Effective Object Detection with Bathymetric Sonar Systems for Post-Storm Response

Val Schmidt and Kevin Jerram

Center for Coastal and Ocean Mapping

Durham, New Hampshire

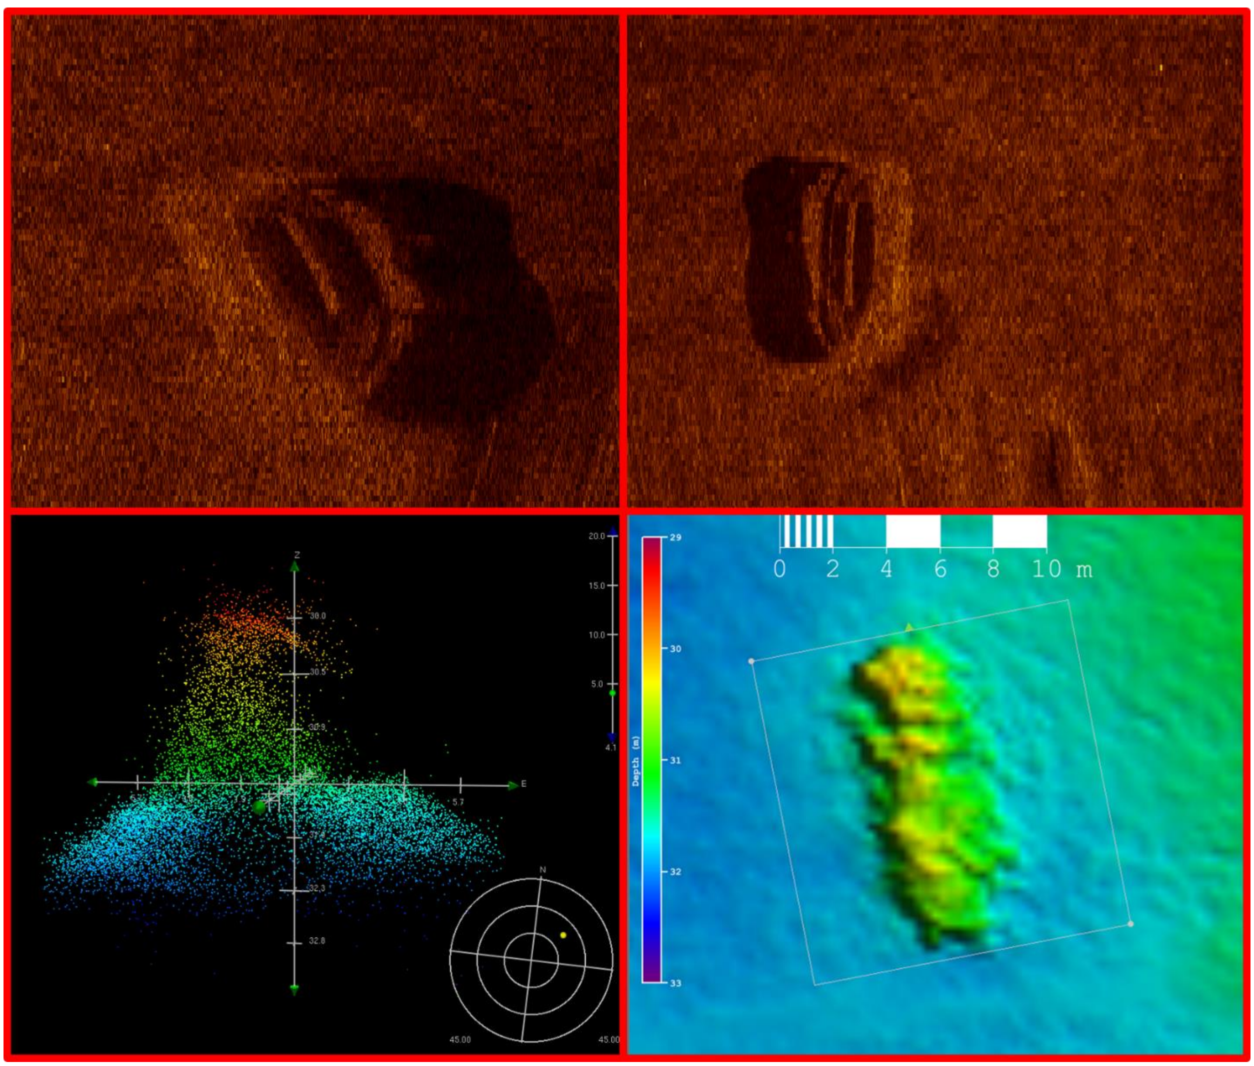

A sunken vessel is evident in several data products from a phase-measuring bathymetric sidescan (PMBS) sonar system. These data products include sidescan amplitude imagery from two adjacent survey passes (upper two images), raw soundings (bottom left), and gridded bathymetry (bottom right), providing different contexts and details for object detection and identification.

Introduction

In the wake of a major storm event, it is important to quickly survey navigable waters to assess the impacts of the storm on the seafloor, to identify marine debris for remediation, and to detect potential hazards to navigation. The response to Super Storm Sandy included the collection of several hydrographic data sets in impacted areas which, when combined with modeling of sonar operations and other empirical tests, provide a mechanism for establishing a set of best practices for data collection and processing in these scenarios.

This document describes the 'lessons learned' from examining these datasets with the goals of detecting objects and identifying the smallest observable changes in seafloor features. Importantly, these evaluations provide an estimate of the minimum size object or seafloor feature that is likely to be detected and accurately measured under a variety of survey setups. Although multibeam echo sounders (MBES) are the workhorses of NOAA hydrographic operations, phase-measuring bathymetric sidescan (PMBS) systems have also been evaluated. PMBS systems have advanced in recent years to improve the quality of bathymetric data and may offer post-storm survey efficiency gains either alone or in conjunction with traditional multibeam surveys.

An important consideration for effective and efficient data analysis is the ease of user interaction with the data. Routine steps in the acquisition and processing stages often require multiple software packages to examine soundings (e.g., raw or binned depth estimates provided directly by the echosounder as well as water column and sidescan data), apply filtering and gridding methods (e.g., CUBE processing), and review data products (e.g., depth, shoal, uncertainty, and hypothesis count layers). This disparity of processing steps prevents a holistic view of the data and its processing stages necessary to optimize the detection of hazards to navigation. To address the issue, this document also includes a conceptual user interface that could streamline analysis for object detection in a storm response scenario by presenting these data throughout the acquisition and processing stages, improving the user's ability to correlate bathymetry and sidescan imagery, and minimizing or eliminating transitions between software packages.

Best Practices for Post-Storm Operations

Here the best practices are listed concisely for ready reference with further discussion and analysis provided in the following sections.

1. **Collect acoustic water column or sidescan imagery.** Interpretation of bathymetric data sets for object detection is significantly enhanced by the availability and use of full resolution acoustic imagery, such as sidescan backscatter amplitude for PMBS systems and water column backscatter data for MBES. Where the inspection of water column data is possible, it is preferable to seafloor backscatter data from MBES systems.

2. **Examine data in real time.** Small objects that are readily apparent in acoustic imagery frequently are not reliably captured in bathymetric measurements. For this reason, operators must scrutinize imagery in real-time or near-real-time, identifying potential hazards quickly enough to provide opportunities for additional data collection at a variety of headings and athwartship distances to better quantify their shapes and least depths.

3. **Scrutinize the outer swath.** Both dual-head MBES and standard PMBS systems provide the capability to generate swath widths beyond 65° (~4.25 times water depth, or WD) from nadir. When refraction conditions are favorable, swath widths meeting IHO Order 1 standards for bathymetry are possible exceeding 75° (~7.5 WD), and object detection capabilities are enhanced by acoustic shadowing at these shallow angles. However, the object measurement capabilities of both systems typically suffer in the sounding data beyond 65°, such that these portions of the swath cannot be used reliably for object quantification. A strategy is required that carefully scrutinizes the outer swath for hazards to navigation and provides opportunity to revisit them in more favorable geometries for proper quantification.

4. **Utilize systems with real-time uncertainty and in-ping averaging techniques.** When using NOAA's standard metrics for meeting IHO requirements for hydrographic surveys, PMBS systems providing the capability to estimate real-time measurement uncertainty and to bin and average the raw bathymetric data across each ping have been shown to perform well and fit readily into the existing data processing pipeline. Bin averaging can be used to reduce sounding uncertainty and data density to acceptable levels while still maintaining significant features. In addition to examining the raw data, the ability to vary the bin size or method for processed soundings is also advantageous when investigating suspected targets. Acoustic imagery such as sidescan backscatter should always be retained at the full resolution for scrutiny of targets.

