

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) 15-09-2016		2. REPORT TYPE Final		3. DATES COVERED (From - To) 15-12-2015 / 31-08-2016	
4. TITLE AND SUBTITLE Development of Theoretical and Computational Methods for Single-Source Bathymetric Data				5a. CONTRACT NUMBER N00014-16-1-2035	
				5b. GRANT NUMBER 11893686	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Calder, Brian, R				5d. PROJECT NUMBER 100001828	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Office of Sponsored Research, University of New Hampshire 51 College Road, Service Building 107 Durham, NH 03824-3585 USA				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research 875 N. Randolph Street Suite 1425 Arlington, VA 22203-1925				10. SPONSOR/MONITOR'S ACRONYM(S) 342	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Unclassified Unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Providing a coherent navigational picture from multiple, inconsistent, source documents can be time consuming, subjective, and prone to error. A method is outlined for fusing the information inherent in such source documents, at different scales, into a single picture for the marine navigator which includes indicators for safely navigable areas, those to be treated with caution, and those known to be unsafe for navigation. A measure of algorithm reliability, which reflects the degree of inconsistency of the source documents, is also provided. A conceptual outline of the method, and a preliminary, limited demonstration in two areas around Norfolk, VA, are given.					
15. SUBJECT TERMS Bathymetric Modeling, Single-Source Bathymetric Database, Safe Operating Contour, Navigable Area, Bathymetric Uncertainty, Decision Fusion, Algorithm Reliability					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			BRIAN R. CALDER
U	U	U	UU	16	19b. TELEPHONE NUMBER (Include area code) +1 (603) 862-0526.

Development of Theoretical and Computational Methods for Single-Source Bathymetric Data

Brian R. Calder¹

Center for Coastal and Ocean Mapping
University of New Hampshire, Durham NH 03824

Award Number: N00014-16-1-2035

Introduction

When planning for navigation of surface or sub-surface vessels, many different sources of information must be taken into account in order to establish the safe operating envelope for the vessel. Sources include electronic charts at different scales, gridded bathymetry of different ages and quality, special purpose vector products, and auxiliary information such as notices to mariners. Not all of these sources, however, are necessarily consistent making it especially challenging to provide a coherent picture for planning.

This project looks at means to provide a consistent planning picture from disparate source data, using only the sources that would be available to the fleet while at sea. Since the planning purpose can be relatively complex, the project has focused on generating a single output product, with uncertainty, in the form of a safety contour, set at a depth to be determined by the vessel commanding officer. This provides a simpler target for initial efforts, and a gateway product to flush out the difficulties in generalizing to be more complex analyses.

This report details the first year's work on the project, intended to be a foundational year that established the basis for methods to approach the project's goals.

Goals and Objectives

The overall goals for the project are to provide to the navigator tools that support planning activities, and the means to determine the limitations of knowledge surrounding those tools.

The main objectives for this first year of the project were to establish an approach to treating the available data sources and their known uncertainties so that a safety contour with associated uncertainty estimate could be established; and to consider mechanisms for expressing the uncertainty of the reconstructed safety contour such that the navigator would be informed as to the quality of the reconstruction.

¹ Principal Investigator. T: (603) 862-0526, E: brc@ccom.unh.edu

Approach

Recognizing that vessels at sea do not have, and are unlikely to ever have, the entire source database for the area in which they are operating, the methods here are restricted to the types of products that will be available, attempting to form a coherent picture of any given area using only the chart-scale data, with at most data from DBDB-V to provide a higher-resolution product.

The work has focused on the idea that an experienced navigator would use more than just the soundings and depth areas on a chart to form a mental model of the configuration of the seafloor in the area of interest. Although it is generally not well quantified, a navigator interprets the nature of the chart, the configuration of the soundings and contours, the known or estimated character of the seafloor, the age of the data, and other information in order to come up with an impression of what the seafloor is doing in the areas between the soundings, varying between optimistic and pessimistic outlooks, and therefore where it is safe to navigate. The approach therefore has been to ask if it were possible to attempt to teach an algorithm to do something similar. That is, is there a system that would allow for a reconstruction of the data from the various sources which could be influenced between a pessimistic and optimistic estimate of depth based on a parameter set that could be adjusted to reflect the opinion of expert observers, environmental events within the area, the current tactical situation, and other factors? If so, is it possible to use this to construct a consistent estimate of the location of a safety contour, and its associated uncertainty?

