Multi-dimensional Visual Representations for Underwater Environmental Uncertainty

Greg S. Schmidt[1] Sue-Ling Chen[1] Aaron N. Bryden[1] Mark A. Livingston[1] Bryan R. Osborn[2]
Lawrence J. Rosenblum[1]

[1]Virtual Reality Laboratory, Naval Research Laboratory, Washington, D.C.
[2]Metron, Inc., Reston, VA

1 Introduction

As portrayed in Tom Clancy's novel "The Hunt for Red October" (and the subsequent movie), precise localization and tracking of submarines in the ocean is difficult. The difficulty lies in the inherent environmental uncertainties that affect acoustic propagation in the ocean. In active acoustic detection a source device (sonar) transmits an acoustic signal into the surrounding environment, and the acoustic returns obtained from a set of receivers are processed to extract echoes from underwater objects. These echoes provide estimates of the position and, over time, the heading of a target.

However, numerous factors limit the accurate reconstruction of the target's location and heading. Uncertain temperature-salinity depth profile information results in inaccurate prediction of the curvature of the acoustic rays. Uncertain bathymetry (knowledge of the ocean bottom) includes misregistered maps, data gaps in maps, temporal changes to the ocean bottom, and limited knowledge of the mineral content of the ocean bottom. These factors again limit the accuracy of position estimations based on the acoustic returns. Finally, local oceanic effects (e.g., internal waves, thermal currents, fine-structure) also impact the acoustic return and thus add to the uncertainty surrounding the target's position. The resulting uncertainty has four dimensions: latitude, longitude, depth, and signal strength. The sea bottom data contains multi-variate physiographic information (region data associated with soil types) and depth measurements that are subject to multi-dimensional uncertainty. The target state prediction process (i.e. determining the position, heading, etc. of a target) involves data that spans multiple dimensions caused by the uncertainties discussed above.

We investigate how to represent the resulting multi-variate information and multi-dimensional uncertainty by developing and applying candidate visual techniques. Although good techniques exist for visualizing many data types, less progress has been made on how to display uncertainty and multi-variate information, and this is especially true as the dimensionality rises. At this time, our primary focus is to (1) develop the statistical characterizations for the environmental uncertainty (described only briefly in this paper) and (2) develop a visual method for each characterization. The mariner community needs enhanced characterizations of environmental uncertainty now, but the accuracy of the characterizations is still not sufficient enough and therefore formal user evaluations cannot take place at this point in development. We received feedback on the applicability of our techniques from domain experts. We used this in conjunction with previous results to compile a set of development guidelines (some obvious, others not).

We first review a select number of applicable visualization techniques. We then describe the investigations for representing the bathymetric information in Section 3 and the target state estimation uncertainty in Section 4. Section 5 describes briefly the display-system architecture we developed to demonstrate the visual candidates. Finally, in Section 6 we summarize the results and compile a list of guidelines for developing visualization techniques for multi-dimensional uncertainty and multi-variate information.

2 Review of Applicable Visualization Techniques

Pseudo-coloring is a popular approach for environmental [3] and meteorological [2] data in which colors are assigned based on the type of information and uncertainty. Much visualization for geographic data rely on glyphs (e.g., isobars, tufts). It is natural to map the display parameters of the glyphs to the information associated with the data. Parameters that have been demonstrated include color, thickness, opacity, and scale [6].

Multi-sensory modalities such as haptics and sound [5] have been utilized when the visual approaches have been exhausted—typically the case for multi-dimensional data. Modifications to geometry [6] such as surface displacements, fat surfaces, and bumps show uncertainty about interpolated surface locations. Animated surfaces [1, 6] can show the range of locations over which a surface might lie as well as the range for other data types.

For a more complete review see [4, 6].

3 Representing Sea Bottom Information

Our research investigated representations of depth measurements containing multiple dimensions of uncertainty and

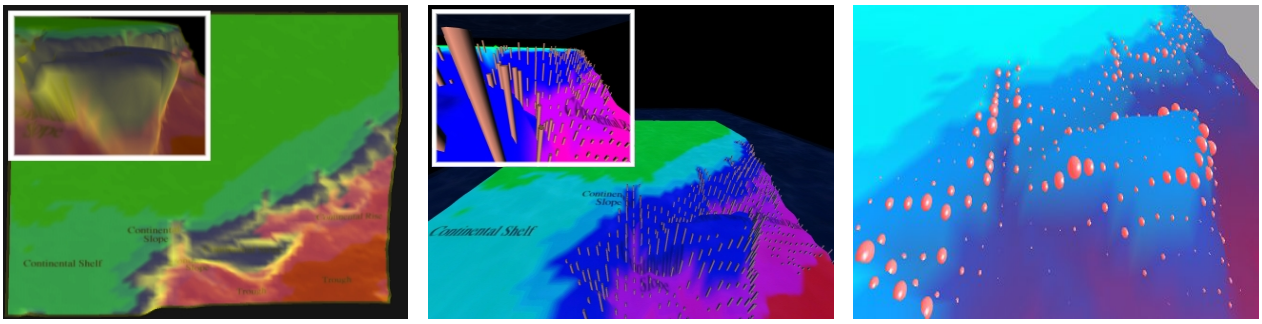

Figure 1: Images show from left to right a combination of physiographic region data and 1D depth uncertainty using the following visual techniques: (a) color, (b) cylinder glyphs and (c) sphere glyphs. In (a) shades of yellow are used for the uncertainty and the remaining colors for the physiographics regions, which are also labeled and a key is available not shown in the image.

how to combine them with the physiographic region data associated with soil types. We then explored adding additional feature information to the uncertainty representation. This section presents results for the 1D, 3D, and 4D cases for uncertainty and also shows how to represent multi-dimensional uncertainty for multi-variate information.