5. **Use PMBS systems without real-time uncertainty and binning only with great care.** PMBS systems without real-time uncertainty of their bathymetric measurements and binning and averaging of data can still be used effectively for object detection and meeting of IHO standards for hydrographic survey. However, the methods used for processing these data sets and their quantification for IHO purposes is fundamentally different than other systems. Specifically, because the data itself must be used to estimate the uncertainty, great care must be taken to filter data for outliers and not the tails of a noisy measurement distribution. Moreover, because of the noisiness of the data and one's inability to quantify the quality of any individual measurement other than statistically when taken as part of a whole, one must use gridded surface estimates in lieu of individual soundings to quantify potential hazards to navigation.

6. **Generate CUBE surfaces using NOAA guidance for shallow water.** MBES and PMBS data should be gridded according to NOAA's current standard practices. Specifically, it is

recommended that the CUBE algorithm be applied at 0.5 m grid spacing for water depths less than 20 m. NOAA's standard CUBE processing parameters ensure that no single sounding contributes to more than one grid node, thereby avoiding any 'smoothing' effects that could obscure detection of small objects.

When these parameters are used, it is estimated that the lower limit of object size detectable by routine visual inspection of a bathymetric surface of this type is 1m (the standard IHO requirement for object detection), *only* when the bathymetric uncertainty of the seafloor (taken as the standard deviation of the combined seafloor roughness and depth measurement) is less than 40 cm (1-sigma).

Discussion

The following section provides further discussion and quantification of the best-practices identified above.

1. **Collect acoustic water column or sidescan imagery.**
   The utility of MBES water column data and PMBS sidescan data along with bathymetric soundings cannot be overstated in the search for marine debris and hazards to navigation in post storm assessment. These acoustic imaging methods often better define objects than soundings alone, as they more clearly delineate acoustic shadows indicating the presence of objects proud of the seafloor. Moreover, they provide indication of small objects that may go undetected by standard bottom detection routines.

   Figure 1 illustrates the ability to pick out pilings from dock piers clearly in sidescan data that are undetectable in the corresponding bathymetry. Note that MBES seafloor backscatter or snippet data is of limited utility for object detection. These data are extracted from amplitude measurements associated with the bottom detection in each beam, which are more likely to represent the acoustically hard bottom rather than objects protruding from the seafloor. Furthermore, unlike acoustic shadowing evident in sidescan data, MBES seafloor backscatter is typically free of acoustic shadows. While MBES seafloor backscatter and snippet data can be useful in recognizing deviations from natural seafloor geology, in general water column data from these systems is preferable for detection and quantification of marine debris.

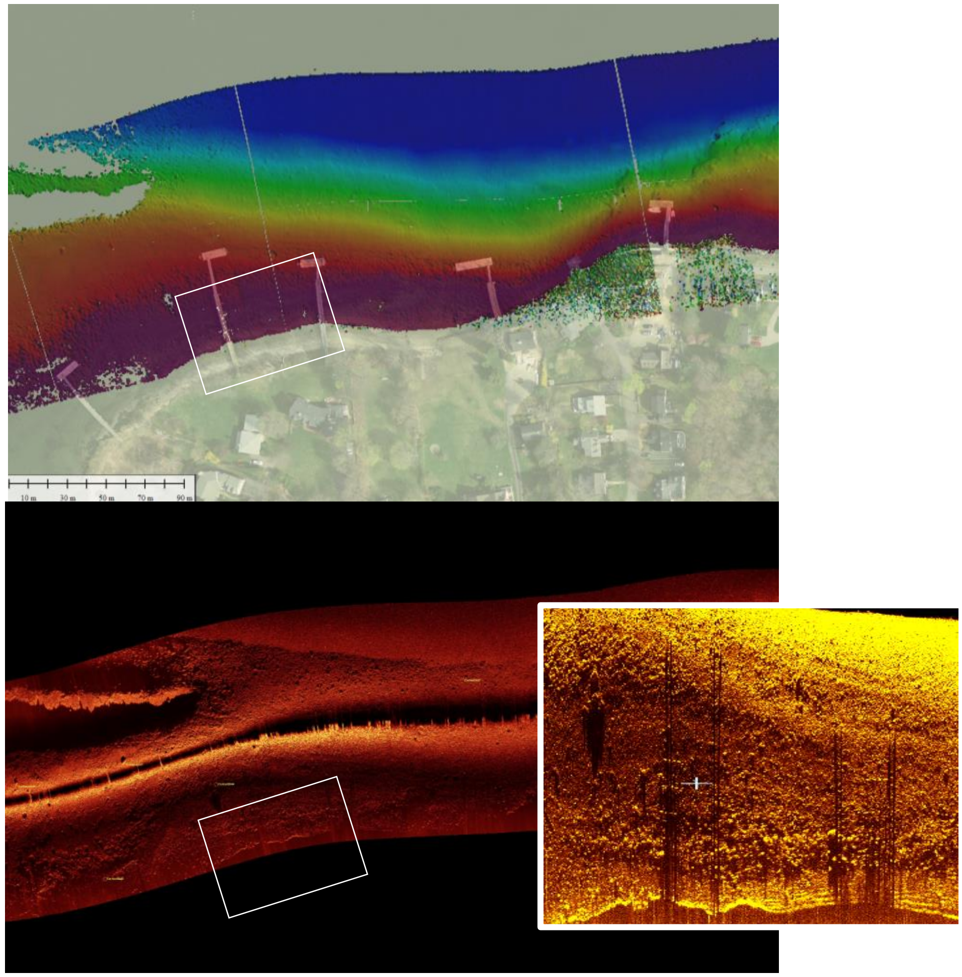

Figure 1. In this figure pier pilings, undetectable in the bathymetry data set (top), are clearly evident in the sidescan imagery (bottom and inset).

2. **Examine data in real time.**
   It is difficult for any bathymetric echosounder to reliably quantify objects proud of the seafloor under all geometries. A balance must be struck in sonar design and operation to capture dynamic changes in seafloor features that may represent hazards to navigation and to also to prevent copious false positives from fish, bubbles and other outliers. Further, objects that are strong acoustic reflectors are often captured in the sidelobes of adjacent beams. All too frequently, sidelobe detections of this type produce soundings both above and below the object. This effect can be seen in the test tank data illustrated in Figure 2a, where a manufactured rock of approximately 1 m in dimension is ensonified by a MBES. When the object is on the order of a grid cell size, the soundings above and below the seafloor tend to average out giving little indication of the object in the final bathymetric surface. When raw soundings are inspected in an area-based editor, where the geometry of the measurement is not clearly evident, these soundings may be dismissed as noise.

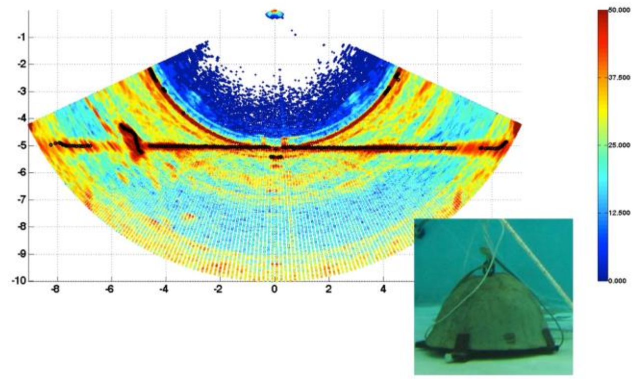

a)

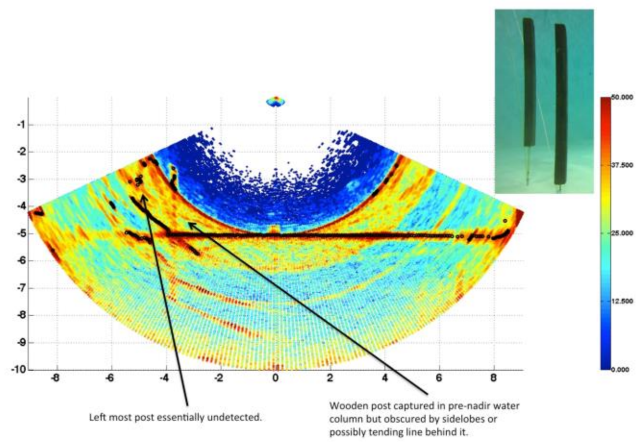

b)

Figure 2. These images depict MBES water column data overlaid with bottom detections (black points) operating in a test tank facility. Targets ensonified by the MBES include a 1-m synthetic rock (a) and two 6" x 6" x 48" posts. Sidelobe returns from the rock on the left side of the upper image cause bottom detections both above and below the nominal seafloor. In the lower image, the tank walls and a rope behind the posts tend to be more strongly represented than the posts in the bottom detections, especially when outside the nadir sidelobe ring.

Objects that are poor acoustic reflectors are typically not detected at all against a stronger acoustically reflecting seafloor. This effect is illustrated in Figure 2b, where two 6" x 6" x 48" posts, suspended from a test tank floor, are ensonified by a MBES. Only the nearest post is detected and, even then, only the portion of the post outside the nadir sidelobe return. Other portions of the tank or ropes holding the post in place are detected in lieu of the pole itself.