These questions have been approached in a two stage process: first build the reconstructions from each source, then fuse the reconstructions while maintaining the uncertainty estimate for the safety contour. The reconstruction method envisions a parametric surface being formed from the available bathymetric data, with the parameter modulated to allow for expert opinion, source metadata, present tactical situation, etc., allowing the formation of both a pessimistic and optimistic reconstruction. Cutting these surfaces at a given safety contour depth splits the spatial area into "go", "caution", and "no go" regions, leading to a ternary classification of the area of interest. The ternary classification surfaces from each source can then be fused, with weighting, to provide an overall ternary classification. The "caution" region provides uncertainty of the location of the safety contour; analysis of the fusion results provide some evidence for second-order uncertainty in the placement of the classification boundaries.

Work Completed

In collaboration with NRL Stennis researchers Dr. Paul Elmore and Brett Hode, the first year of the project has seen development of theoretical models for surface reconstruction and ternary classification fusion, along with initial development of a test system to explore these ideas.

Theoretical Model

For the surface reconstruction problem, the model developed was to provide a parametric surface that, taken between two endpoints along a transect (Figure 1), could be adjusted between a simple slope (the most optimistic view a navigator might have) and preserving the shoalest depth for as long as possible before switching to the deeper reconstruction depth (the most pessimistic view that might be considered) as the parameter varied in the range $[0,1]$. A number of different parametric reconstructions could be considered, but the simplest is a Non-Uniform Rational B-spline (NURBS) where the control points can be adjusted using a suitable control function to give the desired parameter range. In each case, a parameter range would be established providing the most pessimistic and optimistic reconstruction that the operator would be willing to accept for each data region (Figure 2).

The parameterization of the reconstruction allows the method to adapt to different conditions in different sources. More dubious sources, for example, could have their parameter range narrow and high (e.g., $[0.9,1]$), while a recently updated chart of a simple area might have a narrow and low range, but a recently updated chart of a complex area might have a wider range to reflect the idea that although the data is expected to be good, the reconstruction from the limited chart-scale data might not be in the right place.




Figure 1: Example of extremes of pessimistic and optimistic reconstruction given two endpoints of a transect, (A, B). Adjustments of the controlling parameter cause the reconstruction to move smoothly from one extreme to the other, allowing the reconstruction or range of possible reconstructions to be modulated by external constraints.




Figure 2: Example of a plausible range of parameters that might be accepted as usable reconstructions of the depth along a transect between the ultimate optimistic and pessimistic reconstructions.

The parameterization also allows for the monitoring of the behavior of data over time. A schedule of adjustment of the parameter from its default state to a more pessimistic reconstruction could be established as a function of time (Figure 3) with the option of point-wise increases in parameter as environmental effects (e.g., storms, tsunamis) are observed in the historical record post-survey.

Finally, the parameterization can also be used to provide calibration on the degree of pessimism appropriate to any given chart configuration. That is, although it is possible to make the reconstruction at any location on the pessimism-optimism scale, it is difficult to provide a valid means to predict the appropriate location for any given configuration of the chart. It is therefore envisioned that expert opinion would be garnered for a variety of chart configurations by providing a chart configuration, the predicted output location of the safety contour, and a control over the reconstruction, and letting the experts adjust the reconstruction until it meets their expectations, given the survey metadata. Statistics of these data then provide the guidance on how to reconstruct particular configurations, and can be correlated with configuration descriptors to provide a constructive parameter selection method.

Not all operators will agree on the reconstruction, of course, and particular tactical situations might skew the preferred reconstruction in either direction. It is therefore likely that the end user will have some small adjustment of the parameter available, within the overall range of parameter estimates established by the expert panel.




Figure 3: Example of possible time-sequence adjustment of reconstruction parameter for more/less stable areas, with effects of storms. More geologically stable areas might be discounted at a slower rate, while areas that change dynamically (e.g., the North Sea) might more rapidly move to pessimistic reconstructions. Storm events, tsunamis and the like would provide point-wise increase in the level of pessimism about the relationship of the pre-event observed data to the current configuration of the seafloor.




Figure 4: Example mock-up of the reconstruction algorithm showing “no go” (red), “caution” (yellow) and “go” (green) regions derived for a keel depth of 14m.

The concept of optimistic and pessimistic reconstructions in each area directly leads to a model of uncertainty for the safety contour. Specifying a safety depth for the ship cuts the transect model twice (Figure 4), and it is clear that in the region shoalward of the optimistic reconstruction it will be unsafe to travel; in the region

deepward of the pessimistic reconstruction it will be safe to travel; but in between these regions there is a grey area. The size of the “caution” area reflects the uncertainty in the reconstruction engendered by the data.

Applied spatially, this method provides, for each source, a ternary spatial classification of the area of interest. The ternary classification also provides the means to fuse the different sources into a composite product for display to the navigator, while still maintaining the uncertainty expressed by the “caution” region, including the inconsistency between the different sources. Each data source will generate a ternary classification, which will be to greater or lesser degree inconsistent with the others.