3.1 One-Dimensional Uncertainty

We began with a sea-bottom dataset (described in Sidebar 1) provided by NRL-Stennis. One of our goals was to develop visual techniques that combine the bathymetric data with the uncertainty in a manner that allows cross-correlation of the uncertainty with the proper locations of geographical content. Since the bathymetric dataset ranges over several physiographic regions, each with differing depths and soil types, we used a coloring and text-labeling scheme to distinguish the different physiographic regions. We generated a surface from the depth data and applied the texture composed of the colored regional data. We scaled the depths differently than the latitude and longitude since the range of depths (100 m to 2000 m) is quite small compared to the range of latitude and longitude (each spans over 200 km). Three visual techniques were developed for representing the 1D uncertainty in the depth. Uncertainty value thresholds assured that no representations were shown at sample points where the uncertainty was less than a specified threshold value.

The first method applies a range of intensities of one color to a sample point based on the degree of uncertainty at that point. To avoid confusion, we selected colors that do not conflict with the colors used for the phyiographic regions. The intensity of the color is mapped proportionally to the degree of uncertainty. We use a picture-in-window approach to show a wide-area and close-up view of the technique in Figure 1a, using yellow for the uncertainty.

The second approach uses cylinder glyphs placed at the corresponding sample point, where the height of the cylinder corresponds to the degree of uncertainty and is aligned in the dimension associated with the depth data. We show the technique applied to the sea bottom data in Figure 1b. The cylinder works well for distinguishing the degree of uncertainty when observed from the side. However, it is difficult to interpret the glyphs (and associated uncertainty) when observed from above. The choice of shape for the glyphs is an important decision to be made for the design of a technique.

The third technique we developed avoided viewpoint dependencies. We determined that sphere glyphs (Figure 1c) would work well for the 1D uncertainty since the sphere allows observations to be performed from all viewpoints. A drawback to this approach is that the dimension in which the uncertainty lies can not be identified from observing the sphere glyph. This may produce conflicting perceptions to what the glyph refers. However, our domain experts preferred the sphere glyph over the others because it allowed them to observe the 1D uncertainty from all viewpoints.

3.2 Three-Dimensional Uncertainty

To develop visualization techniques for the three-dimensional case consisting of uncertainty in latitude, longitude and depth, we followed three goals. The first goal is to provide cross-correlation between the uncertainty information and the underlying geographic data. The second goal is to reduce view dependence in techniques when possible. The third goal takes into consideration the domain knowledge of the task to determine if wide-area or local-scale observations (or both) are required. Thus, if a user wants to gain an understanding where the most significant uncertainty lies, visual techniques that work best for wide-area observations should be used. If the user wants to learn the value of an attribute or differentiate between two, local-scale visualization techniques should be used.

We generated a dataset containing uncertainty of the 3D surface shape by extending the 1D, directed uncertainty in the bathymetric dataset. That is, by applying the 1D uncertainty in the bathymetry data, we can generate the uncertainty in the 3D surface from which the bathymetric measurements were taken. The latitude and longitude uncertainty estimates were generated by keeping consistent the essential features (scale and distribution) of the true uncer-

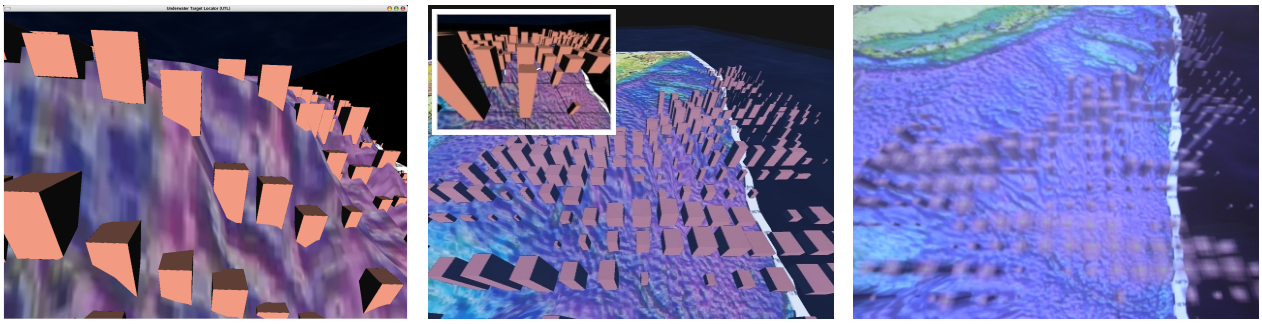

Figure 2: Images show three approaches using box glyphs for 3D uncertainty. (a) Left shows boxes placed at surface sample points. (b) Middle shows regularly-spaced boxes placed on plane above surface. (c) Right shows snapshot of animation approach with vibrating boxes based on the degree of uncertainty in latitude and longitude.

tainty in the sea bottom measurements. For the first investigations, we consider the case where the uncertainty is symmetric in latitude and longitude directions about sample points lying on a regular uniform grid. A later investigation considers the true situation where the centers for the uncertainty estimates do not lie on a regular grid.

We used a glyph technique that maps depth, latitude and longitude uncertainty to the dimensions of a box. We applied the box-glyph approach to the symmetric 3D-uncertainty dataset. The box glyphs are generated by centering a unit-size cube at each sample point and scaling it by a factor proportional to the uncertainty in each corresponding dimension. The scale of each box glyph is adjusted consistently across all dimensions to prevent overlaps with neighboring boxes while conserving the proper correlation between the error in each dimension. Figure 2a shows a side-to-top view of the box glyphs as they are applied to the dataset. We replaced the color-coded regional map of the physiographic regions with a low-resolution bottom texture of the sea-bottom terrain.

We were not able to satisfy the view-independence goal with the box-glyph approach, but were able to satisfy the other goals. Observations of the depth uncertainty can be made from the side, and the latitude and longitude uncertainty from the top. Wide-area observations give a general idea where the uncertainty is most predominant, but local-scale observations are hindered by occlusions between the glyphs and the surface. The occlusions occur most frequently where there are sloped surfaces, making it hard to distinguish variations in the uncertainty from the side and top views. Besides the occlusions, it is also difficult to distinguish depth variation when there is no common visual reference to compare heights.