It is for these reasons that real-time scrutiny of MBES water column and/or sidescan data is recommended as a critical component of an object detection survey. Small objects may not be detected or adequately quantified on a single survey pass, requiring operators to scrutinize this data sufficiently quickly (if not in real-time) to afford the opportunity to conduct additional surveys without the undue expense of a second mobilization.

3. **Scrutinize the outer swath.**
   When complemented by water column or sidescan data, increasing swath widths beyond 65° can significantly increase the efficiency of surveys intended to locate hazards or areas for further study. For example, increasing the effective swath width from 65° to 75° by examining multiple data products results in nearly a doubling of the effective coverage on the seafloor as a function of depth beneath the sonar (assuming a flat seafloor). The value of this increase in coverage is especially high in shallow waters of 20 m or less, where swath geometry (not acoustic loss) is typically the limiting factor.

   The wider swath width provided by PMBS and dual-head MBES systems is largely due to the increase in SNR that results from the more favorable geometry relative to broadside of the array when measuring signals at the edges of the swath. However, for PMBS systems, because errors in the direction of arrival of the returning pulse are amplified by the long lever arms of a large swath, the variance in the soundings can increase beyond the capability to detect objects. MBES suffer from this effect too, although to mitigate the increased measurement variance, MBES systems average increasingly larger sets of measurements within each beam moving away from nadir. Whether by an increase in sounding variance or an increase averaging, both types of systems suffer a corresponding decreasing resolving capability in the outer swath.

   In addition, both system types use differential phase measurements for sounding determination and these measurements can suffer severely from multipath acoustic effects. Multipath conditions from the surface become more likely outside 60° from nadir and become common in general when objects present sharp corner geometries with the seafloor facing the echosounder. This effect has been shown to obscure the measurement of objects at the outer edges of a swath.

   For both system types, the along-track beamwidths also begin to subtend significant seafloor areas at longer ranges. While larger along-track beamwidths can be beneficial in object detection for imagery, as a single object may be observed in multiple pings, it can obscure one’s ability to quantify objects in sounding data.

   In addition, when operating with wider swath widths, all sonars are sensitive to under-sampled sound speed conditions and the refraction artifacts that result. Refraction artifacts of this kind generally present as characteristic “smile” or “frown” biases across the swath; when many swaths are combined into a single surface, the combined biases (especially in overlapping outer swath areas) can significantly hamper object detection. (Within a single line, however, these biases usually have little negative effect on object detection relative to other factors.)

All of these factors contribute to a general difficulty to reliably quantify objects that lie proud of the seafloor when measured beyond 60° from nadir. Nonetheless, mapping beyond 60° can still be done effectively both for detection of marine debris and safety of navigation, achieving IHO Order 1 and possibly better, when refraction conditions are favorable. These statements may seem contradictory, however other factors combined with a careful survey strategy make this possible.

For example, the extended acoustic shadows that can result from the geometry of the sonar and objects proud of the seafloor beyond 60° greatly enhance the ability to detect objects in sidescan and water column imagery. This effect affords an opportunity to detect objects in imagery for further investigation through additional data collection with the target closer to nadir.

The increasing beam width with across track range serves to both increase and hinder object detection capability in imagery depending on the object. For example, objects with low SNR that cast a small shadow relative to the ensonified area can go undetected, as the seafloor return tends to wash-out the shadow. This effect increases as the object's relative contribution dwindles relative to the rest of the ensonified area. This effect is unlikely to be seen in modern sidescan systems whose beam width is often as small as 0.3°. In water depths of 20 m, beams of this size would subtend only 40 cm at 80 m range. For objects with relatively large SNR, their detection capability can be enhanced with range from nadir, as one tends to observe strong returns from the object multiple times in the overlapping ensonified areas of adjacent pings.

Therefore, to operate with an extended swath width, extra care is required to scrutinize imagery data in the outer swath. Potential hazards to navigation must be identified prior to completing the survey and, if necessary, further quantified with sounding data acquired in more favorable geometries.

4. **Utilize systems with real-time uncertainty and in-ping averaging techniques.**
NOAA's standard process when evaluating surfaces for hydrographic compliance is to generate a CUBE surface. For sonars such as PMBS systems, the ability to create a CUBE surface that includes sounder uncertainty in the total uncertainty model is currently only possible if the sonar system itself provides real-time uncertainty for its measurements. Although models exist for PMBS sounding uncertainty, they have proved difficult to implement in a general way for the wide variety of processing methods in use by commercial systems.

   In addition, PMBS systems that reduce the data volume and data uncertainty by binning and averaging soundings within each ping more easily meet NOAA's metrics for verification of IHO uncertainty requirements. To understand this requires a brief explanation of NOAA's methods.

When a CUBE surface is generated, several ancillary surfaces are also generated. Among these the *uncertainty* layer and the *standard deviation* layer are used for verification of IHO requirements. The *uncertainty* layer is calculated as the mean of the estimated uncertainty for the individual soundings contributing to the selected hypothesis at the grid node. The *standard deviation* layer is calculated as the standard deviation of the soundings contributing to the selected hypothesis at the grid node. Therefore, the former characterizes the estimated uncertainty of the soundings while the latter characterizes the measured uncertainty of the soundings. The larger of the *uncertainty* layer and the *standard deviation* layer is used to verify IHO compliance. Ninety-five percent of the grid nodes must meet IHO compliance for uncertainty from these metrics.

These metrics were developed with MBES systems in mind, with the understanding that that MBES bottom detection methods average many individual measurements together to produce low-noise soundings (at the cost of reduced resolution, particularly in the outer beams). Accordingly, the standard deviation of the MBES bottom detections tended to be relatively low compared to that of the raw, pre-average measurements. These metrics were also developed when motion sensors and other correctors were less precise and there was considerable danger in assuming the measurements to be statistically independent. Therefore, the average of the estimated uncertainty in the soundings (the *uncertainty* layer) was used to assess IHO compliance rather than the uncertainty of the (weighted) mean depth. The uncertainty in the weighted mean depth decreases with the addition of measurements (when they are statistically independent) and in high-density datasets would approach zero even when the seafloor is sloped causing correlated measurements. Therefore, the average sounding uncertainty was deemed more appropriate, as it conservatively does not decrease with the addition of more measurements.

The result is that relatively noisy, densely populated data sets are penalized when attempting to meet IHO standards, as the addition of more data tends to increase the values reported in both the *uncertainty* layer and the *standard deviation* layer. For this reason, PMBS sonars that produce real-time uncertainty to support CUBE processing, and bin and average data to reduce the data volume and uncertainty, are recommended for surveys attempting to meet IHO specification using NOAA's metrics.

Several methods are possible for binning and averaging data. Generally, measurements should first be filtered for gross outliers and low SNR (< 10 dB, as a typical rule). Post-filtered soundings are then sorted into either fixed size bins by across-track range or fixed angular sectors and the median, mean or uncertainty-weighted mean is returned for each bin.

Two additional methods have been considered in this work. The "Most Probably Angle Algorithm" (Schmidt et al., 2012) returns either the most-probable angle within each range bin or the raw soundings falling within a window around the most-probably angle.

The window is chosen based on the size of the phase ambiguity sector for the largest aperture in the sonar. This method was found to produce high resolution soundings generally free of outliers, or low resolution soundings with very low variance.

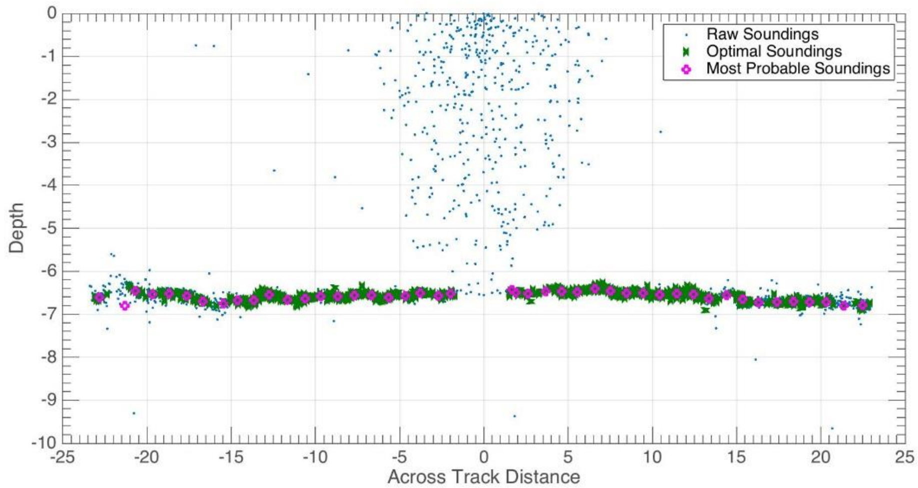

Figure 3. In this example, raw soundings from a PMBS system are plotted with those created from the "Most Probable Angle Algorithm", in which soundings are binned and only the most probable depth is returned (magenta diamonds). Also shown are the "Optimal Soundings" (green crosses) in which raw measurements (after filtering) are combined in an uncertainty-weighted mean to obtain a sounding result meeting a pre-defined desired uncertainty.

