6-2003

Automatic Processing of High-Rate, High-Density Multibeam Echosounder Data

Brian R. Calder
University of New Hampshire, Durham, brian.calder@unh.edu

Larry A. Mayer
University of New Hampshire, larry.mayer@unh.edu

Geochemistry Geophysics Geosystems G3
AN ELECTRONIC JOURNAL OF THE EARTH SCIENCES
Published by AGU and the Geochemical Society

Technical Brief
Volume 4, Number 6
11 June 2003
1048, doi:10.1029/2002GC000486
ISSN: 1525-2027

Automatic processing of high-rate, high-density multibeam echosounder data

B. R. Calder and L. A. Mayer
Center for Coastal and Ocean Mapping and Joint Hydrographic Center, University of New Hampshire, 24 Colovos Road, Durham, New Hampshire 03824, USA (brc@ccom.unh.edu; larry.mayer@unh.edu)

[1] Multibeam echosounders (MBES) are currently the best way to determine the bathymetry of large regions of the seabed with high accuracy. They are becoming the standard instrument for hydrographic surveying and are also used in geological studies, mineral exploration and scientific investigation of the earth's crustal deformations and life cycle. The significantly increased data density provided by an MBES has significant advantages in accurately delineating the morphology of the seabed, but comes with the attendant disadvantage of having to handle and process a much greater volume of data. Current data processing approaches typically involve (computer aided) human inspection of all data, with time-consuming and subjective assessment of all data points. As data rates increase with each new generation of instrument and required turn-around times decrease, manual approaches become unwieldy and automatic methods of processing essential. We propose a new method for automatically processing MBES data that attempts to address concerns of efficiency, objectivity, robustness and accuracy. The method attributes each sounding with an estimate of vertical and horizontal error, and then uses a model of information propagation to transfer information about the depth from each sounding to its local neighborhood. Embedded in the survey area are estimation nodes that aim to determine the true depth at an absolutely defined location, along with its associated uncertainty. As soon as soundings are made available, the nodes independently assimilate propagated information to form depth hypotheses which are then tracked and updated on-line as more data is gathered. Consequently, we can extract at any time a "current-best" estimate for all nodes, plus co-located uncertainties and other metrics. The method can assimilate data from multiple surveys, multiple instruments or repeated passes of the same instrument in real-time as data is being gathered. The data assimilation scheme is sufficiently robust to deal with typical survey echosounder errors. Robustness is improved by pre-conditioning the data, and allowing the depth model to be incrementally defined. A model monitoring scheme ensures that inconsistent data are maintained as separate but internally consistent depth hypotheses. A disambiguation of these competing hypotheses is only carried out when required by the user. The algorithm has a low memory footprint, runs faster than data can currently be gathered, and is suitable for real-time use. We call this algorithm CUBE (Combined Uncertainty and Bathymetry Estimator). We illustrate CUBE on two data sets gathered in shallow water with different instruments and for different purposes. We show that the algorithm is robust to even gross failure modes, and reliably processes the vast majority of the data. In both cases, we confirm that the estimates made by CUBE are statistically similar to those generated by hand.

Components: 10,227 words, 13 figures, 1 table.

Keywords: Multibeam echosounders; data processing; bathymetry estimation.

Index Terms: 4536 Oceanography: Physical: Hydrography; 4594 Oceanography: Physical: Instruments and techniques; 3094 Marine Geology and Geophysics: Instruments and techniques.

Received 2 December 2002; **Revised** 14 March 2003; **Accepted** 16 March 2003; **Published** 11 June 2003.

Calder, B. R., and L. A. Mayer, Automatic processing of high-rate, high-density multibeam echosounder data, *Geochem. Geophys. Geosyst.*, *4*(6), 1048, doi:10.1029/2002GC000486, 2003.

1. Introduction

[2] Advances in technology and more exacting requirements for hydrographic and marine geological and geophysical surveys have led to ever increasing data rates and densities for multibeam echosounder (MBES) data sets. These data sets can provide data for constructing nautical charts and, in different depth ranges, provide insight into the processes that are forming and re-forming the earth. From the study of ocean ridge formation [*Grevemeyer et al.*, 2002; *Fox*, 1996] to the formation of continental margins [*Gardner et al.*, 2001a; *Eichhubl et al.*, 2002; *Lee et al.*, 2002] and mapping for ocean exploration [*Dziak et al.*, 2001], baseline maps of the ocean floor provide the geospatial foundations for most ocean studies, and a source of innovation for new investigations [*Jakobsson*, 2002; *Bacon et al.*, 2002].

[3] While our ability to gather bigger and denser data sets has increased dramatically, our ability to process and make sense of these data sets has not. Unlike single beam sonar data sets, the complex geometry and sensor integration associated with multibeam sonars leads to demanding processing requirements. The time and effort involved in manual processing (the current most common method) is increasing in proportion to the increase in data rates. While the hydrographic community is drowning in data from shallow water systems, the oceanographic community is faced with the problem of turning data into information in real time so that logistical and scientific decision can be made during a cruise; clearly, an automatic approach would be beneficial. However, automatic processing leads to questions of reliability, robustness and (in certain domains) safety and liability. In this paper, we propose a new scheme for automatically processing multibeam bathymetry that attempts to answer these concerns. Specifically, we examine the fundamental question of uncertainty in predicting the depth assigned to a particular point. Our three goals are to provide a robust method of processing MBES data, objectively and without human intervention to the stage of a preliminary product; to provide "value added products" to this data set which indicate the expected quality of the data and any locations which require further investigation; and to do so at least as fast as the data can be collected, and preferably in a real-time mode (i.e., where the data is processed as it is gathered, rather than having to wait for all of the data to be available).

[4] There have been a number of approaches to the task of automatically processing high rate bathymetric data. The simplest examples include simple depth and angle gates (i.e., to reject a sounding shallower or deeper than reasonable limits based on a general knowledge of the target area, or from outer beams, where refraction effects are more significant). Slightly more complex examples include filtering based on angles between points, cross-track sounding distance, local gradient, etc., as implemented in the HIPS, MBsystem and SwathEd processing suites [*Gourley and Dodd*, 2000; *Caress and Chayes*, 1995, 1996; *Hughes Clarke et al.*, 1996], with more recent approaches including multiresolution tiling and surface fitting [*Gourley and DesRoches*, 2001]. All of the methods are driven by the need to identify soundings that do not correctly measure the depth (i.e., "outliers") and hence remove them from further consideration. This is primarily a concern of the shallow water community where surveys range on the order of 10^7–10^{10} soundings. Identifying the outliers by hand in this case is typically the most significant time expenditure of any survey by a very significant amount (see, e.g., *Calder and Smith* [2003]) and hence is the primary candidate for process improvement.

[5] Where data are subject to random variation, it is often more appropriate to consider a stochastic description of the problem, and a number of

statistical cleaning techniques have been proposed. *Ware et al.* [1992] suggest a technique to compute mean and standard deviation estimates for entire data sets, and therefore to eliminate some samples based on their deviation from the weighted mean estimated surface. *Varma et al.* [1989] use binned data and an efficient database engine to implement much the same scheme. An alternative technique is to consider hypothesis tests based on pseudo-variance estimates, which (combined with a leave-one-out testing scheme) has been used by *Eeg* [1995] to detect spikes in dense MBES data. All of these schemes rely on an estimate of point statistics of the data in a small area (either geographically, or swath by swath). It is also possible to estimate sounding density, and hence attempt to determine modes corresponding to outliers, from a histogram. *Du et al.* [1996] use this technique to construct an automatic processing scheme that is intended to simulate how a human operator edits data, and is shown to be functionally equivalent to human editing on a small data set. Robust data fitting techniques have also been used in geographic mode [*Debese*, 2001] in order to determine which soundings are consistent within a local area. This method has recently been extended to shallow water data sets [*Debese and Michaux*, 2002].

[6] Most of the automatic editing schemes that have been proposed operate in either swath mode (i.e., causally as the data is collected), or in spatial mode (i.e., geographically after the data has been geo-referenced). One exception to this is the multi-pass filtering of *Lirakis and Bongiovanni* [2000], which is incorporated into the Naval Oceanographic Office's area based editing scheme [*Depner and Hammack*, 1999]. This starts with data in swath mode, processes analogously to many of the filters considered above, and then converts the data into geographic mode for further filtering. This system also adds the concept of a modifiable classification attribute for each sounding, so that a depth can be marked ''Unknown'', ''Good'' or ''Bad'', and many of the tests implemented revolve around transitions between these states.

[7] These techniques all have the same ''data flagging'' paradigm. That is, each sounding is to be preserved, but some may be marked as ''not for use'' through some criteria, automatic or manual. We depart from this standard model by focusing on the essentially contiguous surface that we are trying to measure, i.e., the seabed. The question of interest, then, is not which soundings are subjectively ''good'' or ''bad'' (all of them are noisy to some extent) but how well we can determine the depth at any given point of interest, including the confidence with which we can report the depths estimated.

[8] The advantages in focusing on a surface are threefold. Firstly, working at a point of interest reduces the complexity of estimation; if we are at a fixed point, we should see only one real depth and hence we only have to estimate a constant. This allows us to take advantage of a mature body of estimation theory, and to readily and simply address questions of robustness in the presence of outlier data points. Properly treating suspect data gives us, in essence, a stochastic method for data ''cleaning'' in the sense above, but without the subjectivity. Stochastic processing also allows us to develop estimates of data variability and hence of the reliability of the estimates that are produced. In real-time processing, these value added products can be used in survey planning and data quality assurance.

[9] Secondly, using the concept of a continuous surface, we can take advantage of arguments of required continuity in the depth estimates to discriminate between valid and erroneous data at the level of an estimation node. This emphasis on a 'best estimate' surface simplifies our processing stream, since we do not need to apply multiple levels of hydrographic rounding algorithms (where soundings are rounded to the next shoaler quantum of depth in order to ensure safety of navigation) as we work. We maintain and archive our best estimate, at the best resolution available and/or required, rather than a surface modified multiple times from the original data.

[10] Finally, working on a surface, or grid, of estimation nodes allows us to produce a data product which is more readily manipulated than the common alternatives. For example, a gridded surface is easily manipulated to generate smooth

contours automatically, at any interval, and can be readily prepared for database maintenance. A surface is also ideal for future digital end products, e.g., Electronic Navigational Charts, [*Smith et al.*, 2002].

[11] In the remainder of this paper, we outline the theoretical foundation of the method and illustrate its operation on two data sets gathered with different instruments, in different depth regimes, and with different processing requirements. One of the data sets is a massively over-sampled data set [*Flood et al.*, 2000], which allows us to undertake some statistical comparisons of the surfaces estimated by alternative methods. We use this data set to show that the proposed method is consistent with simpler and more intuitive (but also more memory-hungry and non-real time) algorithms. The second data set is a typical commercially driven survey operating at normal data densities; we use this to investigate the fidelity of CUBE by comparing the algorithmic results against a hand-edited data set. We then conclude with some perspectives on future developments of the method.

2. Theory

[12] We present here an intuitive description of CUBE and its component parts, leaving the mathematical details to the appendix. CUBE is fundamentally about answering the question "what is the true depth at this point, given that all measurements have errors in all three dimensions?", with auxiliary question "How sure are we of that estimate?". The extent to which we can answer this is a function of the level of noise in the data, both stochastic and systematic (for example, due to failures of the MBES to resolve the bottom correctly, poor tide correctors, etc.)

