.

[12] G. Strong, O. Hoeber, and M. Gong. Visual image browsing and exploration (Vibe): User evaluations of image search tasks. In *Proceedings of the International Conference on Active Media Technology*, pages 424–435, 2010.