Another method was proposed in an attempt to optimize the tradeoff between resolution and variance in the measurements. This "Optimal Sounding" method (Schmidt et al., 2013) considers soundings and their uncertainty one at a time, comparing them to a desired uncertainty level and combining them into an uncertainty-weighted mean until the combined estimated uncertainty falls below the desired level. The results are variably sized bins producing soundings whose variance is constant across the swath. Soundings are produced at the maximum resolution of the system when signal-to-noise and other factors are favorable, reducing the resolution automatically when they are not. Figure 3 illustrates raw soundings, "most-probable soundings" and "optimal soundings" for a single ping of data.

5. **Use PMBS systems without real-time uncertainty and binning only with great care.** PMBS systems that do not provide real-time uncertainty metrics and/or within-ping binning and averaging can still be useful tools for object detection and hydrography in assessing post-storm impact. However, the methods used to process the bathymetry data and assess its IHO compliance are different from NOAA's normal methods.

   In real-time and in post-processing, it is wholly appropriate to filter PMBS data based on low SNR, poor coherence, high estimated uncertainty (if provided), a manufacturer-defined quality factor, or even unlikely measurement geometries. These soundings have little to no information for the purposes of object detection and it need not be a hydrographic conundrum to omit them. [When raw soundings are viewed without the context of their SNR or other metrics, it can be far more difficult to discount them.] It is helpful in later processing stages when this filtering/flagging can be done during acquisition.

What is not appropriate, although tempting to do, is to use aggressive along-track or across-track filters to omit soundings that deviate some fixed amount from the mean or median seafloor estimate measured over some window.

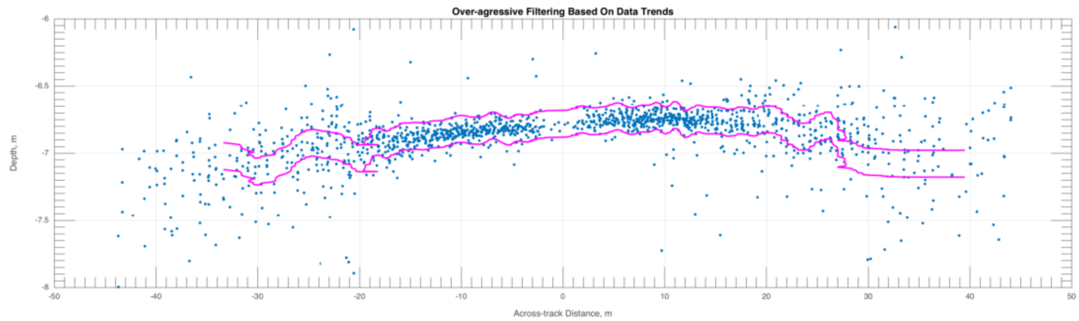

Figure 4. Here over-aggressive fixed-window filtering (shown in magenta) loses the trend of the seafloor when sounding uncertainty increases. Note the vertical scale is exaggerated in this image.

Figure 4 illustrates the case, in which an overly-aggressive fixed window filter from the mean seafloor fails to estimate meaningful depths when the sounding uncertainty grows too large. Not only does the window fail to capture the true depth, but any uncertainty estimate based on the retained soundings will grossly underestimate the uncertainty of those measurements.

The effect can be looked at quantitatively as shown in Figure 5. For this example, 20 synthetic measurements (vertical lines) are scattered about the “true” depth (0 m, center), with a 1 m standard deviation. A histogram of those measurements is shown in brown. Measurements falling less than +/- 25 cm (1/4 sigma) from the mean depth are colored blue, illustrating those measurements that might be retained in an overly aggressive filter. The standard error of the mean depth as calculated from all measurements is shown in magenta and the shaded area under this curve that falls outside +/- 25 cm represents the probability that the true depth has been excluded by this filter.

When this simulation is repeated many times, one finds there to be a 32% chance that the true mean of the data falls outside the filter window on average. If the measurements have correlated errors, such as those that might result from a sloping seafloor, sound speed error or other bias, the chance that the true mean falls outside the filter window increases to 79%.

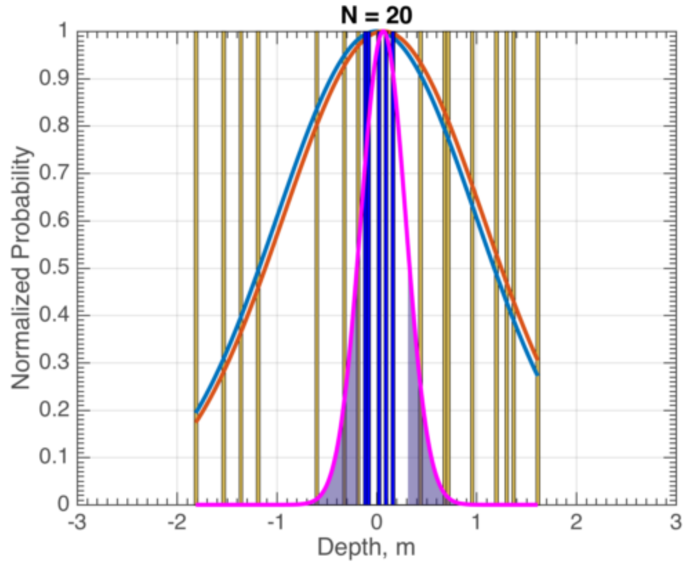

Figure 5. Here a histogram of 20 measurements (vertical lines) having a 1-m standard deviation from the true depth (0 m, center) is plotted along with the probability distribution of those measurements (blue). The standard error (the probability of its deviation from the true-mean) is plotted in magenta. Soundings falling within ¼ sigma from the calculated mean are shown in blue. The shaded area below the standard error curve and outside the filter window indicates the probability that the filter window fails to include the true depth.

Therefore, great care must be taken to avoid overly aggressive filtering of data that are not outliers but simply deviate from the mean because their uncertainty is high. *In other words, one should not aggressively filter the tails of the distribution of seafloor measurements.* When distribution tails are excluded, the true seafloor is not likely to be captured and the result will have artificial measurement statistics.

Note that these points at the tail of the measurement distribution are different from *outliers*, which we define as measurements of other things (fish, the surface, etc.) or measurements having very low SNR, indicating that they are, in fact, measures of nothing. It is perfectly appropriate to omit outliers during data processing.

When appropriately filtered to remove outliers rather than distribution tails, the data may be gridded by any number of mechanisms. While CUBE may seem an inappropriate method when no sounder uncertainty is available, it turns out to still be a formidable tool for gridding and outlier rejection, especially when other uncertainty metrics are appropriately specified. (Note: CARIS HIPS allows generation of CUBE surfaces without sounding uncertainty, but other software may not.) CUBE's multi-hypothesis solution will generate numerous hypotheses resulting from outliers that are inconsistent with the nominal seafloor and may be readily rejected, leaving a reasonable seafloor hypothesis in most cases. One must be careful to note that the calculated TPU, having no sounder uncertainty in these cases, may be inordinately and incorrectly small. When used in this way, CUBE acts to some extent as a median filter might, separating outliers from the most likely seafloor depth, and providing a reasonable estimate of that depth.

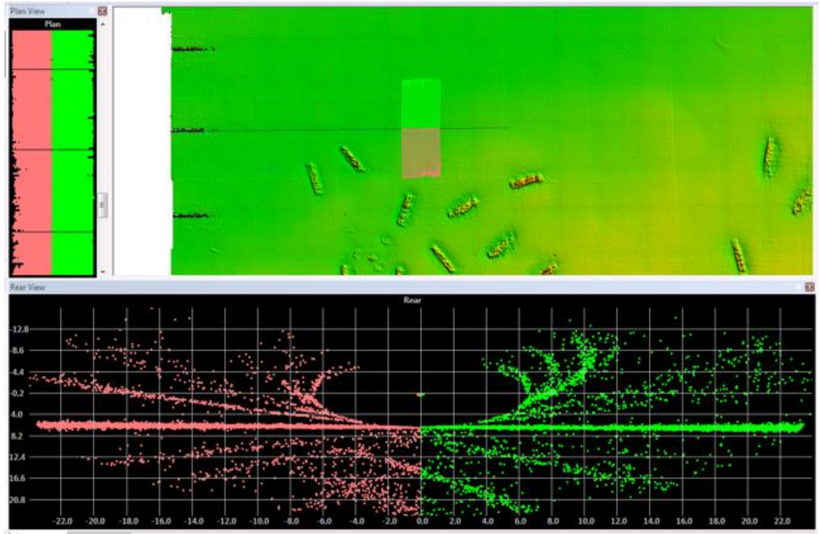

a)