To reduce the limitations we revised our approach by placing box glyphs on a planar surface elevated above the highest point in the bathymetric surface as shown in Figure 2b. Elevating the glyphs reduced the visual correlation between the uncertainty and data unless it is viewed using an orthographic projection from above looking down. This lack of correlation is most obvious from the side views, but the tradeoff is a significantly improved distinction of the heights (and correspondingly the depth uncertainty) of the glyphs. Having the glyphs placed next to each other on the same plane also improves the distinction of the latitude and longitude uncertainty variation when viewed locally.

A third method was tried that used animation to represent the uncertainty. This method maps a range of motions to the glyphs based on the degree of uncertainty. A large range of motion corresponds to high uncertainty, and a small range to low uncertainty. We applied the approach for the latitude and longitude uncertainty by translating the box glyph randomly from the sample-point center proportionally to the uncertainty in each dimension for each frame. We attempt to demonstrate the method by showing a photo of the screen with the animation snapped by a camera with a long exposure time in Figure 2c. We found that the areas where the largest range of motion occur attracts the user's attention the most. The technique works very well for differentiating regions containing different degrees of uncertainty in latitude and longitude.

The planar approach was most effective of the three. We further improved the planar technique by adding a 2D grid to the plane that provides correspondence with the sample-point centers of the measured data. A user is now able to easily differentiate the degree of uncertainty in alike dimensions (latitude and longitude) between sample points by using the grid lines as a reference. We show a top-view of the improved approach in Figure 3a.

Next we addressed the case where the uncertainty values lie on an irregular grid. We applied the planar approach used for the regular-grid data to the irregularly distributed dataset. We modified the technique by shifting the box glyphs according to the "actual" center (the midpoints) of the range of uncertainty in latitude and longitude. We indicate the center of each data element by drawing a cross-hair marking on the tops of the boxes. The approach is shown in Figure 3b. Notice the change in position of the boxes with respect to a view of the same range-difference symmetric data shown in

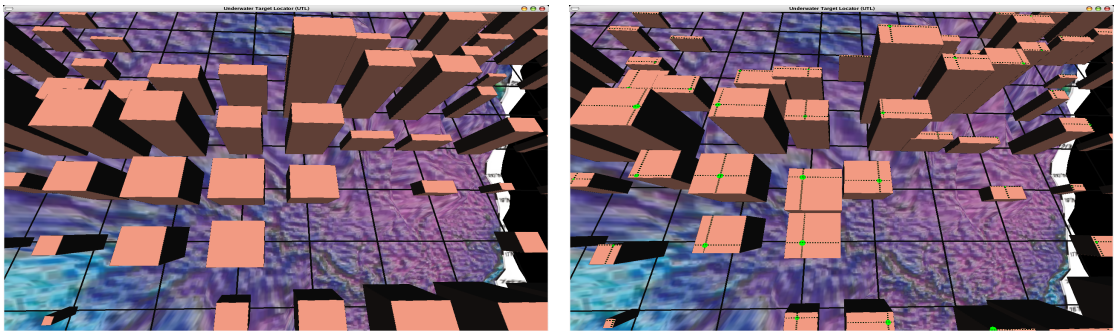

Figure 3: Box-glyph approach applied for: (a) symmetric latitude and longitude uncertainty and (b) non-symmetric uncertainty. Boxes are shifted in the latter case. Cross-hair markings indicate sample-point centers.

Figure 3a.

This method still works well for visualizing the variation of the uncertainty from a wide-area viewpoint. On a local scale, the secondary visual cue (cross-hair markings) effectively indicates how much the data is shifted in both dimensions. A minor drawback is the data must be down-scaled in order to avoid having overlapping boxes.

3.3 Four-Dimensional Uncertainty

The four-dimensional uncertainty was visualized by treating the 3D measurement uncertainty as a function of time. Datasets are sometimes compiled from multiple scans of the bottom, which may take place several years apart. In some cases, the bottom may have shifted, especially in areas with high vertical gradients. The use of different sonar scans can also produce multiple datasets for the same location. We treat this data as being four-dimensional to visualize the associated uncertainty.

We first consider the case when multiple non-overlapping datasets (in latitude and longitude) are connected together to make a larger contiguous dataset. The dataset is four-dimensional since the time elements vary between separate pieces. We created a 4D dataset by splitting the symmetric 3D dataset into four groups and assigning a unique time element to each. We developed different-shaped 3D glyphs (each has the scaling properties of the box glyphs used in previous representations) for each dataset group. Figure 4a shows the following generated shapes: pyramid, upside-down pyramid, wedge and modified box. This technique is effective locally, since the properties of the glyphs are easy to distinguish. However, there can be confusion in determining the regions of highest uncertainty for a wide-area view, since the uncertainties associated with each glyph type may have different scale distributions.

When the elements from separate datasets overlap in the 3D space, the visualization task is more challenging. Some rules for reconciling the overlapping information have been suggested by Smith et al. [7] (e.g., a node with lower uncertainty supersedes a node with greater uncertainty, a newer node supersedes and older node (particularly when the old data is known to be inadequate) and a shoaler node supersedes a deeper node). The data could also be combined by averaging. We did not fully address this case, but suggest the following two visualization strategies be applied. Show the combined (or selected) 3D component of the data using a primary visual cue. Then show the elements of the data that were lost and/or parameters of the dataset that were used in the combining operation with secondary visual cues.

3.4 High Degree of Multi-variate Information

Here we address the case when the degree of information to be presented exceeds the number of effective visual approaches that can be used to represent the information. When combining the four dimensions of uncertainty for the sea-bottom measurements with the information about the physiographic regions, landmark features, map information and more, the quantity of information to be presented may exceed the number of effective visualization channels for representing the information. This opens the door for exploring non-visual approaches (e.g., touch and sound) and combinations of the senses (multi-modal visualization). Combining senses to convey information may be superior to what could be gained with only one modality. The reason being that one modality may reinforce the other, or a user may be able to effectively fuse different information coming from different modalities. There has been little research to date on multi-modal "visualization" of uncertainty.