Treating the ternary classifications as a stack (Figure 5) makes it clear that at any point in the area of interest, a vertical transect of the stack provides a number of estimates of the true state of a ternary classification for the point. Consequently, it is possible to treat the fusion problem as a distribution estimation problem, using all of the sources to estimate the class probability mass function at each point.




Figure 5: Starting position for the fusion process with a stack of different ternary classifications for the same area from different source data, each of which is not necessarily consistent with the other classifications. Fusion resolves to estimating the probability mass function for the ternary classes at each point of interest.




Figure 6: Example end-members of fused categorical distributions showing (left) a case with consistent sources and (right) inconsistent sources. Cases where there are inconsistent sources would be displayed to reflect the greater uncertainty in the classification, which would typically occur in border regions between the different classes.

A Bayesian solution to the estimation problem is envisioned using the classical Dirichlet conjugate prior for the categorical distribution. A semi-empirical prior is expected to be constructed from the best-available gridded bathymetry (typically DBDB-V), although higher-resolution gridded data could be integrated if available, and each piece of source data will then be assimilated to form the overall estimate of the probability mass of the ternary classes. Maximum *a posteriori* reconstruction will then result in a ternary classification that best represents all of the available data, which still preserving the “caution” area that provides information on the uncertainty of the safety contour’s location. In addition, analysis of the shape of the estimated categorical distribution as each location can be used to determine the degree of consistency of the various sources (Figure 6) and therefore provide some evidence as to the stability of the reconstruction, which can also be displayed to the user as qualifying information for the algorithm and data.

The information that can be gathered from each source is expected to vary, possibly considerably: some sources might be considered more canonical than others, irrespective of their data or metadata. It is expected, therefore, that the fusion scheme will have to accommodate each source being weighted; the Bayesian scheme allows for this. Most of the evidence for these weights is going to be derived from more qualitative assessment of source data, which is a more linguistic processing problem less suited to classical statistical methods. We have therefore collaborated with researchers at NRL Stennis, who will provide methods to assess these weights and their influence on the fusion scheme.

Similarly, there are many features of charted information which are not directly reflected in the bathymetric components of the chart. Features such as local notices to mariners, chart notes (“discolored water observed”), etc., could change the interpretation of the chart by an experienced analyst, and should therefore also be reflected here. The case is envisioned where these “extra information” notes would be applied as local adjustments to either the bathymetric parameterization during re-




Figure 7: Example areas selected for demonstration of the reconstruction and fusion techniques, in the vicinity of the Chesapeake Bay, VA. Background chart is a portion of the NOAA ENC of the area.

construction, or directly to the ternary classifications, shifting the classification rather than entirely replacing them. The mechanisms required for this are expected to be relatively simple to implement, but the translation of the source data into the desired adjustments is expected to be a more significant task. Again, this is being tackled in collaboration with colleagues at NRL Stennis.

At this stage of the project, the visualizations being considered are relatively simple. However, models of uncertainty that can be readily applied in two-dimensional displays have been considered, including the use of cross-hatching, shading, and textures to represent areas of uncertainty in the reconstruction.

Example Reconstruction

In order to illustrate the theoretical model, consider two regions around the entrance to the Chesapeake Bay, VA, as shown in Figure 7. Area A is a complex area around the pilotage location for the entrance, including the split of the main channel between the approaches to Norfolk and the southern bay, and that for Baltimore and the northern bay; area B is a nominally simpler area in the outer approaches. Each of these areas is covered by different scale bands of ENCs from the NOAA catalogue: area A is covered by scale bands 3 (coastal), 4 (approach), and 5 (harbor), while ar-

ea B is covered by scale bands 2 (general), 3, and 4. NOAA indicates that General ENC's have scales from 1:600,000 to 1:500,000; Coastal from 1:150,000 to 1:600,000; Approach from 1:50,000 to 1:150,000; and Harbor from 1:5,000 to 1:50,000. Each ENC is generalized to the appropriate scale, and therefore provides a different view of the area, and in particular a different (and not entirely consistent) implication on the available under-keel clearance; Figures 8-9, with area B showing more variability than area A. Even in the absence of any distinct difference in the safety contours (e.g., in area A), the different density of supporting soundings can significantly affect the ability to estimate depth in the area – for a navigator, or for an algorithm.

Each of the ENC's for each area was reconstructed by the same process. First, a Delaunay triangulation of the available sounding information and depth contours was constructed. For each triangle edge that connected a depth above and below the required safety contour depth, the NURBS-based interpolation scheme was applied to determine the position of the safety contour depth along the edge using a given optimism level; the position was then added to the data being manipulated as a pseudo-sounding. An optimistic and pessimistic solution was constructed in each case. The optimistic and pessimistic point clouds were then gridded using an inverse-square distance nearest-neighbor algorithm. The ternary output was then generated by setting the grid to "not navigable" (numerical value 0) where the optimistic reconstruction is shoaler than the safe operating contour depth, "navigable" (numerical value 2) where the pessimistic reconstruction is deeper than the safe operating contour depth, and "caution" (numerical value 1) otherwise. The grid was marked "no data" (numerical value 3) if either input is lacking data.