[13] We note that focusing on an attempt to estimate the "true" depth is a significant departure from traditional hydrographic practice, where only an actual sounding is acceptable as a depth measurement for charting. This attitude is rooted in pre-MBES practice, mandated by the issue of safety of navigation: the shoalest sounding should be preserved. In MBES systems, however, that shoalest of all accepted soundings is simply the shoaler tail of the sampling distribution for the MBES; charted soundings are often significantly shoaler than the true depth in the area. Indeed, building a histogram of selected sounding origin as a function of beam number for hydrographic surveys typically shows that the (less reliable) outer beams of the sonar are heavily preferred in the product that is the primary archive of the survey. This significant asymmetry occurs because the outer beams are most affected by noise due to low grazing angle and refraction, etc. They are, therefore, more frequently shoaler than other measurements in the same area, and are preferentially chosen by typical (shoal-biased) hydrographic data reduction algorithms. A statistically justified estimate of true depth allows us to bypass these biases, with the caveat that the depth estimates constructed may have to be adjusted for navigational safety should they be used for hydrographic purposes [*Smith et al.*, 2002].

[14] Our requirement for a "true" depth implies that we have to make our estimate at a point in space, since any area will have depth variation on some scale. Taking a point with absolutely known horizontal position significantly simplifies the estimation process, since at a point there can only be one depth (taking the shoalest in the case of overhangs). If the position is assumed to be known precisely, then the only residual error is in the vertical axis. If we then cover the spatial area of interest with a sufficiently dense network of these estimation nodes, we can estimate the depth over an entire survey area. From the point estimates, we can then compute a continuous surface if required. All of the theory for CUBE is based on a single estimation node, with the understanding that a network of such nodes will be required. CUBE does not place any limitations on the spacing, regularity or location of the estimation nodes, as long as they are sufficiently dense to adequately represent the detail implicit in the surface. Each sounding may contribute information to more than one node, so there is also no loss of intrinsic resolution (as, for example, with a binned estimate). In the worst case, using too many nodes will just lead to inefficiency and sparse estimates (a node will report no depth if no data has been received).

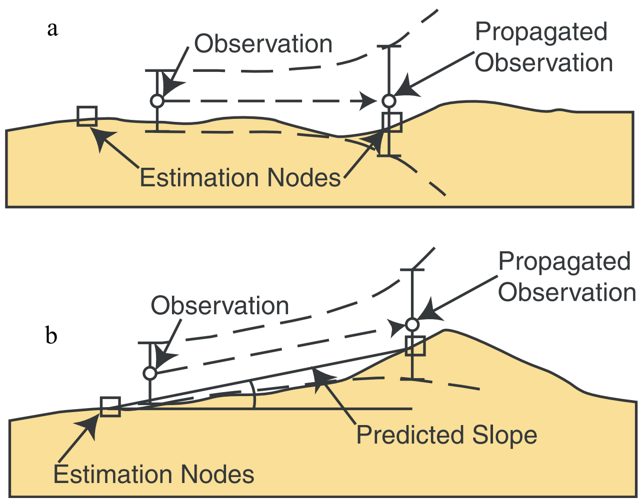

Figure 1. Propagation strategies for referring the sounding data to the estimation nodes. In regions that are essentially flat (a), we might use a zero-order prediction of depth, increasing the uncertainty associated with the data as a function of distance between sounding and estimation node; in regions with significant slope (b), we might make a first-order prediction of depth, with suitably modified confidence bounds.

[15] Around each estimation node, at least locally, we can assume that the soundings will report a noisy but unbiased estimate of the same depth, where we correct for predicted slope if required (Figure 1). However, not all of the data is available at any one time and we have to allow the estimate of depth to evolve as more data is gathered. We formulate the estimation problem by notionally assembling all of the soundings around the node into an ordered, pseudo-time, sequence (although the ordering is essentially arbitrary, e.g., the order in which the soundings are read from file). Our assumption of local unbiasedness implies that the sequence should simply oscillate about a constant value, namely the depth to be estimated (Figure 2). We may then make the system capable of running in real-time by implementing a robust causal sequential estimator of this constant value.

[16] The theory of sequential estimates is encapsulated here by a Bayesian Dynamic Linear Model (DLM) [*West and Harrison*, 1997], configured to estimate a constant value. In this model, only the current estimate of depth is retained at any time; as data is collected, this estimate is updated and the original data is "discarded" (typically, inserted into a database for further analysis and archive). A critical advantage of this scheme is that the uncertainty in the estimate (i.e., the posterior variance of the estimate) is also tracked, which can then be used for confidence checks on the depth track (Figure 2b). Estimates of accuracy for each individual sounding are computed using the model of *Hare et al.* [1995], and depend on detailed knowledge of the survey system, auxiliary sensors, and configuration of the survey platform. In essence, however, the model is an application of the principle of propagation of variance to the fundamental equations for depth computation using an MBES. More complex and complete models of the component that deals with the MBES bottom detection method exist (e.g., [*Lurton*, 2000]), and would allow other factors, e.g., signal to noise ratio (SNR), to be taken into account. However, these are typically not known and the predictions of these more complex models are well represented by the current one as long as the SNR is approximately 10dB or more. If the MBES is going to report a depth at all, this is typically the case, and we have not pursued the added complexity further.

[17] Not all data exhibits solely measurement noise; Figure 3b shows a typical problem, where groups of soundings are present that are mutually inconsistent, but exist in sets that are internally consistent. It is essential that these groups of

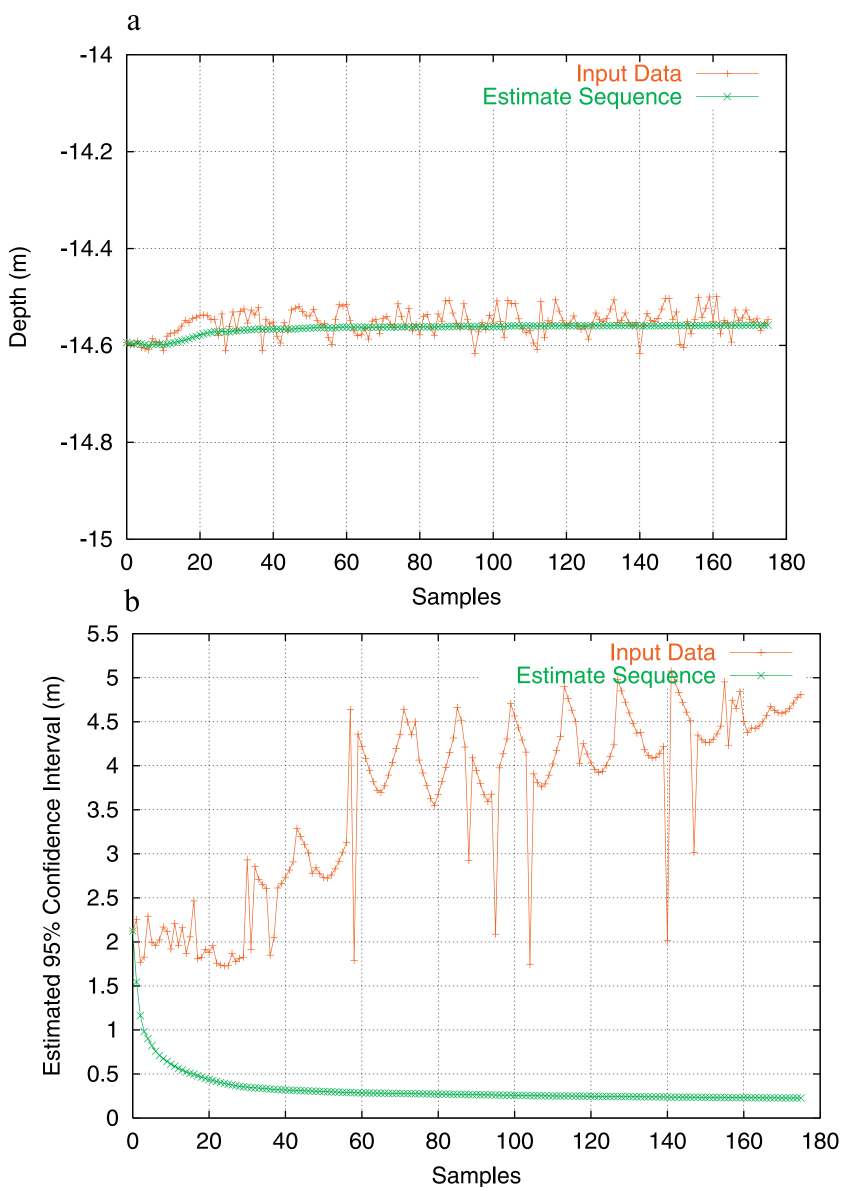

Figure 2. Trace of depth estimate and uncertainty for data in a locally flat area. Note smoothing of the depth estimate and the reduction in uncertainty as the number of samples increases. The input data uncertainty is combined propagated uncertainty $\sigma_j[n]$ (see section A4), scaled assuming a normal distribution, and hence includes both horizontal and vertical uncertainty.

soundings are segregated from each other to avoid cross-contamination of estimates; in the process, we ensure that outliers are not mixed in with "true" soundings. This is a question of model robustness, since these mutually inconsistent groups of data are deviations from the model of constant depth proposed above. Although segregating the burst mode noise (or other outliers) like this may appear to have no intrinsic value for estimating the true depth, it is the key element in making such estimation possible. If we do not remove the outliers by hand (a massively laborious part of any survey and especially so in shallow water environments), then any estimate of depth in this sort of region would be heavily affected by the outliers. Removing the burst data into a sacrificial alternate hypothesis allows us to estimate the true depth and ignore outliers simultaneously.