To gain an understanding of how effective these approaches are, we investigated haptics and combinations of haptic and visual approaches. Haptic approaches utilize the user's sense of touch. We utilized Immersion's CyberForce/Grasp system[1] shown in the lower-left corner of Figure 4b to provide force feedback to the user's hand. The device can be programmed to map the interactions of an on-screen virtual hand to a user-worn exo-skeleton device (CyberGrasp), which applies forces to the user's fingers. The user then has the perception that he/she is feeling the virtual object.

We began by combining haptic input with the 4D glyphs

[1] Immersion, www.immersion.com/3d/products/cyber_force.php

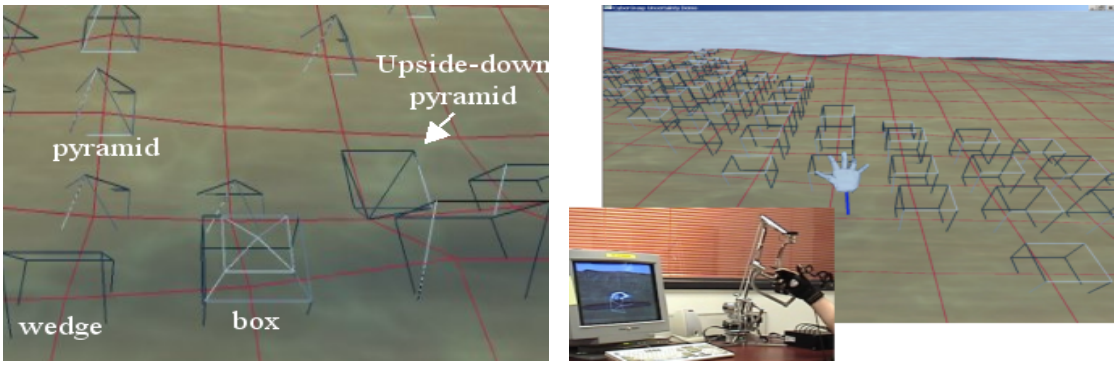

Figure 4: Methods for datasets of four or more dimensions. (a) Left shows multi-shape glyphs with characteristics for three dimensions of information. (b) Right shows a haptic device being used to perceive information about a dataset, which can be applied for one or more dimensions of information.

from the previous investigation. We mapped one of the dimensions of the 4D uncertainty to the haptic channel. Specifically, we mapped the depth uncertainty to the degree of stiffness of the glyphs. Large values of uncertainty then produced greater force-feedback to the user's fingers. The box glyphs corresponding to the very low uncertainty values were thresholded so they cannot be felt. This mapping of uncertainty to stiffness seems inverted (i.e. should it not be the case that certain objects are stiffer?), but our goal is to use the haptics only when necessary and in this case to highlight where the features have the most uncertainty. A dataset is haptically observed on a general scale using an arm sweeping motion. A perfect dataset should allow the user to perform complete sweeps through the dataset without interference. We only want to indicate where the degrees of uncertainty are highest (infrequent numbers of locations).

We utilized the stiffness characteristics by developing a set of interaction methods that allow the user to observe the data on both a local and wide-area scale. The forces applied to the fingers by the CyberGrasp provide the ability to inspect the surface structure of objects locally. The forces applied at the wrist by the CyberForce when combined with an arm-sweeping motion allow the user to gain a rough understanding as to where the most prominent uncertainty lies in the dataset. We estimated from our investigations that three to five different ranges of uncertainty could be differentiated with the arm-sweeping approach, which is less than with visual methods.

On a local scale, the user can distinguish two objects by inspecting each individually with the fingers. The fidelity of the haptics equipment is critical for the differentiation. From what we have observed with the CyberGrasp and the shapes we used, there was only a small range of distinguishable forces. We could differentiate between the glyphs with pointed tops (e.g., pyramid and wedge) from the glyphs with flat tops (e.g., box and upside-down pyramid), but we could not easily differentiate the pyramid and the wedge. Overall, our investigations indicate that the use of haptics has promise for conveying information about data, but in order for haptics to be used effectively there needs to be improvements in the fidelity and usability of the hardware.

4 Representing Target State Estimation Uncertainty

One method for estimating the target state (e.g., position and velocity) of submarines is the use of multi-static active acoustic sensors. A number of buoys are distributed over an ocean area. The charges on the source buoys are detonated sequentially, one every few minutes. Between detonations, the hydrophones in the receiver buoys listen for echoes of the shockwave as it scatters off objects, potentially including a target submarine. The time between the reception of the direct blast and an echo produces an ellipse of possible locations for the echo producer. By accumulating a number of these echoes from the target over time, it is possible to identify a distribution for the position and heading of the target.

Present systems include signal processing algorithms that have been developed to automatically process the time-series at the receiver hydrophones and identify detections. These algorithms produce a set of time values for each hydrophone where the signal or matched-filter output exceeds some threshold. Each of these "detections" comes from one of three things: random stochastic fluctuations in the noise signal (false alarms), clutter echo, or a target echo.

In order to use these automatic detections to estimate target state, we adapted the Likelihood Ratio Tracking (LRT) algorithm [8]. The LRT algorithm uses a recursive Bayesian framework to update prior target state information with newly received sensor information. In the standard framework, the algorithm begins with a prior state over a gridded state space with position and velocity as dimensions. We extended the state space to include an environmental dimension that accounts for uncertainty in the environmental predictions. As sensors report information, their information content is incorporated into the posterior likelihood ratio over the state space by the use of a measurement likelihood