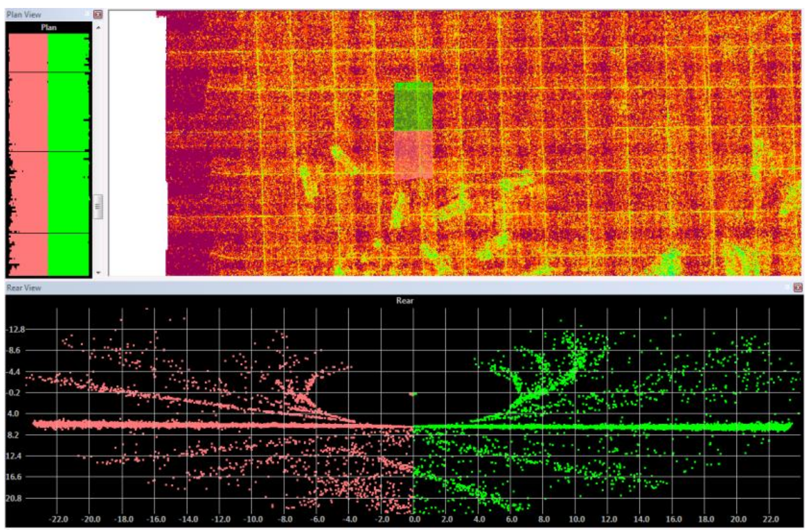

b)

Figure 6. In a) above, a portion of a CUBE surface from a GeoAcoustics GeoSwath survey of Red Bird Reef off the Delaware coast is depicted, along with an across-track profile of several pings below it. Total propagated uncertainty for soundings contributing to the CUBE surface was calculated without echosounder uncertainty. Nonetheless, CUBE's ability to generate multiple hypotheses at each grid node provided a powerful filter to estimate the seafloor with little manual editing. However, the method produces multitudes of hypotheses shown in b) where the color of the surface indicates number of hypotheses (yellow = 4-5 hypotheses).

Figure 6 illustrates what can be done with data from Red Bird reef off the Delaware coast. Large objects on the seafloor in these images are subway cars, part of an artificial reef laid over several years beginning in 1996. In this data set, only an amplitude filter set at the 10^{th} percentile was applied. As can be seen in the across-track profile plots, this filter only served to remove the water column data. Nonetheless, when gridded with the CUBE algorithm, the surface is realistically represented with few outliers as shown in the left image in the upper plan-view. However, because of the many outliers, multitudes of hypotheses are generated that prevent the common strategy with other data sets of using multiple hypotheses as a guide for further data inspection. The number of hypotheses are shown on the bottom image, where purple indicates 1 and yellow indicates 4-5 hypotheses. Because of the volume of data, many outliers, and increased noise of raw soundings, it can be challenging to determine least depths from the raw soundings. *When processing raw, unbinned PMBS data, it is more accurate to use the gridded surface itself rather than individual soundings to establish least depths.*

When raw PMBS data are processed with CUBE in this manner, the resulting surface's "Uncertainty" layer cannot be used for meeting IHO uncertainty requirements. The values in this surface depend on the predicted uncertainty of the depth estimate, which lacks sounder uncertainty and is therefore too low. The preferred method for meeting IHO compliance is to add the standard deviation of the soundings contributing to the accepted hypothesis at each grid node to the other sources of error that would not be captured in this empirical measure of survey uncertainty, namely, tides, heave and draft. Uncertainty from each of these components can be combined in a root-mean-square sum (i.e., adding of the variances of each component) to estimate the total uncertainty for the survey. This method will over-estimate the uncertainty on a flat seafloor, as no reduction due to increased data density is assumed. However, it will correctly approximate the uncertainty in depth over a seafloor that changes appreciably within a grid cell.

6. Generate CUBE surfaces using NOAA guidance for shallow water.

NOAA's Field Procedures Manual specifies guidance for creation of CUBE surfaces for NOAA survey teams. This guidance states that for water depths less than or equal to 20 m, gridded surfaces should be created at 0.5 m grid node spacing. Further, the CUBE parameters specified in Figure 7 are utilized to optimize the opportunity for object detection. These parameters set the "Capture_Distance_Min" field to .707 times the grid node spacing (0.5 m) to mitigate the chances that a sounding will contribute to more than one grid node and the "Capture_Distance_Scale" to 0.5% of water depth effectively removing the depth-dependent influence on capture radius. Other parameters are specified as shown below and similar guidance for other surface types can be found in NOAA's Field Procedures Manual.

Estimation of object detection probability in seafloor mapping data products is a complex topic, involving the physics of the sonar measurement, the data processing methods, the features of the ambient seafloor, and the data visualization methods. As such, a complete analysis of the topic is largely beyond the scope of this work. However, in an attempt to place a bound on what is possible and to gain some intuition regarding the answers, synthetic seafloor surfaces were created containing a 1-m cube object and presented to knowledgeable operators for identification of the object in the surface under typical visualization conditions. An example surface is shown in Figure 8, where an object approximately 1 m in all dimensions located at (73, 13) is barely identifiable.

```
 <ParameterSet Configuration_Name="NOAA_0.5m">
       <Comment value="The 2010 field season parameters for NOAA field units for the CUBE
algorithm on a 0.5 meter surface."/>
       <Distance_Exponent value="2.0"/>
       <Queue_Length value="11"/>
       <Quotient_Limit value="255.0"/>
       <Discount_Factor value="1.0"/>
       <Estimate_Offset value="4.0"/>
       <Bayes_Factor_Threshold value="0.135"/>
       <Run_Length_Threshold value="5"/>
       <Capture_Distance_Scale value="0.5"/>
       <Capture_Distance_Min value="0.354"/>
       <Horiz_Error_Scalar value="1.96"/>
       <Density_Strength_Cutoff value="2.00"/>
       <Locale_Strengh_Max value="2.5"/>
       <Locale_Radius value="1"/>
       <Null_Hypothesis_Min_Neighbours value="3"/>
       <Null_Hypothesis_Ratio value="3.0"/>
       <Null_Hypothesis_Strength_Max value="2.5"/>
       <Enable_Null_Hypothesis value="False"/>
    </ParameterSet>
```

Figure 7. NOAA CUBE parameters for 0.5 m surfaces.

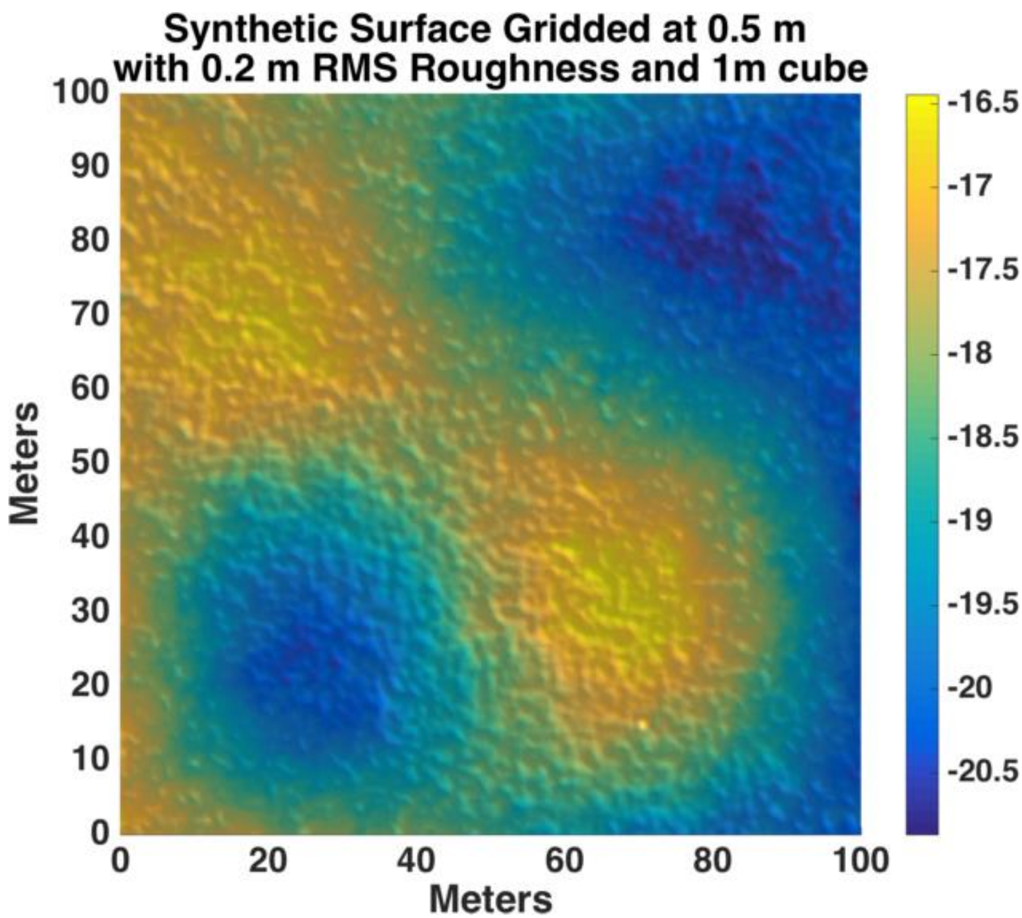

Figure 8. A synthetic seafloor with 20 cm RMS roughness with a 1 m cube object centered at (73, 13). The object is very difficult to detect visually, illustrating a practical limit to detecting objects in gridded surfaces.