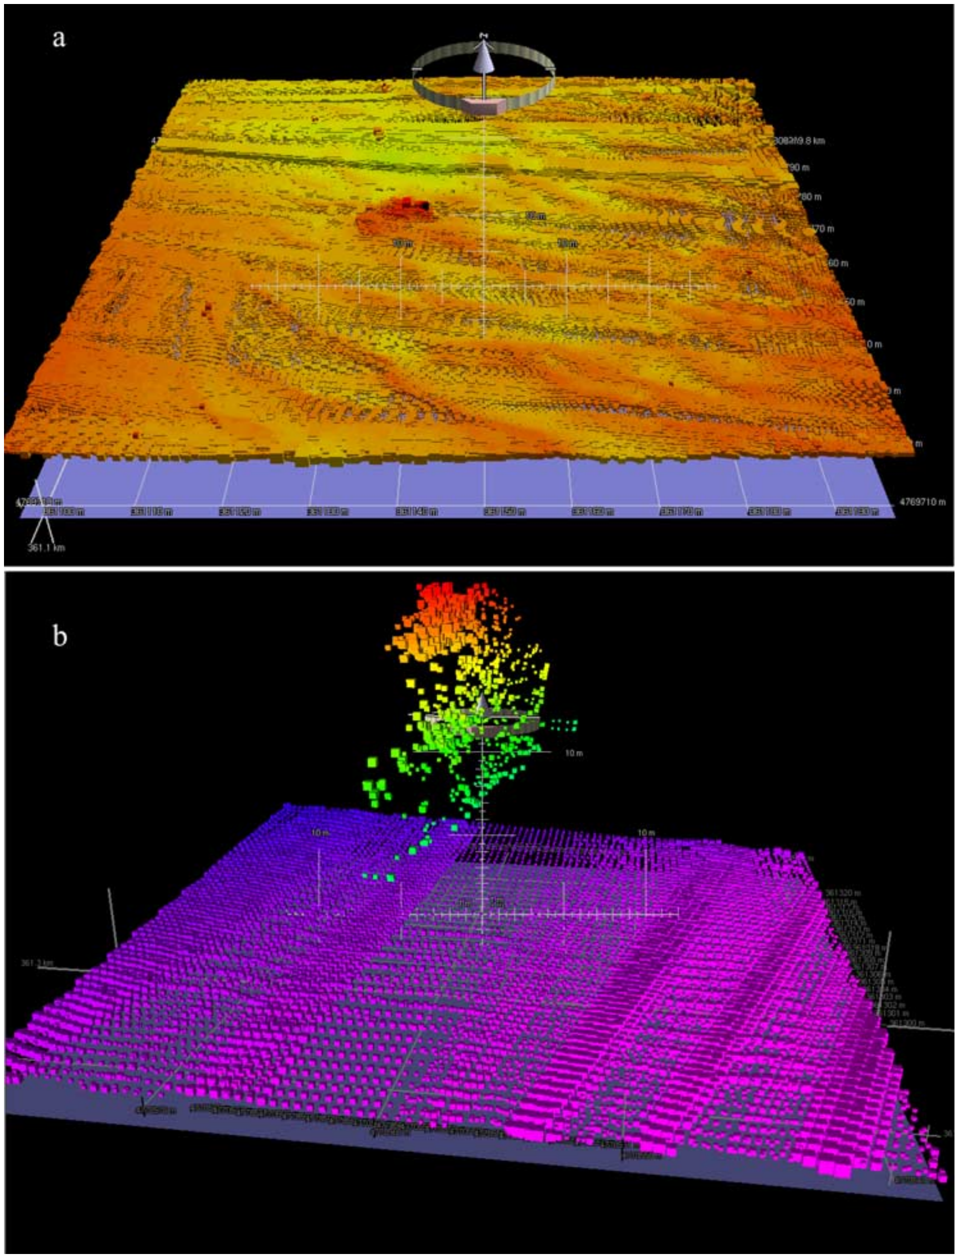

Figure 3. Raw MBES data from a Reson 8101 survey in Portsmouth, NH [*Glang et al.*, 2000]; (a) normal operation and (b) burst-mode failure. Each sounding is represented by a cube scaled to indicate predicted vertical error (see section A3), and color-coded according to depth (hot colors are shallower).

[18] Robustness is implemented by extending the DLM model to allow multiple potential depth solutions to be tracked simultaneously (we call this Multiple Hypothesis Tracking, MHT). As each sounding becomes available, it is compared to each extant hypothesis in turn using a minimum forecast error argument (section A5), and the track to which it is closest is selected for further testing. If the sounding is compatible with the depth estimate in the sense of pointwise and sequential Bayes factors, and given the prediction of uncertainty in the data point and track, then it is assimilated into the depth track. Otherwise, it forms the genesis of a new hypothesis. In this way, we significantly improve the robustness and memory length of the estimator by checking data for validity before

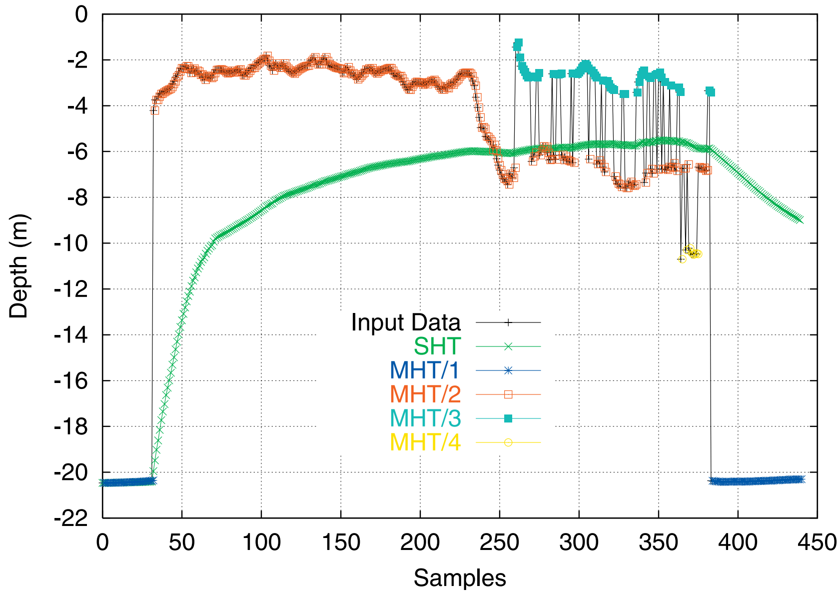

Figure 4. Example of Multiple (MHT) and Single (SHT) Hypothesis Tracking in real data with a burst mode failure; estimates are color coded according to the hypothesis that generated them. Note that the first hypothesis is not corrupted by noise data, and continues to integrate data points when the noise burst disappears around sample 380, and that the SHT track fails to estimate any valid depth. Tracking of data in the sacrificial hypotheses (2–4) representing the noise is significantly different to that of the real data due to differences in statistical properties of the noise.

assimilation and allowing the estimator to backtrack to previous estimates if the data starts to change (Figure 4). The trade-off is slightly increased complexity, memory and computational time.

[19] CUBE's state of knowledge about the data can be summarized through the list of depth hypotheses at each estimation node. Each hypothesis contains an estimate of the depth, the Bayesian posterior variance of this estimate, and a count of the number of soundings that have been assimilated into the track. However, while the extension to multiple hypotheses gives us a significantly more robust estimator, it also introduces ambiguity about the true depth at any node; given a number of potential depths, how do we determine which is the correct one? We resolve this issue using one of a number of disambiguation rules, typically based on a measure of local context (i.e., assuming that the true depth will probably be about the same as those around it). How to best establish this contextual guide is still an open research question, although we are currently using either the closest node with only one hypothesis (i.e., no ambiguity), or a (suitably interpolated) low-resolution estimate of depth (see section A6).

[20] The output of CUBE is thus a set of vectors that represent the algorithm's best estimate of the true depth at each estimation node, the posterior variance of the estimate and the number of potential estimates that exist. A fusion of these outputs is used to inform the user about reliability of the estimates, and problems with the data.

[21] We observe that CUBE cannot be expected to make the correct hypothesis choice in every case: if there is only noise to choose, all choices are in error. Consequently, we wrap CUBE in a human-driven remediation scheme that allows us to investigate the data and estimates in a 3D environment. Currently, we correct those areas where CUBE cannot reliably determine the true depth by flagging erroneous data as "not for use" in the traditional way; a re-run of CUBE assimilates these corrections, resulting in a finalized database. (This may be repeated as more data is added.) This flagging paradigm may not be the best solution; a

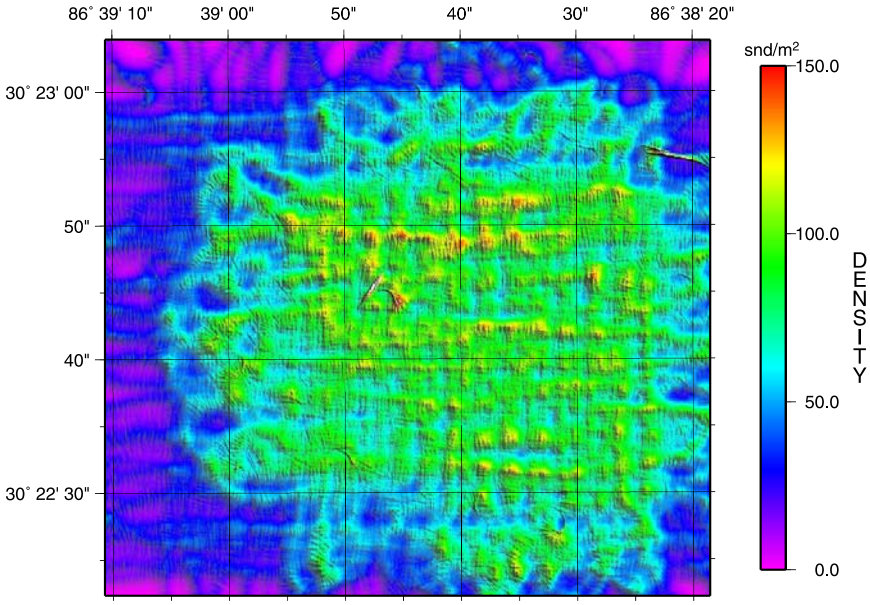

Figure 5. Data density in the SAX'99 high-resolution bathymetric survey center section. Data is colored in soundings/m^2 (white indicates the density is greater than 150 soundings/m^2). Projection: UTM (zone 16N) on WGS-84.

tool that interacts with CUBE's multiple hypothesis space and allows the user to remove, merge, or nominate hypotheses may be more efficient. Such tools are currently in development [*Depner et al.*, 2002].

3. Experiment

3.1. A "Reference" Surface: SAX'99

[22] SAX'99 was a research initiative sponsored by the Office of Naval Research. Mainly concerned with sediment acoustics [*Richardson et al.*, 2001; *Thorsos et al.*, 2001; *Chotiros et al.*, 2001], the data set collected also contained a massively dense multibeam bathymetric data set [*Flood et al.*, 2000] collected in the Gulf of Mexico using a Simrad EM3000 shallow water MBES. The data density is much higher than would be typically collected (Figure 5), with the 1.2 $\times$ 1.2 km study area being surveyed repeatedly over eight days. Data were collected in tracklines running north-south, east-west, northwest-southeast and southwest-northeast, with data being collected at a rate controlled by the MBES controller (typically around 10 Hz).

[23] We constructed two reference surfaces by computing median and iteratively trimmed mean (outliers rejected at 2σ from the mean on each iteration until the estimated mean stabilizes) surfaces in 2 m bins using all of the data gathered and measured tide correctors. All soundings were reduced to mean lower low water (MLLW) and computations were done in UTM coordinates (zone 16 N) on the WGS-84 derived horizontal datum. Visual inspection of the resulting surfaces shows no obvious outliers or other anomalies. To quantify the agreement between mean and median surfaces, we use the International Hydrographic Organization's standards for Order 1 survey, viz. that the 95% confidence interval (CI) for soundings should be $c = 1.96\sigma = \sqrt{0.5^2 + (0.013d)^2}$ m where d is the reported depth. This is the predominately quoted standard for hydrographic survey [*IHO Committee*, 1996], and is used in various adapted forms by all of the U.S. Federal mapping agencies, and hydrographic organizations around the world. The mean and median estimates agree within these limits. Since the mean (Figure 6) and median agree, and the median is known to be robust in noise, we conclude that this is a reasonable estimate of depth in this special case of massively dense data.