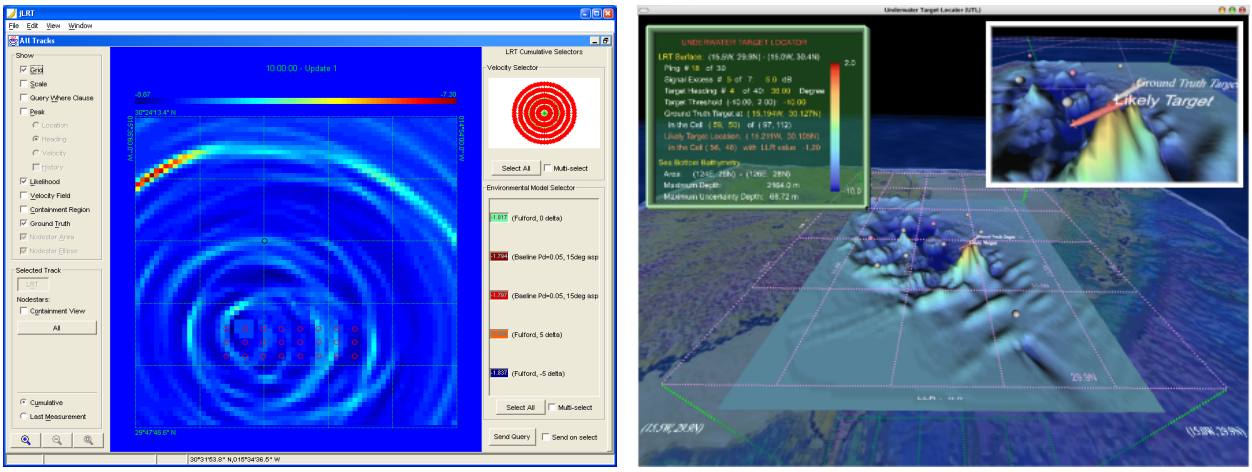

Figure 5: (a) Left shows the LRT display system with one frame from a likelihood image sequence. Multiple colors are used to represent the range of likelihood values. (b) Right shows the 3D desktop software demonstating the 3D color-height map approach. A red submarine is placed and oriented at the position of predicted maximum likelihood, and a white submarine at the ground truth position and heading. (See inset at upper right.) The ground truth was recorded during the capture of acoustic echo data.

ratio function. Between sensor measurements, the likelihood ratio function evolves according to a motion model prescribed for the target.

We developed visual representations for the extended Bayesian state space which includes a dimension of environmental uncertainty parameterized by mean signal excess prediction error. The state space for the likelihood ratio function $\mathcal{L}(x, y, t, theta, SE)$ is parameterized by five dimensions where x and y are the latitude and longitude spatial components, θ is the target heading, SE is the mean signal excess prediction error and t is time. We first focus on what is primarily meaningful to the end-user: the positions (x, y) and headings θ corresponding to the maximum likelihood or maximum or mean aggregations of likelihood values over the state space at a given time t. To present this information we first consider a projection of the state space that utilizes no more than the maximum number of display dimensions that can be shown at once. We include time with the standard three Euclidean spatial display dimensions since the dataset is generated as a function of time and therefore can be presented as an animation over the time dimension. Then we investigate higher-dimensional subsets of the state space which map to more than the four display dimensions (i.e. three spatial dimension and time denoted by (X, Y, Z, T)) and require sophisticated visual techniques to present the information. We divide our investigations into sub-space projections of the state-space dimension variables: 3D, 4D and 5D.

4.1 Three-Dimensional State Data

We start by developing visual techniques for the 3D subsets of the state space that include the position (x, y) and time t dimensions. When we consider mappings of the state space to display variables, the values of the likelihood function $\mathcal{L}$ are included. Therefore, the display mappings will entail sub-space projections on the 6D space $(x, y, t, \theta, SE, \mathcal{L})$. A single projection of the display space of the form $(subspace\ vars) = proj_{subspace\ vars}(x, y, t, \theta, SE, \mathcal{L})$ we denote using the notation $g(subspace\ vars) = proj_{subspace\ vars}(\mathcal{L}(x, y, t, \theta, SE))$. Notice we include the range of the function encoded in the function notation for each single projection. The set of sub-space projections encompassing the remaining variables not used in the projection is denoted using the function notation $f(remaining\ vars)$. For example, the 3D sub-space projections of the state space (or 4D sub-space projections of the display space) that we address here on position (x, y) and time t (for display space representation include $\mathcal{L}$) are $f(\theta, SE) = \{proj_{x,y,t}(\mathcal{L}(x, y, t, \theta, SE))\}$, where θ and SE are typically in the ranges $[0, 360]$ degrees and $[-20\ dB, +20\ dB]$, respectively.

To represent the space of data $f(\theta, SE)$ a specific value for SE and θ can be specified by the user interface or aggregation algorithms can be performed across the two dimensions. Figure 5a shows a frame from the LRT engineering display system that presents 2D animations of the spaces $f(\theta, SE)$ for specific θ and SE. The image is color-coded based on a logarithmic scale of the likelihood values $\mathcal{L}$ at each position (x, y). Log likelihood ratio (LLR) values are greater than zero. The likelihood values are mapped in descending order using a "ROYGBIV" color scale which maps the highest values to red and lowest ($LLR < 0$) to violet. The warmest colors (reds and oranges) represent the areas where it is likeliest that a target is located.

The LRT engineering display system uses two of the three Euclidean display dimensions and time (X, Y, T). The an-

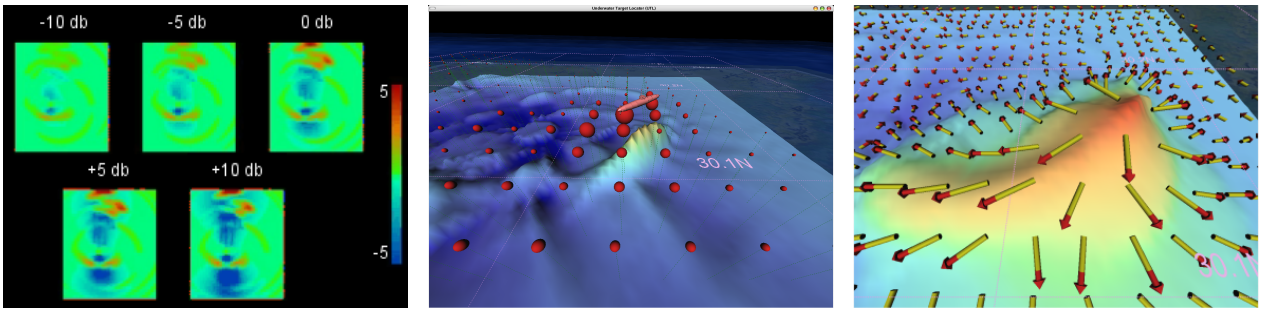

Figure 6: First visual approaches for 4D sub-space projections of 5D dataset. (a) Left shows frame of 2D-image sequences placed side-by-side. (b) Middle shows mean likelihood across SE using color-height map and variance using varying-size sphere glyphs. (c) Right shows mean likelihood across heading using color-height map, maximum likelihood across heading and heading using arrow length and direction, and difference of maximum and mean using red color.