In this cursory investigation, operators found that 0.5 m resolution surfaces having an RMS roughness of 20 cm provide a limit to reliable visual detection of 1 m objects. It is important to note that this roughness measure is the combined effect of the ambient seafloor and the measurement process (e.g., the sonar, motion sensor, etc.). In some instances, the seafloor is far rougher than the variance of the measurements. In other instances, the variance of the measurements is far larger than the roughness of the seafloor. However, it is their combined effect that impacts object detectability; therefore, object detection efforts over flat seafloors can tolerate more variance in the measurement process.

It is important to note that synthetic seafloor models like the one shown in Figure 8 do not look completely like a real seafloor. Real seafloors tend to have smoother 'lows' due to the sediments that tend to collect in troughs and rougher 'highs' owing to rock outcrops, reefs and other features that tend to erode slowly or resist smoothing. However, by providing a uniform roughness spectrum over all parts of the synthetic seafloor, a worst-case scenario is presented.

Conclusion

In post-storm response surveys, the clearing of waterways must be done efficiently to rapidly open ports and harbors and ensure safe navigation for response vessels and other ship traffic. The lessons learned by scrutinizing Super Storm Sandy data sets and follow-on investigations provide a framework for improving these surveys.

The data reviews presented here demonstrate that water column and sidescan imagery provide critical additional context to soundings by providing measurements that are not directly associated with the bottom detection. As such, they allow operators to better determine shoalest depths that might not have been captured by bottom detection algorithms, and to identify anthropogenic debris. While MBES continue to be the workhorse echosounders for NOAA's hydrographic operations, PMBS systems can provide meaningful increases in survey efficiency over single-head MBES systems by increasing the swath width that is useful for object detection and providing collocated sidescan for object identification. However, regardless of the sonar type, while increased swath width provides increased reconnaissance of likely navigation hazards, operators must be vigilant in identifying these hazards during acquisition (if possible) and revisiting these areas in more favorable measurement geometries.

PMBS systems that provide real-time uncertainty and options for binning and averaging data within each ping more readily fit into the existing hydrographic workflow. Operators of PMBS systems that do not provide these features may still achieve increases in survey efficiency and effectively detect hazards, but must be careful to avoid overly aggressive data filtering, being sure to omit only outliers and not tails of a noisy distribution of seafloor estimates. Filtering the tails of seafloor measurement distributions increases the risk of excluding the true seafloor depth, skewing the measurement statistics, and complicating the empirical estimates of uncertainty for IHO compliance.

Appendix 1: An Interface to Maximize PMBS Data Utility for Object Detection

A sonar operator attempting to acquire, process, and scrutinize data during an object detection survey typically requires several software packages to accomplish these tasks. The tools available for data visualization may not be conducive to rapid or simultaneous examination of a target area across multiple data products, complicating the object investigation process. A conceptual model for a data processing graphical user interface (GUI) has been developed which helps to address some of these challenges faced by operators, enabling a user to examine raw soundings, full-resolution sidescan imagery, and processed data layers simultaneously. This conceptual interface integrates several approaches and options currently employed through separate software packages to expedite reprocessing of raw echosounder data in one interface with on-the-fly adjustability for filtering, gridding, and visualization parameters.

Important for storm response surveying in shallow waters, PMBS systems offer sidescan imagery that can be used to extend the useful swath width for detection of objects at longer ranges and lower angles of incidence on the seafloor, where depth data from both MBES and PMBS systems typically become least reliable for this purpose. The conceptual data interface presented here highlights the importance of the full-resolution sidescan imagery for object detection and identification. This information is traditionally presented in a 'waterfall' display with no direct context for ship track or geographic location of imaged targets. By contrast, the conceptual interface supports steps toward 'georeferencing' the images for improved ease of correlation with the bathymetry data. The interface features and their purposes are described below, with examples shown in Figures A**Error! Reference source not found.**-AFigure A.

Note that while the conceptual interface and examples presented here have been developed to streamline the review of PMBS data for detection of object and seafloor features, this approach could be readily applied for MBES systems by providing simultaneous evaluation of raw soundings, gridded layers, and full-resolution water column data. Ideally, an interface streamlined for object detection would support the display of PMBS and MBES data together, including tools to review all available depth, sidescan, and water column data seamlessly.

Processing Control

The gridding control panel provides options for rapid reprocessing of raw sonar data into gridded bathymetry products. As a first step, the user selects the following filters to apply to the raw soundings before gridding.

> **Range**: Exclude soundings within a minimum and/or beyond a maximum range to eliminate range-related artifacts (e.g., transmitter ring-down at the start of the sample record or uniformly low-amplitude data at excessive ranges).
>
> **Angle**: Exclude soundings arriving at angles outside the desired swath (measured from nadir) on the port and starboard sides (e.g., expected multipath or other interference).

Depth: Exclude soundings outside an acceptable/expected depth range for the survey; depth is measured from sea surface after accounting for draft and tide, as opposed to echosounder altitude above the seafloor.

Signal-to-Noise Ratio: Exclude soundings failing to meet a minimum SNR threshold.

After filtering, the user may choose to proceed using the remaining set of 'raw' soundings directly from the echosounder or apply a binning technique to generate depth estimates with reduced data volume and noisiness. The simplest approaches are to group the raw soundings into angle or range bins and take the mean or median depth estimate in each bin.

Angle Binning: Bin raw soundings by angle (e.g., 0.5° or 1.0° increments) and generate a single depth estimate at the center of each angle bin from the mean or median range of filtered soundings within the bin.

Range Binning: Bin raw soundings by range (e.g., 30 cm increments) and generate a single depth estimate at the center of each range bin from the mean or median angle of filtered soundings within the bin.

Additional options such as outlier rejection (or 'despiking') within the bins are presently available in other software packages and would be desirable in this conceptual interface. One important consideration for this step is the ability to generate depth estimates across surfaces which may be detected at the same range. These surfaces may be apparent in the raw data (especially for echosounders capable of producing multiple angles to targets for each range sample) but obscured or confounded in the range-binned results. For example, the submerged base of a retaining wall may include flat seafloor and vertical wall surfaces at the same range from the echosounder. Even if detected separately in the raw soundings provided by the echosounder, range-binning would likely lead to erroneous mean or median angle estimates for range bins which included multiple raw data points on the seafloor and wall surface. Angle binning would be preferable for preserving this sort of feature, with the trade-off of decreasing acrosstrack resolution in the outer swath on flat, benign seafloors. The conceptual interface provides the user with options to move forward using raw (direct from the echosounder), angle-binned, or range-binned soundings for the gridding process.

After optional filtering and binning strategies have been applied, the Combined Uncertainty and Bathymetry Estimator (CUBE) algorithm (used widely among existing software packages) is used to generate gridded depth estimates and additional data layers. (CUBE processing parameters such as vessel and sensor uncertainty values, IHO survey order, and disambiguation method would be user-configurable in a separate window.) The user completes the filtering and gridding process for the selected parameters by selecting the 'Process Grid' button.