[24] We then processed the same data using CUBE, with nodes spaced at 2 m and co-located with the

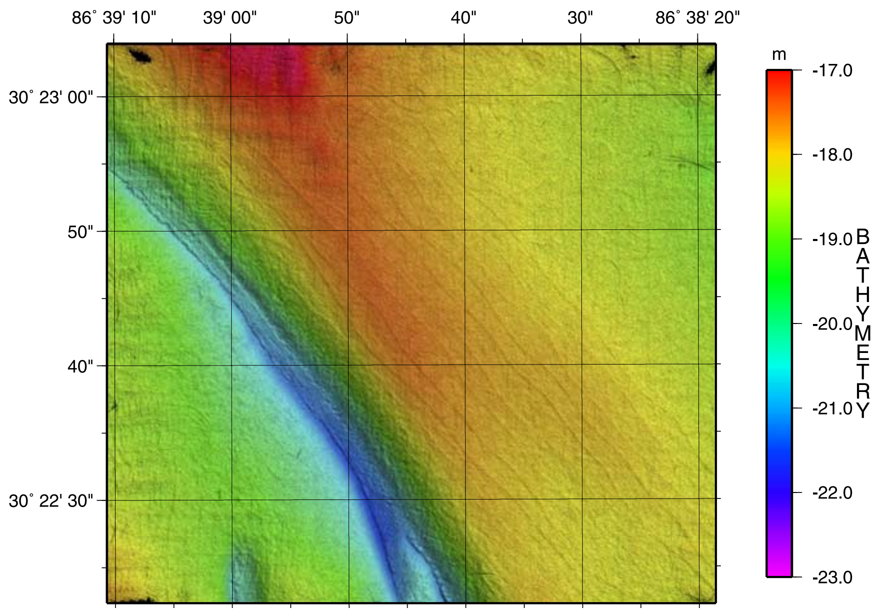

Figure 6. Trimmed mean binned surface estimate for SAX'99. This uses 2 m bins in projected coordinates (UTM on WGS-84), and represents a total of approximately 153 $\times$ 10^6 soundings. Surface is sun illuminated from the northeast. Note small residual artifact in northwest-southeast direction, probably due to residual tide correction errors in lines run in this direction.

centers of the bin estimates. After hypothesis resolution using equation (A13), we computed point-wise differences in depth between CUBE's depth estimate and the mean and median surfaces; the histograms are shown in Figure 7. The difference between surfaces is minimal, with mean absolute difference on the order of a few centimeters. This is both statistically and hydrographically insignificant, and we therefore conclude that CUBE's estimates of depth are hydrographically equivalent to more conventional processing strategies in this case.

[25] The tremendous data density in SAX'99 makes it relatively easy to compute robust estimates of depth; the important question becomes how estimation algorithms degrade as the data density falls to more standard levels. In order to test this, we progressively sub-sampled the data by survey line, and repeated the comparisons above. We found that the errors between CUBE's surface and the reference surface (computed using all of the data) increased, but that the degradation was gradual and remained within IHO limits in all cases down to a mean sounding density of 13.7 soundings/m^2, the lowest density we could construct and still maintain full coverage of the surface. We therefore conclude that CUBE degrades gracefully with decreasing data density.

3.2. Shelf Survey: West Florida

[26] From 3 September until 12 October 2001, a USGS led cruise aboard the R/V Moana Wave mapped segments of the mid-shelf around western Florida with the objective of geological interpretation in support of habitat studies [*Gardner et al.*, 2001b]. The survey system consisted of a Simrad EM1002 MBES, and a POS/MV 320 attitude/position sensor fed with differential GPS from a Satloc MBX-2 receiver. Integration of data, including sound speed profiles, refraction corrections and time-tagging, was implemented in real-time by the Simrad system, although some corrections for latency in navigation were applied in post-processing. Soundings were reduced to MLLW using predicted tides from Panama City (scale factor 0.96, no time shift). Positions were recorded with respect to WGS-84 and computations were done in UTM projected coordinates.

[27] During the cruise, the data was examined by hand and edits were made to flag obviously bad

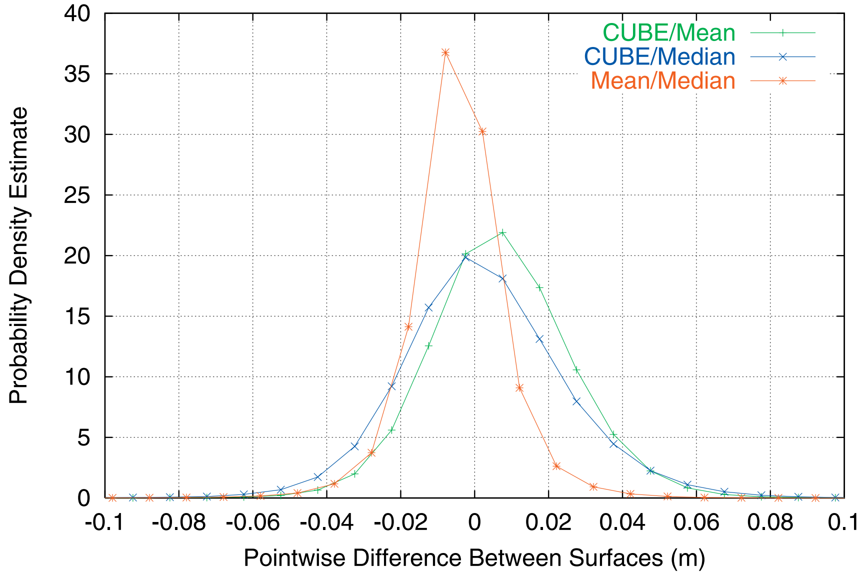

Figure 7. Point-wise difference histograms between surfaces estimated with different algorithms, but using all available data. The mean absolute difference is on the order of a few centimeters in each case, clearly hydrographically and statistically insignificant.

data as not for use. In order to compare the effect of hand edits with CUBE's processing, we ran CUBE twice, once using only the data accepted by the human operators, and once using all of the data collected. This methodology allows us to isolate edit effects from surface generation, since CUBE acts as the surface generator in both cases. Nodes were spaced at 4 m, with depths being in the range 65–140 m; we concentrate on a sub-section of the survey area for conciseness (Figure 8).

[28] Visual inspection of the bathymetry shows some small anomalies in the outer beams of each swath, which are known to be produced by problems with the echosounder mounting position under adverse weather conditions, and which occur in surfaces constructed with and without flags. Pointwise differences of the two surfaces, Figure 9, show remarkable agreement, well within the expected accuracy of the survey, with the most significant errors concentrated in those regions of sparse data known to be of lower quality. Figure 10 shows an example of this type of data, and the surface estimate by CUBE; even where there are gross errors, CUBE successfully maintains surface lock. We therefore conclude that CUBE's estimates are just as good as those generated with the benefit of hand editing, but without the subjective editing or time requirements.

[29] CUBE's auxiliary products, however, provide added benefits. The map of hypothesis counts, Figure 11, clearly shows areas of difficulty, and may be used to improve operator efficiency by concentrating effort where it is required, rather than expending it equally over all of the data. In addition, the map of estimate uncertainties, Figure 12, may be used to judge data quality. Here, it clearly shows three distinct regimes that correspond to three different stages of survey conducted in different sea states. In addition, in the region where conditions were the worst, CUBE shows significant gaps between swaths, indicating regions where the data is very sparse and of poor quality such that the algorithm cannot form a reliable estimate of depth. The default behavior in this case is to leave a blank region. Since both of these products can be constructed as the survey progresses, they could easily be used for survey planning and quality assurance in the field, making it more likely that data of sufficient accuracy and

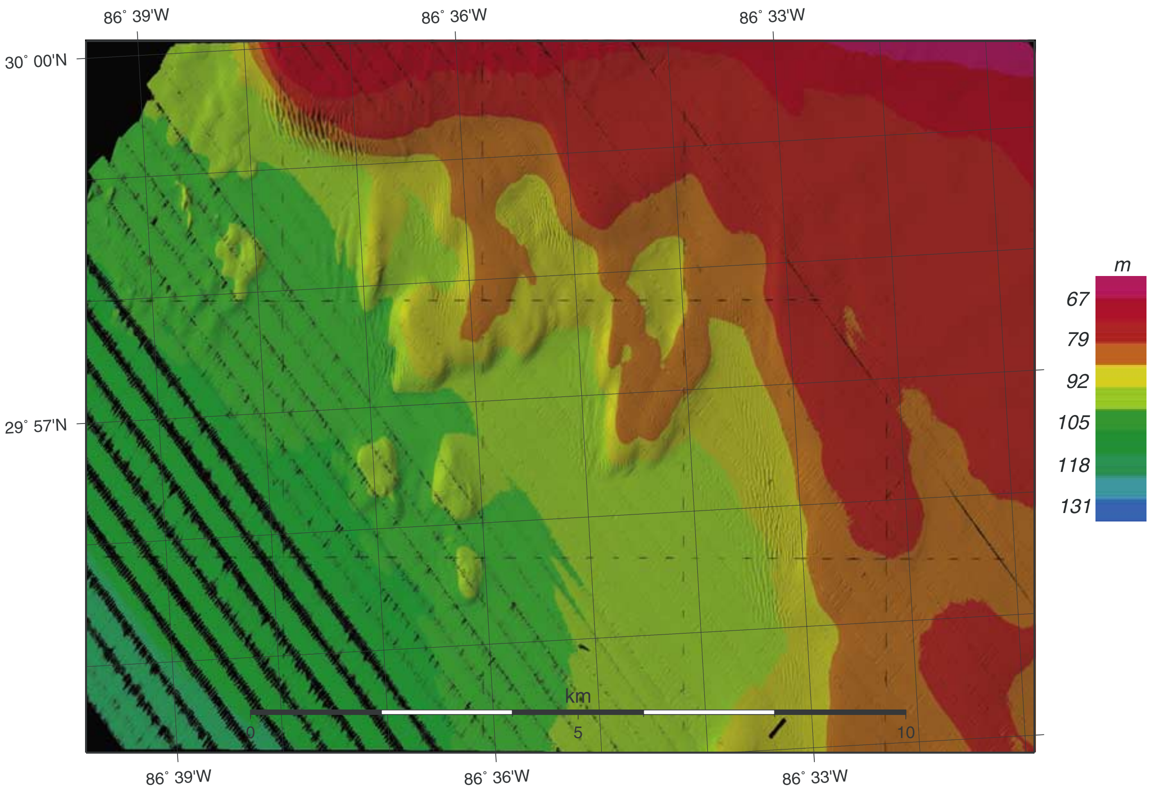

Figure 8. Bathymetric map for West Florida mid-shelf mapping. Automatically constructed by CUBE with no hand edits, nodes at 4 m spacing, and sun illuminated from the north west. Projection: UTM (zone 16N) on WGS-84. Note the track line oriented artifacts, most obviously in the south-west. These are caused by increasing sea-state and attendant reduction in data quality as the survey progressed (c.f., Figures 11–12).

coverage is really "in the can" before leaving the work site.

4. Discussion

[30] Our results show that it is possible to carry out bathymetric data processing, including automatic "cleaning," without having to wait for all of the data to be available, with minimal overhead, and faster than nominal data capture rates. In addition, the robustness of CUBE appears to be sufficient to deal with typical stochastic errors inherent in all surveys, and even some non-stochastic errors such as datum shifts, burst-mode sounder failure and badly ray-traced beams. The use of Multiple Hypothesis Tracking (MHT) to capture mutually inconsistent, but internally self-consistent, sections of data significantly improves robustness, and the additional surfaces resulting from this (particularly the count of hypotheses) give improved insight into the quality of data. The ability of the algorithm to operate in real-time, as the data is being gathered, means that we can use the co-registered products to implement quality assurance in the field.