imations show regions of red and orange form and move across the latitude and longitude dimensions over time. This indicates a lock on a target submarine and shows the heading that the target takes over time. What cannot be inferred very precisely are the values of the likelihood ratios that the colors represent. It is sometimes hard to distinguish if one shade of a color represents a higher likelihood value than another even with a color scale co-located on the interface. It seemed logical to try a mapping of the likelihood values to a geometric scale. So we utilized the Euclidean Z dimension and mapped the log likelihood values to surface height in addition to the colors. The 3D LRT system in Figure 5b shows the new mapping. Note that Figure 5b's view of the dataset does not correspond with the view in Figure 5a.

The height mapping redundantly shows the same property of the dataset as the coloring technique, but we feel that judgements of the data's values can be made more precisely with the geometric encoding. Additionally, the user can draw on a second cue other than warm colors, that being surface peaks, to identify where the likeliest targets are positioned. The user can make clearer distinctions about the likeliest location by observing the data from a side view. One drawback is the surface mapping does have the potential to be misconstrued with representations used to show geographic data. However, the mariner community is warming up to a system called IMAT[2] which combines 3D geometric surfaces for acoustic data with 3D bathymetric surfaces. We believe that we will be able to adapt the color and/or texturing schemes in order to reduce any ambiguities that the geometric structure may produce. We applied the color-height mapping as the primary visual cue for the higher dimensional sub-space projections of the 5D dataset.

We also provide iconographic and text cues to indicate the location and heading of the likeliest targets. We place the text "likely target" and a red submarine oriented with respect to the heading θ at the highest peaks. The system also provides data aggregation in θ, in SE and in the 2D space (θ, SE). In each case, the maximum likelihood values (highest threat) are computed and mapped to the surface and color in the same manner as before. The values for θ and/or SE are recorded for the maximum likelihood values and used for the visualization (e.g., the submarine's orientation is determined from the angle θ at maximum likelihood).

4.2 Four-Dimensional State Data

The set of 4D sub-space projections of the state space are: (a) $f(\theta) = \{proj_{x,y,t,SE}(\mathcal{L}(x,y,t,\theta,SE))\}$ and (b) $f(SE) = \{proj_{x,y,t,\theta}(\mathcal{L}(x,y,t,\theta,SE))\}$. The sets $f(\theta)$ and $f(SE)$ have the form $g(x,y,t,SE)$ and $g(x,y,t,\theta)$, respectively. The first visual approach we developed extracts the raw data image sequences $g(x,y,t,SE_i), i = 1,2,...,n$ for $f(\theta)$ and $g(x,y,t,\theta_i), i = 1,2,...,n$ for $f(SE)$ and displays them side-by-side in one view as shown in Figure 6a. This approach is effective for performing cross-correlation of the data, but will not work for a large number of values for the dimension associated with the separate image sequences due to the limited viewing space.

The next approach was to place a glyph above the spatial position for each data point to represent one of the data elements in the remaining two dimensions (θ, SE). Figures 6b and 6c show examples using sphere glyphs and arrow glyphs to represent values derived from the SE and θ dimensions, respectively. Single and/or aggregated values for the dimension variables or the associated range values for both cases can be applied to the properties of the glyphs. For example, Figure 6c shows the heading θ for each position (x,y) by placing an arrow glyph above the position and orienting it in the direction θ. Some of the aggregations we performed were mean, variance and maxima. The maximum works best for indicating the likeliest locations for targets and the mean and variance operations show a general view of the distribution of the likelihood value across the dimensions.

To represent the sets $f(SE)$ we keep in mind that showing the mean and maximum data is important so the end-

[2]IMAT, www.presearch-inc.com/imat.htm

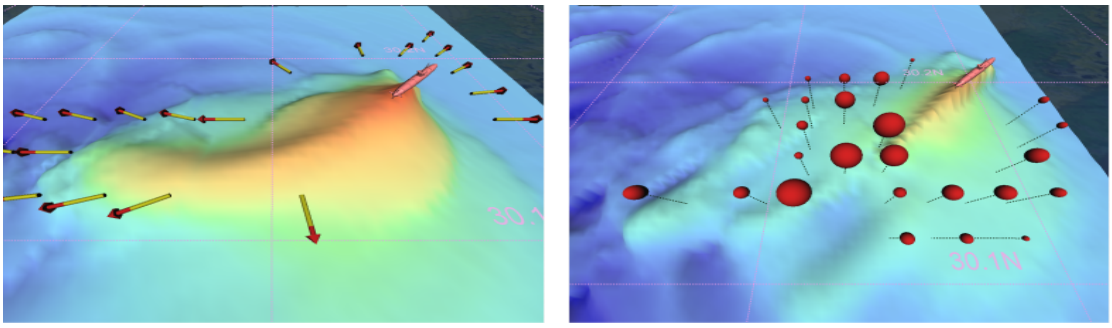

Figure 7: (a) Left shows the arrow-glyph approach with filtering of the most significant mean likelihood and least significant maximum likelihood regions. (b) Right shows sphere glyphs used for difference of maximum and mean likelihood for 4D space across SE with the same filtering as in (a).

user is made aware of threats (likeliest target locations) on a general scale across the whole dataset and for the worst-case scenarios. To represent the sets $f(SE)$ we applied the color-height surface approach for the mean likelihood across θ and used an arrow glyph to represent the heading and maximum likelihood values. The heading and maximum values are mapped to the direction and length of the arrow, respectively. Since the heading information does not directly correlate with the mean data, which is mapped to the height dimension of the color-height surface, we made a frame of reference for the end-user to perform cross-correlations between the two attributes. We used two colors on each arrow to represent the mean and maximum values. We colored the proportion of the length corresponding to the mean value yellow and the remaining portion red, which corresponds to the difference of the maximum and mean values. The heading associated with the mean computation is represented using the orientation of a submarine model placed at the maximum mean-likelihood value. The technique is illustrated in Figure 6c.