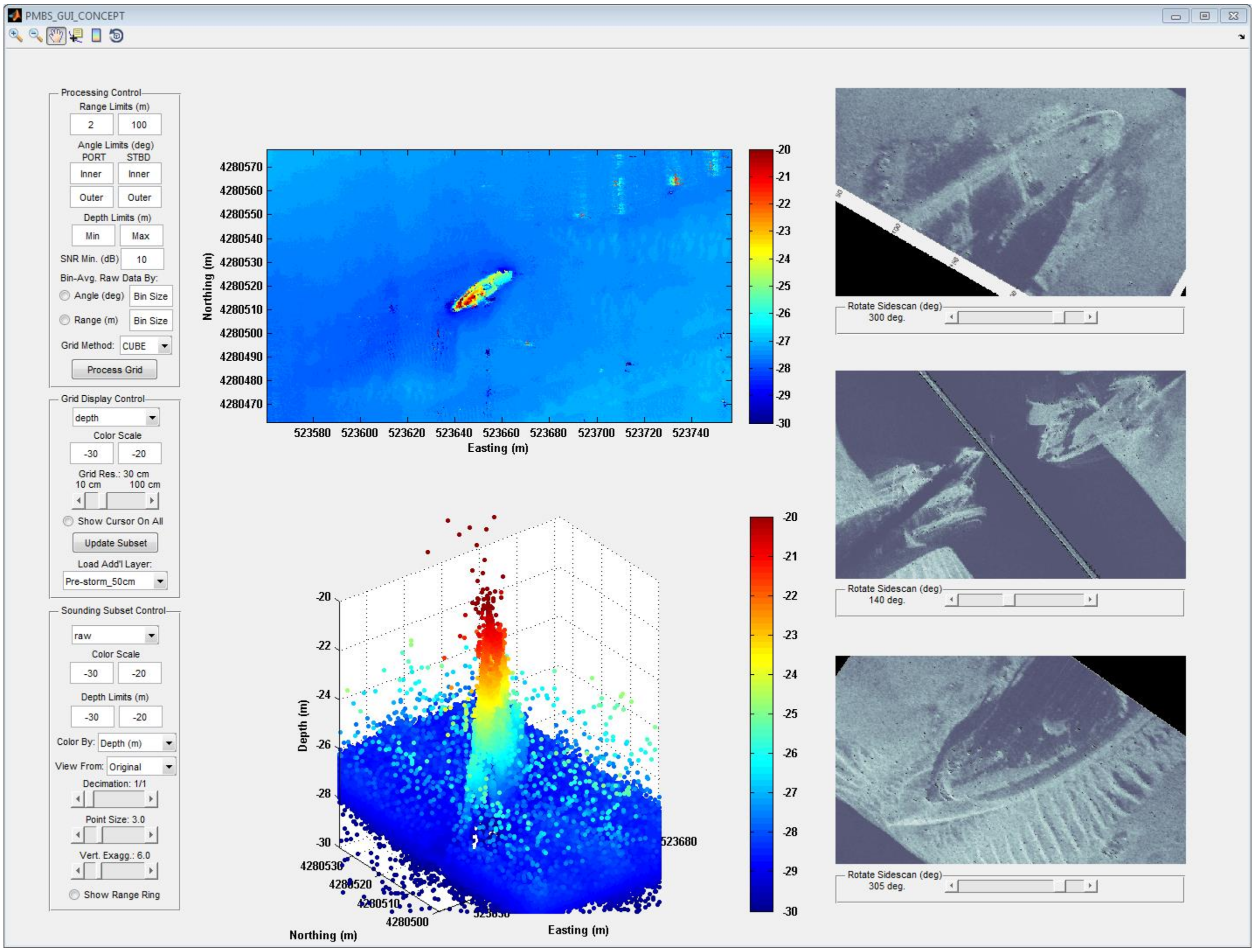

Figure A1. A conceptual PMBS graphical user interface (GUI) presenting raw soundings, gridded bathymetry layers, and sidescan imagery simultaneously for maximum data utility during review for object detection. Data shown here include a shipwreck and sand wave field at Redbird Reef, Delaware, surveyed with an AUV-mounted PMBS system shortly after Super Storm Sandy. This conceptual GUI combines visualisation tools that are traditionally available only in separate software packages, potentially enhancing object detection by streamlining the visual correlation of targets in different data products. Additional GUI features for this correlation are presented below in other figures.

Grid Display Control

The resulting data layers are gridded at a variety of resolutions (e.g., 30-100 cm in 20 cm increments, depending on depths and data density) for initial review. The automatic generation of multiple grid resolutions, grid layer drop-down menu, and grid resolution display slider with color scale control enables the user to rapidly explore trends in depths, sounding densities, hypothesis counts, uncertainty estimates, and shoal and deep layers (among others) to better understand the useful limits of the data under the filtering, binning, and gridding parameters applied.

Selecting 'Show Cursor On All' generates georeferenced cursors on all other display panes to improve the visual correlation between gridded depth estimates, raw data, and sidescan imagery. When a gridded area is identified for closer visual scrutiny (e.g., through depth or uncertainty anomalies), the user may select the 'Update Subset' button to load the raw soundings and any sidescan imagery available for the region in the grid display. Finally, the user may load additional data layers created externally, such as bathymetry from previous surveys in the area to facilitate more rapid identification of apparent depth changes or feature migration.

Sounding Subset Control

While gridded data products are useful for reviewing large survey areas, regions identified for closer inspection require an environment in which the user can more closely inspect the raw data contributing to the gridded layers. The conceptual interface provides a 3-D frame similar to existing data editors for manipulation of the soundings contributing to the gridded area of interest. The user may select the sounding type (e.g., raw soundings provided directly by the echosounder, or other soundings generated after filtering, binning, or other strategies have been applied); adjust the color scale and source (e.g., color by depth, amplitude, or line number) and scale; and exclude data outside a particular depth range.

Other on-the-fly filters could be added, such as a minimum backscatter amplitude threshold which acts independently from that applied in the gridding step. Because the raw datasets may be particularly large or dense, an option is provided to decimate the data for display; this is straightforward ('dumb') decimation by factors of two, though other decimation strategies could be made available. Conventional display options such as point size, vertical exaggeration, and viewing angles along the cardinal directions are also available for rapid adjustment.

In order to provide context and enhance the visual correlation between raw soundings and sidescan imagery, a 'Show Range Ring' option is provided. This feature plots a 3-D constant-range 'ring' in the subset display that corresponds to the ping number, orientation, and range of any point selected in the sidescan display. When identifying apparent shallow targets in the sidescan imagery (which does not include any angle information), this feature provides a reference in the subset plot to better identify nearby or related soundings on the potential targets (Figure A).

Sidescan Imagery

The imagery from all passes over the gridded area of interest and sounding subset are presented simultaneously with corrections for vessel heading on each pass. This approach provides a more comprehensive and intuitive overview compared to the traditional 'waterfall' display of a single survey line in vertical orientation. An option to more rigorously georeference and/or 'stack' the sidescan imagery as user-selectable layers in a single display may further improve this approach. The 'Show Cursor On All' and 'Show Range Ring' options provide instantaneous position updates in all other panes based on the cursor location in the sidescan display.

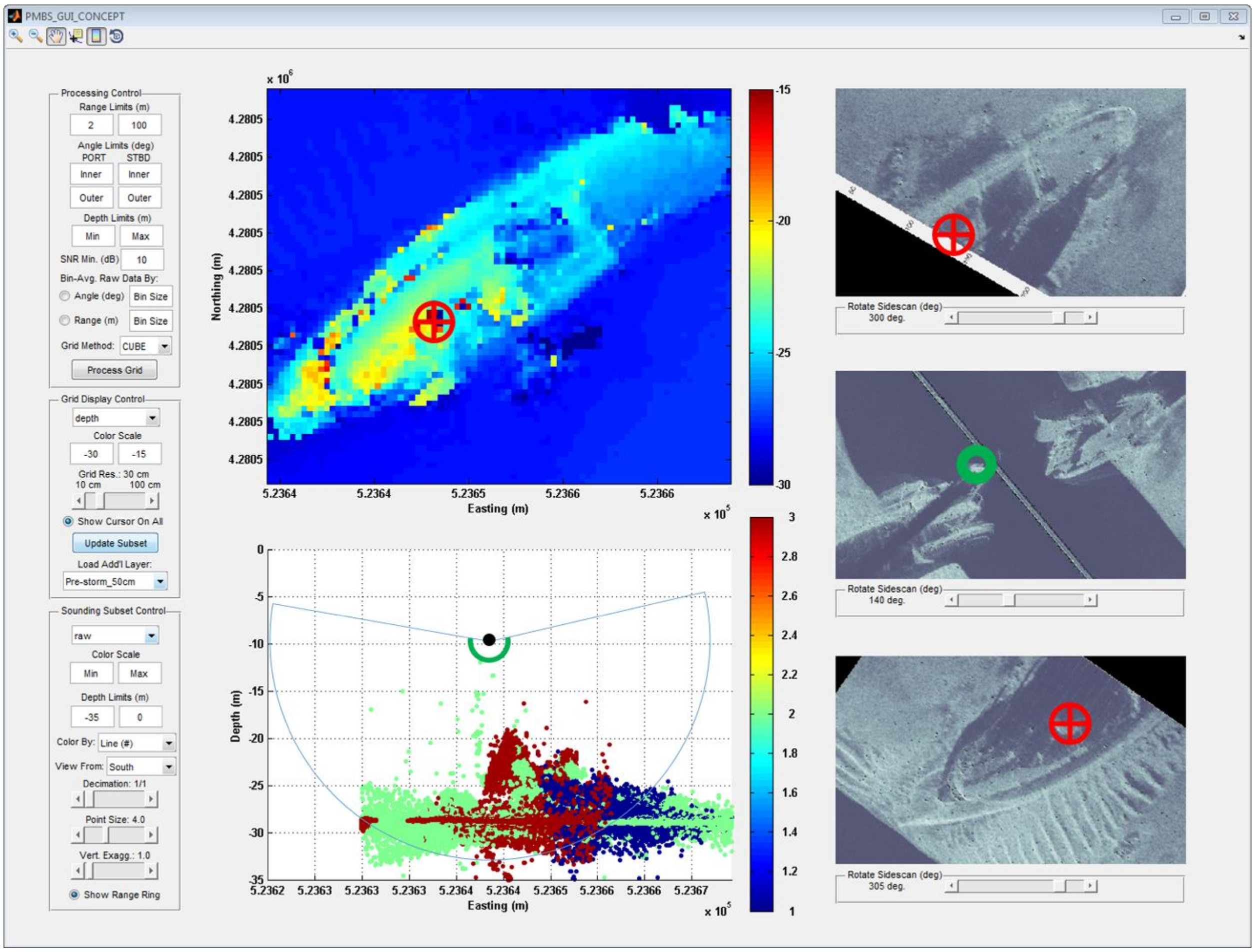

Figure A2. A closer view of the shipwreck with the 'Show Cursor On All' option selected helps the user to identify the same location (red crosshairs) in the gridded bathymetry and the sidescan imagery from multiple passes for better context and correlation among targets. With the 'Show Range Ring' option selected, the range to a high-amplitude target selected in one sidescan image (green cursor, middle right) is depicted in the 3-D subset (green range ring, lower left); this option helps to confirm the relationship between shallow soundings in survey line 2 (light green soundings) and the corresponding shallow sidescan feature.