[31] CUBE may be considered to be a robust weighted surface construction algorithm, although it really forms a collection of point estimates with the surface being implicitly or explicitly constructed from these. However, as an alternative surface construction algorithm, it has much to recommend it. Rather than applying ad hoc weightings, the soundings are individually processed according to their estimated uncertainty, taking account of both horizontal and vertical accuracy. The data assimilation model is intuitive, and readily modified and customized for particular survey conditions; the components of the system are all

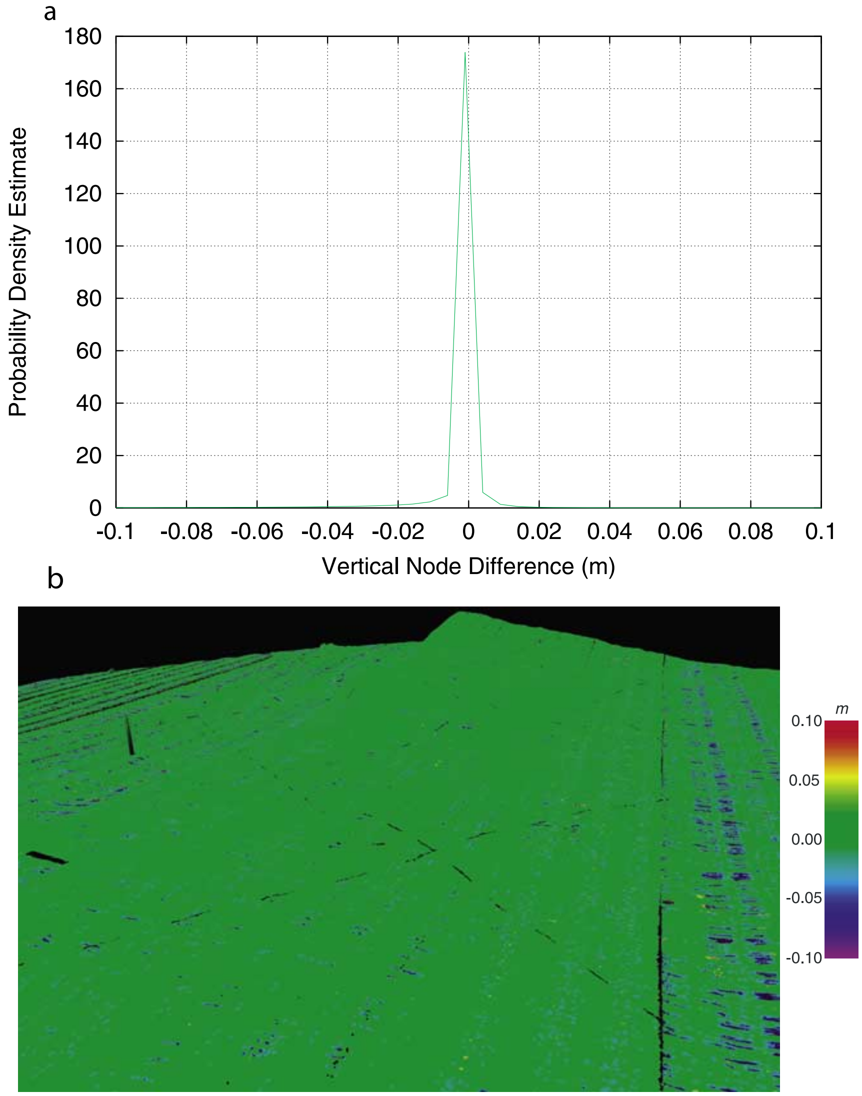

Figure 9. Differences between surfaces generated by CUBE with and without hand edits; (a) point-wide difference histograms and (b) color-coded surface with difference between hand-edited and automatically generated surfaces.

well separated and hence may be readily adapted, modified or replaced as required. Finally, with certain limitations, the algorithm can be used in real-time mode to provide first-pass field quality control and estimates.

[32] In practice, the balance between depth prediction and error propagation may depend on the purpose for which the data is collected. In a strict hydrographic processing chain we might wish to be more conservative, shoal biasing all predicted bathymetry, and making the error bounds increase more rapidly. In a geoscience context, we may try to ensure that small features are correctly represented, even if that means some higher noise levels elsewhere that will require more interaction to correct. In some contexts, we might need to consider areas where our zero-order prediction is no longer valid, and hence incorporate some interaction terms between neighboring estimation nodes to compensate. At present, we specify the propagation terms a priori with at best a nod to a Bayesian subjective prior argument for their values. This is a small (but unavoidable) weakness in the algorithm.

[33] We focus on stochastic uncertainty here, rather than any systematic effects (e.g., an incorrect align-

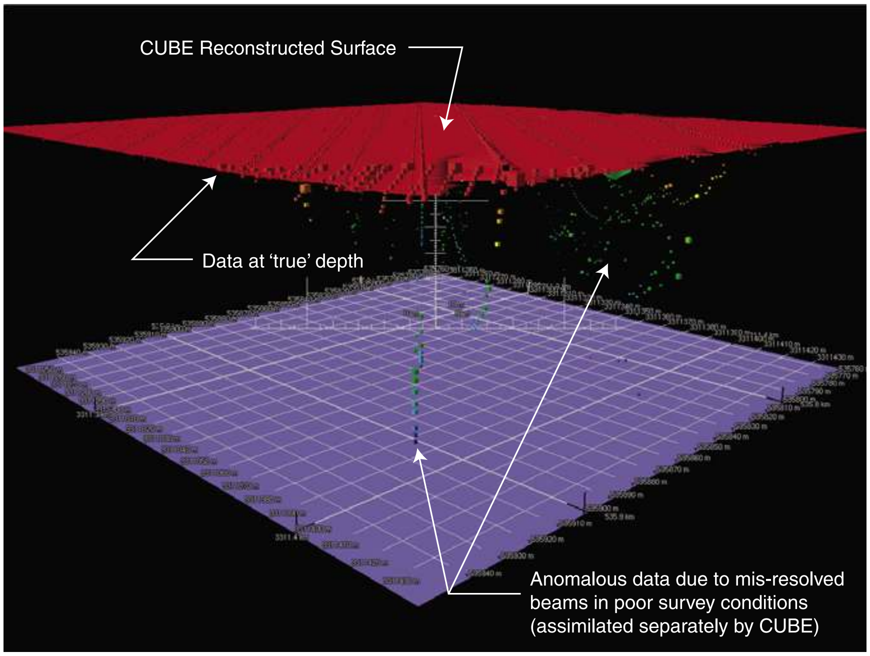

Figure 10. Raw data, with CUBE's estimated surface. Soundings are represented as cubes, color-coded by depth (hot colors are shallower) and sized by predicted vertical uncertainty. Note that even in very significant noise, the surface lock is still maintained.

ment bias term), holding that systematic effects should be taken care of by normal best practice. Although CUBE does not account for them, it is a very effective tool for illustrating the effects, since the soundings from multiple passes with systematic problems will disagree on depth in a manner that is a function of the problem. If the difference is significant then analysis of the pattern of hypotheses can be used as a diagnostic tool. This is in keeping with CUBE's fundamental philosophy of telling the truth about the data, in as much as it is known (e.g., we do not report a depth if no sounding is close enough, and report uncertainties, etc.) We feel it is better to show the 'warts' in the data during the processing stage, and resolve or hide them (as the application demands) at a later stage of the process.

[34] For the hydrographic community, there is a significant paradigm shift required to consider statistically derived surfaces as the most valid description of the data, rather than the more traditional estimates, e.g., a shoal biased selected sounding set. However, we would contend that a significant proportion of the data observed in such conventional estimates represents the upper tails of the sampling distribution of a MBES, and hence are not at all representative of the actual bathymetry. A reliance on such data is probably rooted in the use of physical sounding methods, where if a lead-line showed up shallower within an area than other points, there must have been something there. With indirect sounding methods, this is not necessarily the case.

[35] The addition of MHT and the hypothesis count surface, in addition to the uncertainty surface strengthen the arguments for automatic processing, at least as a preliminary processing stage. We do not believe that any automatic data processing algorithm will compensate for all errors in data, stochastic or otherwise, and that manual data examination may be required when the uncertainties warrant it. However, our approach should provide reasonable estimates for the depth in most places, and hence reduce the processing burden on the operators by focusing attention and effort where it is required. The combination of fast automatic processing in almost all areas and computer directed human decisions (informed by auxiliary data surfaces) should significantly increase processing efficiency.

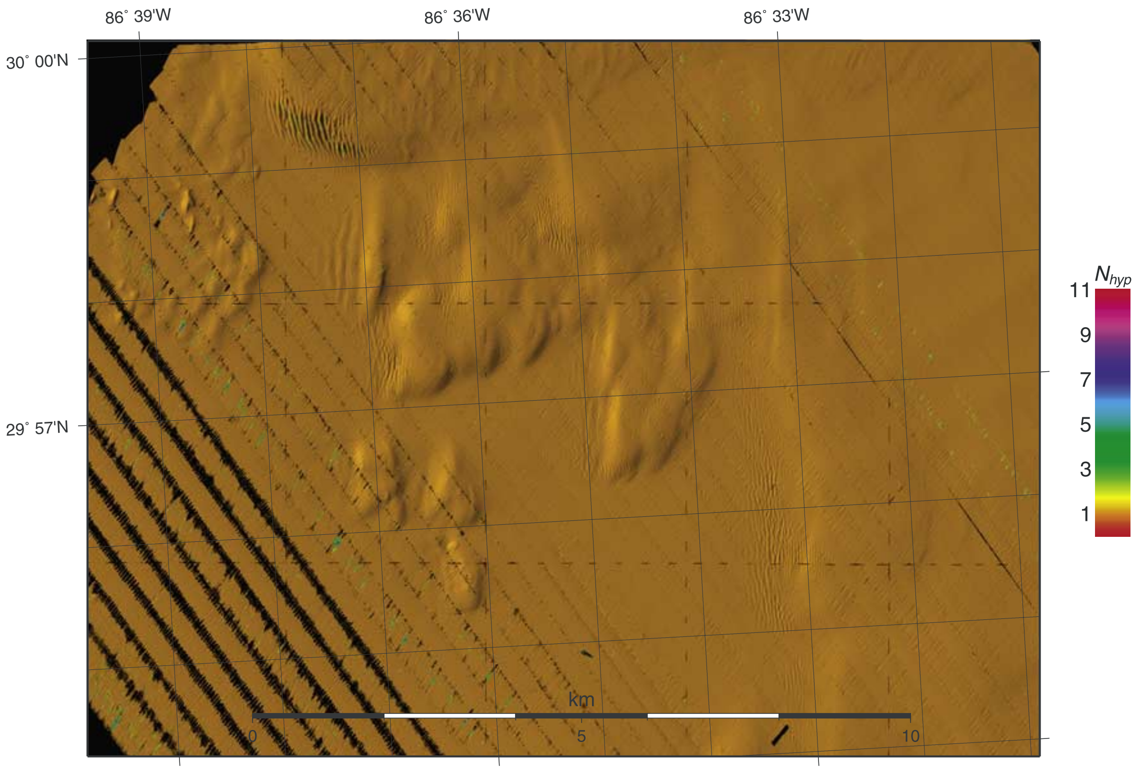

Figure 11. Count of CUBE hypotheses color-coded over CUBE's surface estimate, sun illumination from the northwest. The majority of the survey area is consistently estimated with a single hypothesis, although multiple hypotheses are often seen on outer beam regions. The southwest region is more variable due to reduced data quality caused by worsening sea conditions.