The multi-colored arrow-glyph approach works better for local observations than wide-area for conveying the significant threat information. The glyphs, which represent the information associated with the maximum likelihood values across headings, clutter the display preventing the user from being able to clearly comprehend the meaning of the surface data. Based on these observations, we modified the approach by removing the arrows that are associated with a small difference between the maximum and mean values and overlay the space where the peaks exist for the mean data. We also normalized the lengths of the arrows and raised them to avoid occlusions with the surface. We show the improved approach in Figure 7a.

We used a similar approach to represent the sets $f(\theta)$ but instead used a sphere glyph. We mapped the mean likelihood values across SE to the color-height map and the difference of the maximum and mean likelihood values at those locations to the radius of the spheres. We again apply a threshold to the display in order to show only the most significant threats. We show this technique in figure 7b.

4.3 Five-Dimensional State Data

To represent the complete 5D state space, we discuss and demonstrate two of the approaches we developed here. The first approach takes the maximum of the likelihood values across SE and θ for each given spatial position and time, and displays them using the color-height mapping scheme. This technique, shown in Figure 8a, shows the worst case threat areas for all values of mean signal excess and heading for the target. Notice the additional surface features that appear versus the detail from the previous 4D color-height surface mappings shown in Figures 7a and 7b. The elongated red ridge in the figure shows the possible positions where the submarine could be located for the specific time t. As before, we placed a model and text at the location of maximal value and oriented it with respect to the heading for that location and likelihood value. This approach shows the most significant information throughout the whole dataset at once, however, this does not show all the useful statistical information (e.g., what values for each dimension produced the maximum likelihood) that would be interesting to a statistician developing the models for the state-space uncertainty.

We extended the visual approach to include the important statistical information by combining the sphere and multi-color arrow glyphs to make a new glyph. The sphere size corresponds to the value of SE which is used to compute the maximum likelihood value. The arrow portion of the glyph again maps the correlation between the mean and maximum likelihood values (i.e., total length corresponds to maximum, length of yellow portion corresponds to mean and direction corresponds to heading θ). In the Figure 8b we mapped the mean across the complete dataset at time t to the primary surface. The sphere-arrow glyphs then show the corresponding mean in yellow along the base of the arrow, and the red portion indicates how likely a target could be located at the specified position. The idea is to draw the user's gaze towards the arrows with the longest portions of

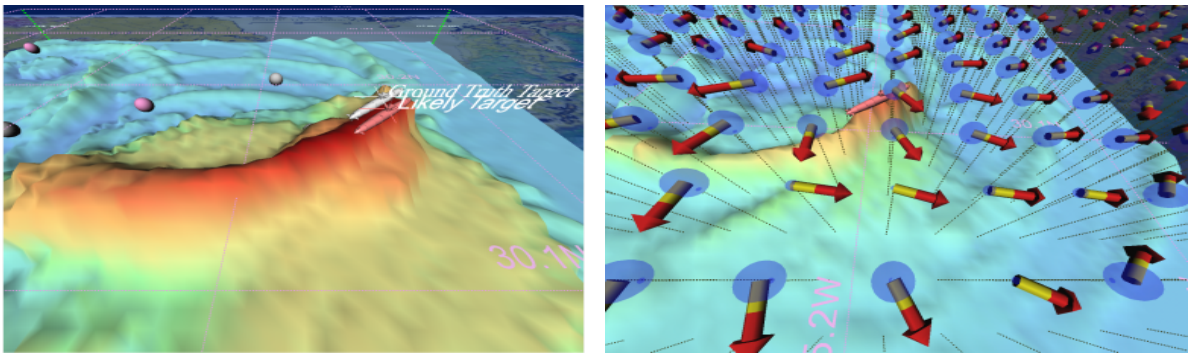

Figure 8: (a) Left shows surface of maximum likelihood across 5D dataset, and models and labels for ground truth and predicted target locations. (b) Right shows sphere-arrow glyph with value for SE mapped to sphere, difference of maximum and mean likelihood, maximum likelihood and heading mapped to arrow and maximum likelihood using color-height map.

red which indicate where the target's are likeliest to be other than what is indicated from the surface color-height mapping. The sphere-arrow glyphs do somewhat clutter the display when not filtered, but still allow the most significant information in the dataset to be conveyed to the end-user. More work will be addressed in the future on this approach and higher dimensional datasets.

5 Display System Architecture

We designed a display architecture that interoperates with multiple virtual reality (VR) and non-VR display systems (e.g., immersive room and desktop). The architecture consists of a desktop and immersive-room interface, a distributed state-based graphics engine and a data synchronization unit which ensures that each display panel has data for the same time state. We are able to run standalone applications on the desktop and in our four-wall immersive room driven by an eight-node graphics cluster. We utilized OpenGL for the graphics engine and GLUT and VR Juggler[3] for the desktop and immersive-room interfaces, respectively.

The desktop interface uses mouse input and graphical pull-down menus. The immersive-room interface utilizes a tablet PC and 6-DOF flightstick in order to provide interaction with a large number of options. The flightstick enables the user to navigate the environment and to select a few interface options using the buttons. The tablet PC contains a Java interface with sliders and buttons for the interface options. We linked the two devices to the distributed state-based graphics engine using VRPN[4] and a wireless CORBA-based network protocol. The graphics engine utilizes a state vector for the display and interface states. The state engine is distributed to each cluster node and updated by the data synchronization unit (part of VR Juggler). Some of the states include shading mode (e.g., wireframe, transparency and lighting), sea-bottom and target-state display methods and toggle switches for information-layer features (e.g., submarine model, text labels and graphic indicators). We show a view of our system operating in the immersive room in Figure 9.