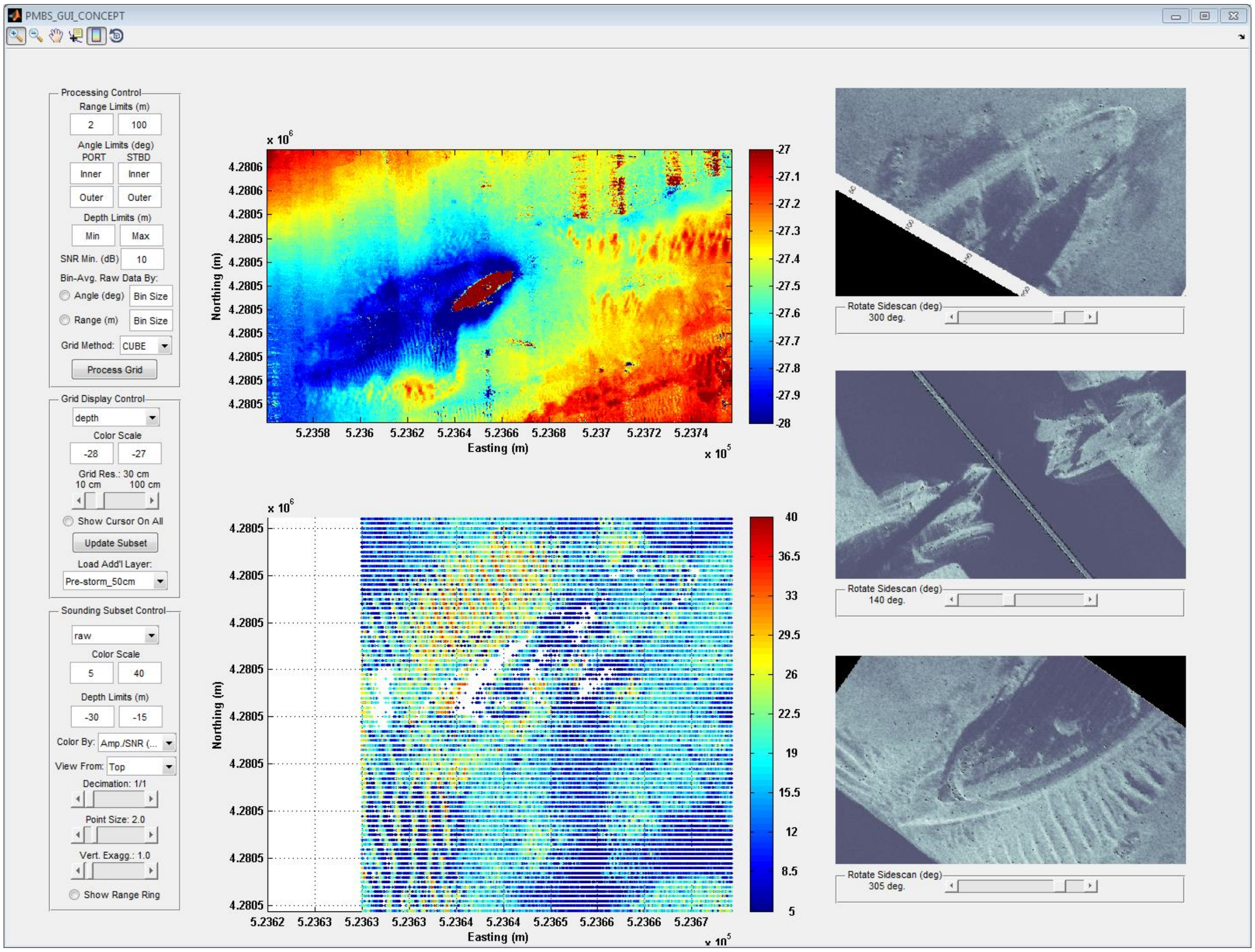

Figure A3. Aside from highlighting the shipwreck, rapidly adjustable display parameters can help an operator emphasize seafloor features such as the nearby sand wave field. Coloring the raw sounding subset by SNR (bottom) provides an outline of the vessel and is particularly useful for identifying the waves across multiple depth ranges, where these features are obscured by the single depth color scale in the gridded depth layer.

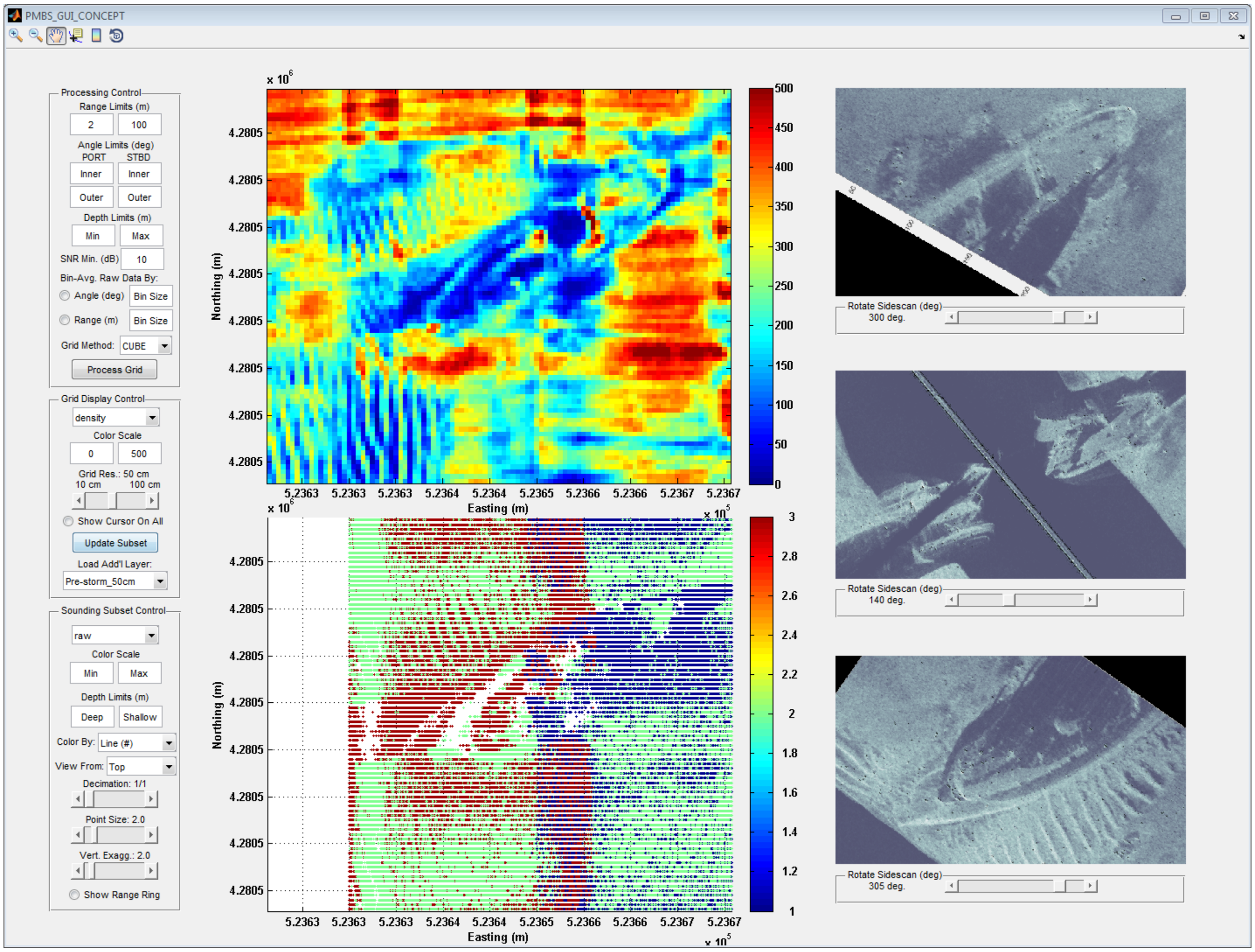

Figure A4. Survey coverage is readily examined by sounding densities in the gridded bathymetry and line coloring in the subset. Regions of low data density (e.g., at nadir, and the northwest faces of the vessel cabin) are evident in the sounding subset and sidescan imagery, increasing confidence in the veracity of this object and indicating limitations of the survey coverage.