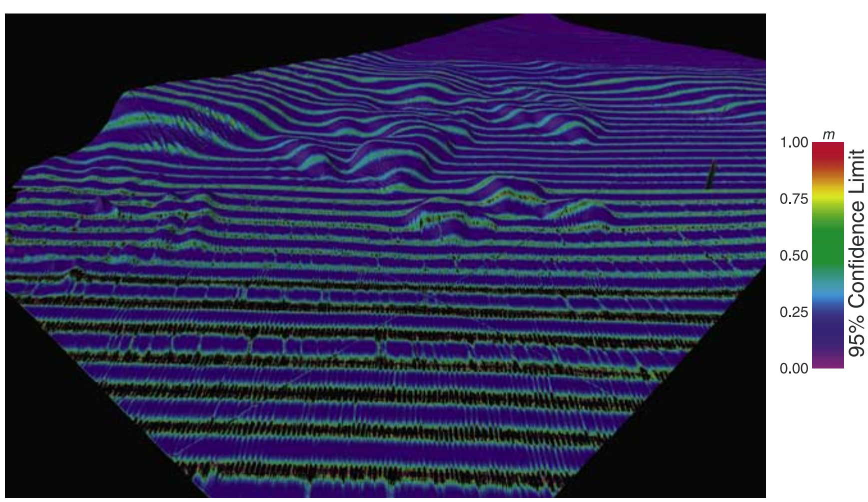

Figure 12. Predicted 95% estimate uncertainty color-coded over CUBE's surface estimate. Three regimes are evident, related to time of data collection and associated sea state. The foreground region was collected under adverse conditions; the outer swath region is of low data density and high predicted uncertainty. Consequently, CUBE declines to generate estimates in these regions, leaving data holidays.

[36] There are a number of directions for further investigation in implementing the algorithm outlined here, particularly with respect to robustness. In real-time mode, we are somewhat limited by the data we can obtain; for example, tide corrections will normally be based on predicted values rather than those measured, and we must work with data as it arrives. However, in post-processing mode, we can make use of all of the data, so that we may, for example, construct surface representations at different resolutions, pushing spatial context from low resolution to high resolution versions. In some environments, the idea of a static depth estimate at a particular point may not be valid under any circumstances (e.g., where there is significant tidally driven sediment dynamics, or where data sets are gathered over a significant period of time). In that case, we would have to implement a full dynamic model, although specification of model dynamics would be complex. Further input from manufacturers of sonar equipment to define and refine error models for particular multibeam (and single-beam) equipment would also be beneficial.

[37] Finally, although the method relaxes the constraint of estimation in a regular grid, we have not pursued the difficulties and benefits of allowing a more general graph of nodes to be specified. In fact, there is no reason for the nodes to be fixed in place (our only assumption is that we know exactly where they are). One possible extension would be to allow nodes to be added or deleted dynamically, to allow us to match the estimated sampling density required for the surface under consideration. Another possibility is to allow the nodes to gravitate to where they are actually required on the surface, rather than having them fixed in place. If the total number of nodes were constrained to be a constant, this would not have a significant computational overhead, although it would require significantly higher investment in bookkeeping.

5. Conclusions

[38] We have proposed a new method for automatic handling of dense bathymetric data. Based on a simple model of the stochastic errors associated with estimates of depth, we have constructed an optimal estimator which tracks an estimate of the depth, the uncertainty associated with that estimate, and the number of self-consistent depths detected at each location. Where multiple self-consistent depth estimates are present, the algorithm selects a "best" estimate according to user-defined rules on what constitutes "best." The method has a number of advantages to recommend it, in particular that it automatically incorporates estimates of 3D uncertainty, allows for optimal combination of estimates, and provides a procedure for updating older surveys with new data (or using older surveys to constrain interpretation of new data).

[39] The method also allows us to build a low memory overhead, real-time, on-line estimator so that we have continually available (as data is added), current best estimates of depth and the associated uncertainty at that depth.

[40] Our results have shown that the estimates generated by CUBE are consistent with other robust estimators and with human-edited data, and hence that in the vast majority of cases we can avoid having to examine every sounding gathered. Our method also provides the means to automate and integrate the human decision making process by linking depth and auxiliary data surfaces.

Appendix A: Theoretical Foundations

A1. Notation

[41] In the following descriptions, bold letters indicate vectors (e.g., $\mathbf{x}$ with transpose $\mathbf{x}^T$); hollow letters are standard sets (e.g., $\mathbb{R}$ for the reals, $\mathbb{N}$ for the naturals, and $\mathbb{Z}$ for the integers), and sans serif letters indicate general sets. $|\mathsf{N}|$ is the cardinality of a set (i.e., the number of members), and $\|\mathbf{x}\|$ represents the Euclidean norm. In the time series description of CUBE's DLM s (Dynamic Linear Models), we use Bayesian conditional notation and square brackets to indicate a discrete time series; hence $\hat{x}[n|n-1]$ is an estimate of x at sample n given information to sample $n-1$, and $\hat{x}[n|n]$ is

the updated estimate including all information to sample n, etc. For random variables, the swung dash indicates a distribution, so $x \sim \mathcal{N}(0, \sigma^2)$ indicates that x is a random variable distributed as a normal (Gaussian) variable with mean 0 (units) and variance σ^2 (units^2).

A2. Foundation

[42] Our task is to estimate depth at a particular geographical location given the sounding data in the immediate vicinity. To quantify our estimate of depth we must also provide an estimate of the uncertainty in the depth. Let $\mathsf{N} = \{\mathbf{n}_j \in \mathbb{R}^2, j \in \mathbb{N}\}$ be the set of locations of estimation nodes, and let $\boldsymbol{\xi}_j = (d_j, \sigma_j^2)^T$ be the vector of depth estimate and uncertainty at the jth node. We have essentially two problems: how to predict, at a node, the depth and uncertainty implied by a sounding even if that sounding is not itself at that location, and how to make the estimation sufficiently robust to deal with MBES failures.

A3. Data Preprocessing

[43] Assume a priori that the ith sounding is valid, and the MBES reports depth ζ_i at nominal location $\mathbf{x}_i \in \mathbb{R}^2$. That is, assume that the depth is correctly detected from the primary seabed reflection through the main lobe of the sonar, and is correctly processed for systematic offsets, refraction and dynamic effects such as platform attitude and dynamic draft. Then, a propagation of variance argument can be applied to the depth solution to estimate, given the measurement accuracies of all of the components of the solution (see, e.g. Table A1), the expected horizontal and vertical errors associated with the nominal sounding ([*Hare et al.*, 1995]). Let $\mathbf{s}_i = (\zeta_i, \sigma^2_{H,i}, \sigma^2_{V,i})^T$ be the vector of computed depth, and horizontal and vertical error variances. A fundamental assumption of our method is that data for a particular location should agree on depth to within these limits. Conversely, if a data point does not agree with our current estimate given this leeway, then we may conclude that we are seeing inconsistent data and take suitable steps. Uncertainties also play a vital role in assimilating data with extant estimates (see section A4).

Table A1. Components of the MBES Error Model[a]

Component	Typ. Value	Units
Spatial offsets between equipment	5.0	mm
Angular alignment of equipment	0.05	deg
GPS positioning (drms)	1.0	m
GPS latency	5.0	ms
Roll/pitch measurement	0.05	deg
Yaw measurement	0.06	deg
IMU latency	5.0	ms
Sound Speed Profile measurement	0.5	m/s
Surface Sound Speed measurement	0.50	m/s
Heave measurement (fixed)	0.05	m
Heave measurement (variable)	5	%
Draft measurement	0.02	m
Dynamic draft	0.02	m
Loading	0.01	m
Speed-over-ground	0.2	m/s
Tide measurement	0.02	m
Tide spatial variation	0.02	m

[a] All values are typical, and are quoted at one standard deviation. Note that these values are not the magnitude of the corrections applied, but the magnitude of the residual after the appropriate corrections (for tide, draft, dynamic draft, etc.) have been made.

[44] While our estimation nodes are fixed in location with respect to some (usually projected) coordinate system, the data soundings are essentially randomly distributed with respect to the nodes (Figure 1). However, continuity and local smoothness of the surface imply that soundings close to a node provide information about the depth there, and we may express this by a function describing the local bathymetric surface, and hence form a prediction of $\boldsymbol{\xi}_j$ given $\mathbf{s}_i$. Let $\delta_{ij} = \|\mathbf{x}_i - \mathbf{n}_j\|$ be the propagation distance. The simplest model is:

$$\mathbf{e}_j(\mathbf{s}_i) = \begin{bmatrix} d_{ij} \\ \sigma^2_{ij} \end{bmatrix} = \begin{bmatrix} \zeta_i \\ \sigma^2_{V,i}\left(1 + \left[\frac{\delta_{ij} + s_H \sigma_{H,i}}{\Delta_{\min}}\right]^{\alpha}\right) \end{bmatrix} \quad \text{(A1)}$$

where s_H is a scale factor for worst expected horizontal error (typically $s_H = 1.96$), $\Delta_{\min}$ is minimum node spacing and α is a user specified exponent (typically $\alpha = 2.0$). The increase of uncertainty with distance δ_{ij}, including the effect of horizontal error, is a reflection of the belief that soundings that are further from the node should be given less credence, as should those with higher initial horizontal or vertical uncertainty. This model also combines separate horizontal and vertical uncertainty into a unified effect by assuming that the sounding could be as much as $s_H \sigma_{H,i}$ m from the nominal location, in the wrong direction.

A4. Data Assimilation

[45] Our ideal model of estimating a constant depth, z, is codified in the Dynamic Linear Model (DLM) [*West and Harrison*, 1997]:

$$z[n+1] = z[n] + w[n] \quad w[n] \sim \mathcal{N}(0, W[n]) \tag{A2}$$

$$d[n] = z[n] + v[n] \quad v[n] \sim \mathcal{N}\left(0, \sigma_j^2[n]\right) \tag{A3}$$

where $w[n]$ and $v[n]$ represent evolution and measurement noise variance respectively. Our model assumes a constant depth, and hence we set $W[n] = 0$ m^2 $\forall n$. To maintain local dependences we select only soundings $\mathsf{S}_j = \{\mathbf{s}_i : \delta_{ij} \leq \Delta_{\max}(i)\}$ for assimilation at node j, where $\Delta_{\max(i)}$ is given by:

$$\Delta_{\max}(i) = \max\left\{\Delta_{\min}, \Delta_{\min} \sqrt[\alpha]{\sigma_{\max}^2/\sigma_{V,i}^2 - 1} - s_H \sigma_{H,i}\right\} \tag{A4}$$

and σ_{max} is the maximum allowable sounding accuracy for the survey in question, computed according to IHO limits ([*IHO Committee*, 1996]).