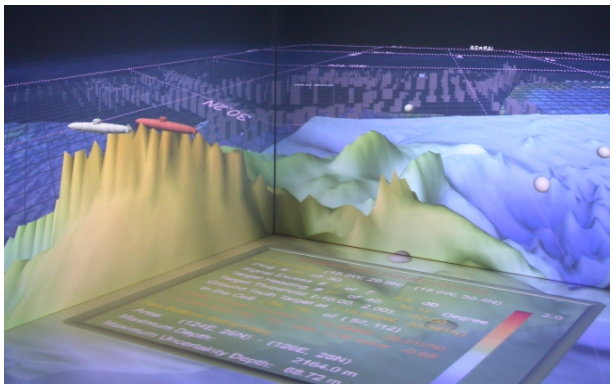

Figure 9: Display System Operating in Immersive Room

6 Summary and Conclusions

We developed visual representations for real multi-dimensional sea bottom uncertainty data, multi-variate information and multi-dimensional target state data. We explored the use of coloring methods, surface representations, glyphs, haptics, iconographic imagery, animation and text to represent the uncertainty, feature information and associated data. We designed an interoperable VR-display system to explore the visualization approaches and the use of haptics for information representations.

Coupling previous work with our own investigations, we compiled the following set of guidelines for developing good visual representations for multi-dimensional uncertainty and multi-variate information.

- Determine a basis for reference when distinguishing values of information (e.g., boxes on plane).
- Cross-correlation of visual representation of error with data is important.

[3] VR Juggler,www.vrjuggler.org
[4] VR Peripheral Network, www.cs.unc.edu/Research/vrpn

- Plan visual approach according to primary and secondary features.
- Plan visual approach for wide-area and local-scale observations where appropriate.
- For 3D representations, the effectiveness of the visual approach should consider the changing viewpoint.
- In most cases, having consistent proportions for the dimensions of the data is desirable. There are exceptions where it is desireable to exaggerate proportions in one dimension to highlight the features of importance.
- Minimize interference between visual representations for uncertainty, multi-variate properties of the data, and the data.
- Presenting the data with multiple approaches or display systems may be better at conveying the information.

These guidelines are general enough to apply to other task domains besides underwater environmental uncertainty.

The visual representations that we developed show a much richer summary of the underwater environmental uncertainty than previous displays. We have presented the visual techniques and software to several U.S. Naval commanders and domain experts in the acoustics field; the initial feedback has been very positive.

Now that the Naval utility of these techniques has been established, demonstration and validation in the fleet are the next steps. This can only be accomplished through a formal user study with potential end-users as the subjects. Based on the initial feedback, we believe that the visualizations whose development we have discussed here provide more complete and accurate insight into the data. We hypothesize that the visualizations will lead to better analysis, and subsequently, better inferencing about the data. We hope that once the LRT system transitions to the U.S. Naval Fleet, it will lead to more formal developments of the visual techniques including end-users in future development from start to finish.

Acknowledgements

We thank Bob Miyamoto and Marc Stewart (APL-UW), Dan Fox and Jim Fulford (NRL-SSC), Larry Stone (Metron), Brian LaCour (ARL-UT), Brian Calder (CCOM-UNH), Erik Tomlin and Alex Pang for research contributions. This work was supported by the Office of Naval Research.

Side Bar 1: Bathymetric Data Source

Bathymetric data exhibits four distinct physiographic regions: continental shelf, continental slope, continental rise and a trough. The continental shelf is the submerged part of a continent, which typically has depths of 100 m to 400 m, a slope around 0.4% and extends on average about 50 km from shore. The continental slope is the outer edge of the shelf, which extends to a depth as great as 4000 m and has an average slope between 3% and 6%. The continental rise begins the transition between the continental slope and the deep ocean basin and generally inclines at about 1.7% [9].

The dataset we used was generated by NRL-Stennis[5] from the Digital Bathymetric Data Base – Variable Resolution (DBDBV) [10]. The data covers a $2^\circ \times 2^\circ$ region with resolution of 30 s $\times$ 30 s. It was then decimated to a resolution of 120 s $\times$ 120 s, for a final mesh of 61×61 grid points. The data has the following depth ranges: 80 m to 180 m for the continental shelf, 180 m to 1500 m for the continental slope/rise, and a trough at depth beyond 1500 m. The data also has uncertainty estimates for the measurements.

References

[1] C. Ehlschlager, A. Shortridge, and M. Goodchild. Visualizing spatial data uncertainty using animation. *Computers in GeoSciences*, 23(4):387–395, 1997.

[2] E. Fauerbach, R. Edsall, D. Barnes, and A. MacEachren. Visualization of uncertainty in meteorological forecast models. In *Proceedings of ICA Commission on Visualization*, August 1996.

[3] D. Howard and A. MacEachren. Interface design for geographic visualization: Tools for representing reliability. *Cartography and Geographic Information Systems*, 23(2):59–77, 1996.

[4] C. Johnson and A. Sanderson. A next step: Visualizing errors and uncertainty. *IEEE Computer Graphics and Applications*, 23(5):6–10, September 2003.

[5] S. Lodha, C. Wilson, and R. Sheehan. LISTEN: Sounding uncertainty visualization. In *IEEE Visualization '96*. IEEE, October 1996.

[6] A. Pang, C. Wittenbrink, and S. Lodha. Approaches to uncertainty visualization. *The Visual Computer*, 13(8):370–390, 1997.

[7] S. Smith, L. Alexander, and A. Armstrong. The navigation surface: A new database approach to creating multiple products from high-density surveys. *International Hydrographic Review*, 3(2):2–16, August 2002.

[8] L. Stone, A. Barlow, and T. Corwin. *Bayesian Multiple Target Tracking*. Artech House Publishers, January 1999.

[9] Paul Tchernia. *Descriptive Regional Oceanography*. John Wiley and Sons, July 1980.

[10] U.S. Naval Oceanographic Office. Digital bathymetric data base MIL-PRF-32030, May 1998. http://www.nima.mil/ast/fm/acq/dbdb_v.pdf.

[5] NRL Stennis, www.nrlssc.navy.mil