[46] Let $\mathsf{E}_j = \{\mathbf{e}_j(\mathbf{s}_i) : \mathbf{s}_i \in \mathsf{S}_j\} = \{\mathbf{e}_j[0], \ldots \mathbf{e}_j[N_j - 1]\}$ (where $N_j = |\mathsf{S}_j|$) be the set of propagated soundings associated with the jth node, enumerated by a fixed but arbitrary permutation (e.g., the order in which they arrive from the data stream). We may consider the soundings as a sequence of estimates arriving at a node, which may in turn be regarded as a pseudo-time sequence, $\mathbf{e}_j[n] = (d_j[n], \sigma_j^2[n])^T$, $0 \leq n < N_j$. At the node, the current estimate $\xi_j[n] = \left(\hat{z}_j[n|n], \hat{\sigma}_j^2[n|n]\right)^T$ may then be updated optimally with new data by the sequence:

$$\hat{\sigma}_j^2[n|n-1] = \hat{\sigma}_j^2[n-1|n-1] \tag{A5}$$

$$\hat{z}_j[n|n-1] = \hat{z}_j[n-1|n-1] \tag{A6}$$

$$G_j[n] = \frac{\hat{\sigma}_j^2[n|n-1]}{\hat{\sigma}_j^2[n|n-1] + \sigma_j^2[n]} \tag{A7}$$

$$\epsilon_j[n] = d_j[n] - \hat{z}_j[n|n-1] \tag{A8}$$

$$\hat{z}_j[n|n] = \hat{z}_j[n|n-1] + G[n]\epsilon_j[n] \tag{A9}$$

$$\hat{\sigma}_j^2[n|n] = G[n]\sigma_j^2[n] \tag{A10}$$

which is equivalent to a simple Kalman filter [*Maybeck*, 1979; *Haykin*, 1995]. This recursive solution of the optimal estimation problem means that we need only hold a current estimate of depth and uncertainty rather than having to database and manipulate all of the data for the survey at one time. This formulation is the basis of our real-time construction. Since the series $\mathbf{e}_j[n]$ comes from the immediate spatial vicinity of $\mathbf{n}_j$, we can interpret equations (A5)–(A10), and particularly the weight $G_j[n]$, either as an area-based "editing" scheme, or as a weighted gridding algorithm where the weights are derived from propagated, combined, error uncertainties. A typical pseudo-time sequence in a locally flat area, and its track depth and uncertainty, are shown in Figure 2.

[47] This model implicitly assumes that the observations taken from different beams in each ping and from ping to ping are independent of each other. Of course, this is not exactly the case since, for example, all of the beams are traced against the same sound speed profile, and are reduced to datum with the same static and dynamic draft measurements. At best, the data are conditionally independent. A more realistic error description could be introduced that accounted for these dependencies at the expense of a much more complex modeling environment. When interpreting CUBE's estimates, it must be kept in mind that an extra error over and above the estimate uncertainty may have to be included. Recall that the uncertainty being reported is an a posteriori distribution variance for the current depth estimate, and not a measure of the standard deviation of the input data.

A5. Model Monitoring and Robustness

[48] All acoustically derived hydrographic data suffers to some extent from problems caused by stray acoustic energy, poor bottom tracking or multiple reflections. To account for these in a robust manner, we must extend the model to provide a "judicious and grudging elaboration of the model to ensure against particular hazards" [*Box*, 1980]. Since the

ordering of data points in $\mathbf{e}_j[n]$ is arbitrary, we immediately improve robustness by passing the data through another permutation that implements a moving median filter [*Cormen et al.*, 1990]. This ensures that any potential outliers are delayed in the filter's queue structure, protecting the DLM as it "learns" about the true depth. The strength of protection is proportional to the length of the median filter, but so is the latency between sample arrival and assimilation. We have found that 11 samples is normally effective. Note that this does not imply that CUBE requires 11 samples to make an estimate. It is possible to "flush" the queue into the estimator proper if required, so that hypotheses are formed before reconstruction. A single sounding is sufficient to generate a hypothesis, although having a small number (e.g., 3–5) within range of the node aids significantly in robustness of the estimation. This does not imply a reduction of achievable resolution, since each sounding may be used at more than one node.

[49] This extension is insufficient when there is significant evidence for more than one depth at any node—essentially a circumstance the model of equations (A2)–(A3) does not allow. We extend the model by allowing more than one potential DLM to be present at any node, indexing the one in use at "time" n by an indicator variable $\theta_j[n] \in \mathbb{Z}^+$, $0 \le n < N_j$. This piecewise description is similar to many fully Bayesian approaches to the problem [*Gerlach et al.*, 2000; *Gamerman*, 1998; *Chib*, 1998; *Richardson and Green*, 1997] except that we continue to work recursively. We initialize each node with no DLMs, adding them as incoming data appears to be inconsistent with all current DLMs in the manner described below.

[50] Let there be $\Theta_j[n-1]$ possible models $\boldsymbol{\xi}_j^{[k]}[n]$, $1 \le k \le \Theta_j[n-1]$ immediately prior to using sample n, where each hypothesis has assimilated $n_j^{[k]}$ samples, so $\Sigma_k n_j^{[k]} = N_j$. Then, we choose the best available model by a minimal forecast error argument:

$$\theta_j[n] = \underset{1 \le k \le \Theta_j[n-1]}{\text{argmin}} \left\{ \left| \left(d_j[n] - \hat{z}_j^{[k]}\left[n_j^{[k]}+1 \middle| n_j^{[k]}\right] \right) \middle/ \sqrt{\hat{\sigma}_j^{2\,[k]}\left[n_j^{[k]}+1 \middle| n_j^{[k]}\right] + \sigma_j^2[n]} \right| \right\} \tag{A11}$$

We then check that $\mathbf{e}_j[n]$ is compatible with this model by one-step and sequential Bayes Factor analysis [*West and Harrison*, 1997, chap. 4]. Our alternative hypothesized model is one with at least a four standard deviation step change in depth (measured with reference to the one step forecast error), and we require a log Bayes factor of $\ln B_j[n] \le -2$ before deciding to reject the null hypothesis that the data and model are compatible. If we accept the null hypothesis, $\mathbf{e}_j[n]$ is assimilated into model $\theta_j[n]$; otherwise, it is used to form the start of a new model track. An example of MHT is shown in Figure 4 along with the track which results from integrating all data into a single track.

A6. Hypothesis Resolution

[51] CUBE's remit is to estimate depths automatically. Where only one hypothesis exists, we reconstruct the surface with $\boldsymbol{\xi}_j[n] = \boldsymbol{\xi}_j^{[1]}[n]$; where we track multiple hypotheses, we must have a method to determine which estimate, $\boldsymbol{\xi}_j[n] = \boldsymbol{\xi}_j^{[k]}[n]$ for some k, we think is most likely. Making this choice correctly is the key to minimizing operator intervention time, since the operator only has to be involved if CUBE has multiple hypotheses, but cannot reliably determine which one to choose.

[52] The simplest solution is to choose by number of samples assimilated (i.e., the longest held hypothesis). Thus we choose hypothesis:

$$k = \underset{1 \le i \le \Theta_j[n]}{\text{argmax}} \left\{ n_j^{[i]} / N_j \right\} \tag{A12}$$

Since $\Sigma_i n_j^{[i]} = N_j$ this may be interpreted as an approximation to the posterior model probability. In most cases, this is sufficiently robust to deal with noise. However, in burst mode failure this may not be the case since the "noise" data may occur more frequently than the "true" data.

[53] To improve matters, we utilize local contextual information on depth, and assume that a single hypothesis node is more reliable than one with multiple hypotheses. Let $g_j = \text{argmin}_k \{\delta_{jk} : r_0 \le \delta_{jk} \le r_1, \Theta_k[n] = 1\}$ be the closest such node to the jth node, with estimate $\boldsymbol{\xi}_{g_j}[n]$. (We use $(r_0, r_1) = (5, 10)$ m; the minimum radius improves reconstruction in burst noise.) A surface continuity argument

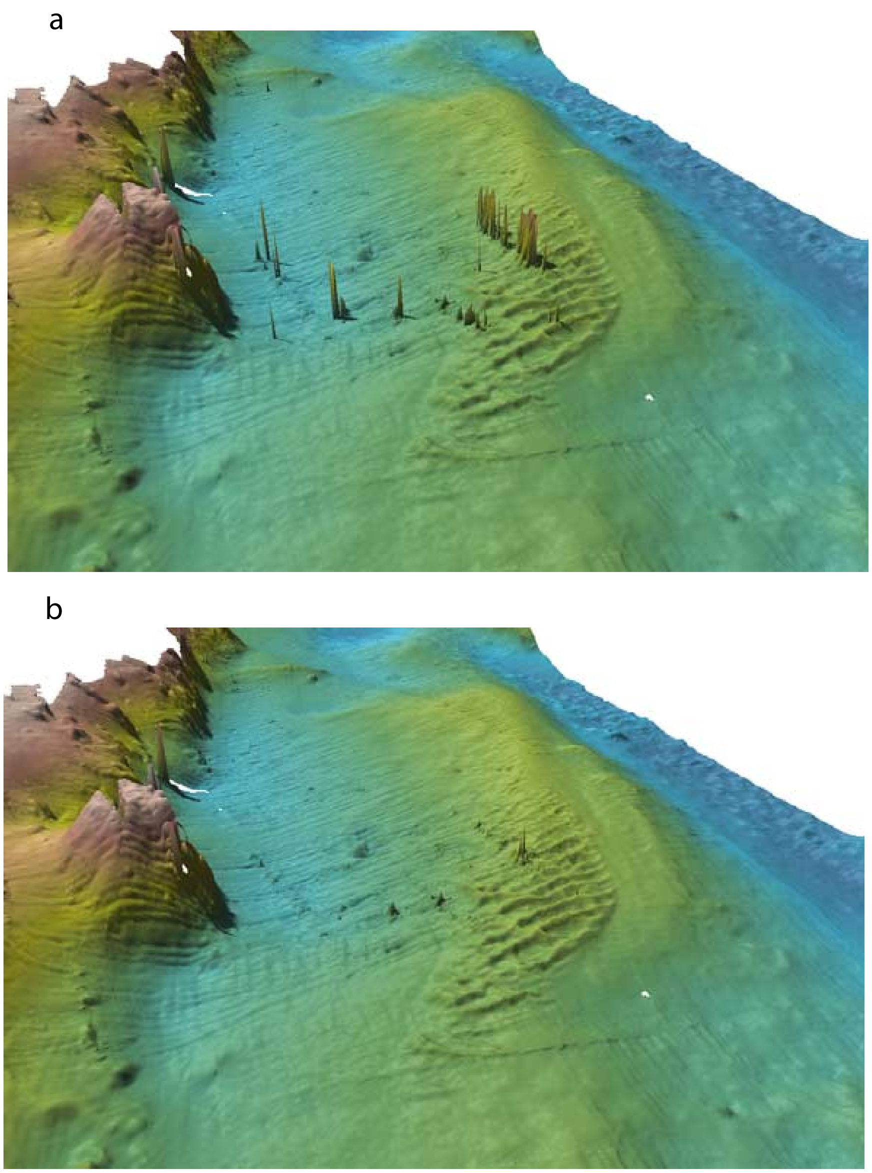

Figure A1. Example of different hypothesis resolution methods. (left) Simple resolution using the number of samples assimilated as a measure of posterior model probability; (right) context-driven hypothesis resolution approximating likelihood of hypothesis given the best certain local node reconstruction. Note that the addition of context significantly improves the number of correctly resolved multiple-hypothesis nodes in the presence of significant noise.

suggests that the depth at the jth node should be similar to that at the g_jth. Hence we choose the hypothesis closest in depth to $\xi_{g_j}[n]$:

$$k = \underset{1 \le i \le \Theta_j[n]}{\operatorname{argmin}} \left\{ \left| \left(\hat{z}_j^{[i]}\left[n_j^{[i]} \middle| n_j^{[i]} \right] - d_{g_j} \right) . \middle/ \sqrt{\hat{\sigma}_j^{2[i]}\left[n_j^{[i]} | n_j^{[i]} \right]} \right| \right\} \tag{A13}$$

[54] Finally, we can combine these two approaches by likening equation (A12) to a prior distribution, and equation (A13) to a likelihood. A pseudo-Bayesian argument then suggests a log "posterior," and selection with:

$$k = \underset{1 \le i \le \Theta_j[n]}{\operatorname{argmax}} \left\{ -\ln \hat{\sigma}_j^{[i]}\left[n_j^{[i]} \middle| n_j^{[i]} \right] - \left(\hat{z}_j^{[i]}\left[n_j^{[i]} \middle| n_j^{[i]} \right] - d_{g_j} \right) \middle/ 2\hat{\sigma}_j^{2[i]}\left[n_j^{[i]} \middle| n_j^{[i]} \right] + \ln n_j^{[i]} \right\} \tag{A14}$$

after some simplification.

[55] Which method to use for hypothesis resolution depends on available time and complexity of data. Determining g_j is costly because of the spatial search required, although the reconstruction using spatial context is typically better. An example from the multibeam survey of Figure 3 in Portsmouth Harbor [*Glang et al.*, 2000] is shown in Figure A1. The data here was particularly noisy due to bubble entrainment from an auxiliary instrument, leading to significant numbers of burst-mode errors. The simple resolution method of equation (A12) has significant difficulties in this case since many of the nodes have more ''noise'' than ''signal.'' However, the context-driven resolution method of equation (A13) improves performance significantly. Note that choosing a ''best'' hypothesis is dubious in any case where there is evidence sufficient to construct multiple hypotheses, and the choice may change as new data is gathered and assimilated. Multiple hypotheses are not necessarily an indication of algorithm failure, merely that an outlier of some significance has occurred; there is only a problem if CUBE cannot determine which hypothesis to choose. Therefore it is essential to interpret the reconstructed surface with reference to the hypothesis count surface and other metrics.

Acknowledgments

[56] The authors gratefully acknowledge NOAA Grant NA97OG0241 and ONR Grant N00014-001-0092 for supporting the research described in this paper. The prototype implementation of CUBE exists as a source code library held at the Center for Coastal and Ocean Mapping, and is being licensed to members of the Industrial Consortium associated with the Center on a non-exclusive, royalty-free basis. Further details of the Industrial Consortium and licensing agreements are available from the authors.

References

Bacon, C. R., J. V. Gardner, L. A. Mayer, M. W. Buktenica, P. Dartnell, D. W. Ramsey, and J. E. Robinson, Morphology, volcanism, and mass wasting in Crater Lake, Oregon, *Geol. Soc. Am. Bull.*, *114*, 675–692, 2002.

Box, G. E. P., Sampling with Bayes' inference in scientific modeling and robustness, *J. R. Stat. Soc. Ser. A*, *143*, 383–430, 1980.

Calder, B. R., and S. M. Smith, A time/effort comparison of automatic and manual bathymetric processing in real-time mode, paper presented at U.S. Hydro 2003, U.S. Hydrogr. Soc., Biloxi, Miss., 2003.

Caress, D. W., and D. N. Chayes, New software for processing sidescan data from sidescan-capable multibeam sonars, in *Proc. IEEE*, *2*, 997–1000, 1995.

Caress, D. W., and D. N. Chayes, Improved processing of Hydrosweep DS Multibeam Data on the R/V Maurice Ewing, *Mar. Geophys. Res.*, *18*, 631–650, 1996.

Chib, S., Estimation and comparison of multiple change-point models, *J. Econometrics*, *86*, 221–241, 1998.

Chotiros, N. P., D. E. Smith, J. N. Pipe, B. K. McCurley, K. Lent, N. Crow, R. Banks, and H. Ma, Acoustic penetration of a sandy sediment, in paper presented at Underwater Acoustics Group Conference, Inst. of Acoustics, Southampton, U.K., 2001.

Cormen, T. H., C. E. Leiserson, and R. L. Rivest, *Introduction to Algorithms*, MIT Press, Cambridge, Mass., 1990.

Debese, N., Use of a robust estimator for automatic detection of isolated errors appearing in bathymetry data, *Int. Hydrogr. Rev.*, *2*, 32–44, 2001.

Debese, N., and P. Michaux, Détection automatique d'erreurs ponctuelles présentes dans les données bathymétriques multifaisceaux petits fonds, paper presented at Canadian Hydrological Conference 2002, Can. Hydrogr. Soc., Toronto, 2002.

Depner, J., and J. Hammack, Area based editing and processing of multi-beam data, paper presented at International Conference on Shallow Water Survey Technologies, DSTO Australia, Sydney, Australia, 1999.

Depner, J., B. Reed, S. Byrne, J. Parker, M. Paton, L. A. Mayer, and C. Ware, Dealing with increasing data volumes and decreasing resources, paper presented at the IEEE Oceans 2002 conference, Biloxi, Miss., 2002.

Du, Z., D. E. Wells, and L. A. Mayer, An approach to automatic detection of outliers in multibeam echo sounding data, *Hydro. J.*, *79*, 19–25, 1996.

Dziak, R. P., C. G. Fox, A. M. Bobbitt, and C. Goldfinger, Bathymetric map of the Gorda Plate: Structural and geomorphological processes inferred from multibeam surveys, *Mar. Geol.*, *22*, 235–250, 2001.

Eeg, J., On the identification of spikes in soundings, *Int. Hydrogr. Rev.*, *72*, 33–41, 1995.

Eichhubl, P., H. G. Greene, and N. Maher, Physiography of an active transpressive margin basin: High-resolution bathymetry of the Santa Barabara basin, southern California continental borderland, *Mar. Geol.*, *184*, 95–120, 2002.

Flood, R., V. Ferrini, and L. A. Mayer, Multibeam bathymetry and backscatter at 300 kHz in the SAX99 study area, west Florida shelf, *Eos Trans. AGU*, *81*(48), Fall Meet. Suppl., Abstract OS62A-17, 2000.

Fox, C. G., Objective classification of oceanic ridge-crest terrains using two-dimensional spectral models of bathymetry: Application to the Juan de Fuca Ridge, *Mar. Geol.*, *18*, 707–728, 1996.

Gamerman, D., Markov chain Monte Carlo for dynamic generalised linear models, *Biometrika*, *85*, 215–227, 1998.

Gardner, J. V., P. Dartnell, K. J. Sulak, B. R. Calder, and L. Hellequin, Physiography and late quaternary-holocene processes of northeastern Gulf of Mexico outer continental shelf off Mississippi and Alabama, *Gulf of Mexico Science*, *20*, 132–157, 2001a.

Gardner, J. V., L. A. Mayer, J. E. Hughes Clarke, P. Dartnell, and K. J. Sulak, The bathymetry and acoustic backscatter of the mid-shelf to upper slope off panama city, florida, northeastern gulf of mexico, *U.S. Geol. Surv. Open File Rep., 01-448*, 2001b.

Gerlach, R., C. Carter, and R. Kohn, Efficient Bayesian inference for dynamic mixture models, *J. Am. Stat. Soc.*, *95*, 819–828, 2000.

Glang, G., M. Cisternelli, and R. Brennan, NOAA ship WHITING data aquisition and processing report, *Tech. Rep. OPR-A397-WH*, Natl. Oceanic and Atmos. Admin., Washington, D.C., 2000.

Gourley, M., and K. DesRoches, Clever uses of tiling in CARIS/HIPS, paper presented at the 2nd International Conference on High Resolution Survey in Shallow Water, Cent. for Coastal and Ocean Mapp. and Joint Hydrol Cent., Portsmouth, NH, 2001.

Gourley, M., and D. Dodd, *HIPS: Hydrographic Information Processing System*, CARIS White Paper 21, Universal Systems, Fredericton, N.B. Can., 2000.

Grevemeyer, I., B. Schramm, C. W. Devey, D. S. Wilson, B. Jochum, J. Hauschild, K. Aric, H. W. Villinger, and W. Weigel, A multibeam-sonar, magnetic and geochemical flowline survey at 14°14′ S on the southern East Pacific Rise: Insights into the fourth dimensional of ridge crest segmentation, *Earth Planet. Sci. Lett.*, *199*, 359–372, 2002.

Hare, R., A. Godin, and L. A. Mayer, Accuracy estimation of Canadian swath (Multi-beam) and sweep (Multi-Transducer) sounding systems, Tech. rep., Can. Hydrogr. Serv., Ottawa, Ont., 1995.

Haykin, S. S., *Adaptive Filter Theory*, 3 ed., Prentice-Hall, Old Tappan, N.J., 1995.

Hughes Clarke, J. E., L. A. Mayer, and D. E. Wells, Shallow water imaging multibeam sonars: A new tool for investigating seafloor processes in the coastal zone and on the continental shelf, *Mar. Geophys. Res.*, *18*, 607–629, 1996.

IHO Committee, IHO standard for hydrographic surveys, *Spec. Publ., 44*, Int. Hydro. Organization, Monaco, 1996.

Jakobsson, M., Hypsometry and volume of the Arctic Ocean and its constituent seas, *Geochem. Geophys. Geosyst.*, *3*(5), 1028, doi:10.1029/2001GC000302, 2002.

Lee, S. E., P. J. Talling, G. G. J. Ernst, and S. J. Hogg, Occurrence and origin of submarine plunge pools at the base of the US continental slope, *Mar. Geol.*, *185*, 363–377, 2002.

Lirakis, C. B., and K. P. Bongiovanni, Automated multibeam data cleaning and target detection, paper presented at the IEEE Oceans 2002 conference, Biloxi, Miss., *1*, 719–723, 2000.

Lurton, X., Swath bathymetry using phase difference: theoretical analysis of acoustical measurement precision, *IEEE J. Ocean. Eng.*, *25*, 2000.

Maybeck, P. S., *Stochastic Models, Estimation and Control*, vol. 1, Academic, San Diego, 1979.

Richardson, M. D., et al., Overview of SAX99: Environmental considerations, *IEEE J. Oceanic Eng.*, *26*, 26–53, 2001.

Richardson, S., and P. J. Green, A Bayesian Analysis of Mixtures with an Unknown Number of Components, *J. R. Stat. Soc. Ser. B*, *59*, 731–758, 1997.

Smith, S., L. Alexander, and A. Armstrong, From high density bathymetry to Next-Generation ENCs, *Int. Hydrogr. Rev.*, *3*, 12–26, 2002.

Thorsos, E. I., et al., An overview of SAX99: Acoustic measurements, *IEEE J. Oceanic Eng.*, *26*, 4–25, 2001.

Varma, H. P., M. Boudreau, M. McConnel, M. O'Brien, and A. Picott, Probability of detecting errors in dense digital bathymetric data sets by using 3D graphics combined with statistical techniques, *Lighthouse*, *40*, 31–36, 1989.

Ware, C., L. Slipp, K. W. Wong, B. Nickerson, D. E. Wells, Y. C. Lee, D. Dodd, and G. Costello, A system for cleaning high volume bathymetry, *Int. Hydrogr. Rev.*, *69*, 77–94, 1992.

West, M., and J. Harrison, *Bayesian Forecasting and Dynamic Models*, 2 ed., Springer-Verlag, New York, 1997.