The ternary grids from the three ENC's in each area were then fused as outlined before, using arbitrary weights of 10, 5, and 1 for the largest, middle, and smallest scale ENC in each area. The ternary grids were all constructed at 100m resolution, although the fusion algorithm applies area-normalized weighting to ensure that higher resolution input grids would not inadvertently re-weight the fusion.

An overview of the inputs, and fused output, for area A is given in Figure 10 using a 12m safe operating contour depth. Due to the lack of sounding values in the Coastal chart (A3, top left), the whole channel to Baltimore is missing, and disconnected in the Approach chart (A4, top right). The fusion shows the channel as connected (bottom right, green areas), but uncertain at the edges (blue), which might give some pause for a navigator opting to take the channel. This is more strongly emphasized in the detailed fusion analysis shown in Figure 11, which includes the estimate for the reliability of the fusion determination. The lower reliability for the conclusion that the northern channel is navigable reflects the difference of opinion in the inputs as to the state of this area, which again might cause concern, or at least elicit a warning, from the navigator.

A detailed view of the fusion within the throat of the connection to the northern channel is shown in Figure 12. Clearly, although the algorithm is indicating that the area is navigable (due to the weighting of the input from the Harbor chart), the conclusions from the other sources have reduced the reliability, making the passage dubious at best.



Figure 8: Examples of ENC data for area A (Figure 7); note that depths are given in meters. The Coastal chart (top) shows only general features of the area, while Approach (middle) and Harbor (bottom) charts show increasing levels of detail and supporting soundings that allow for better reconstruction of the depth in the area.




Figure 9: Examples of ENC data for area B (Figure 7); note that depths are given in meters. The General chart (top left) shows a small pocket of heavily generalized safe water (deeper than the 60' contour), which is much at odds with the same safety contour shown on the Coastal chart (top right) and the detail of the Approach chart (bottom).




Figure 10: Inputs, and fused output, for area A (Figure 7) using a safety operating contour of 12m. The fused result mostly follows the Harbor chart input (bottom left) due to weighting.




Figure 11: Fusion result detail for area A (Figure 7), showing the fused ternary diagram (left), and the estimated reliability of the fused output (right). High reliability is indicated where all of the inputs agree on the state of an area, with lower reliability where the inputs disagree.




Figure 12: Detail of the fusion in the throat of the connection to the northern channel in area A.

The fusion in area B is a more complex situation, since the inputs are radically different from each other about their interpretation of the available underkeel clearance in the area. Figure 13 shows the input ternary diagrams and fused output for the area, using the same fusion weightings as for area A, demonstrating the confusion on navigability due to generalization of the ENC's for scale, particularly in the case of the General chart (top left): due to the generalization and lack of supporting sounding information, at best the chart appears to have two disconnected navigable (green) areas, even though the Coastal (top right) chart shows them connected, and the Approach (bottom left) chart shows the majority of the area as navigable, including a channel to the north.

The inconsistency in the inputs is also reflected in the detail of the fusion, shown in Figure 14. Here, although the fused result follows the Approach chart for the most part due to the weighting applied, the reliability indicates that large areas are driven by inconsistent inputs, so that although it is generally possible to say reliably where it *not* possible to navigate, it is difficult to be sure where the conclusion of navigability is valid. It seems likely that a navigator responsible for planning a passage through such an area might be reluctant to recommend any route through the nominally navigable areas under such circumstances.

The detailed analysis of a "maybe navigable" area to the west of the region (in the connection between the channel to the north and the main navigable region) shown in Figure 15 highlights the robustness and flexibility of the fusion method. Although one of the sources has no information (and therefore is reconstructed as

“not navigable”) and the other two sources disagree, the fusion algorithm has still generated an appropriate output, and rated the reliability low to reflect the confusion between the available sources. Different weightings would of course change the scale of the reliability assessment, although more closely balanced weightings would result in higher confusion, and hence lower reliabilities. Therefore although the algorithm is clearly sufficiently flexible to support the reconstruction and fusion tasks required, obtaining reliable and effective results will rely heavily on appropriate optimism weightings being generated for the reconstruction, and plausible fusion weights being developed from the supporting material and metadata on the charted products.




Figure 13: Inputs, and fused output, for area B (Figure 7), using a safety operating contour of 19m. The fused result (bottom right) mostly follows the B4 (Approach chart) input due to weighting.




Figure 14: Fusion result detail for area B (Figure 7), showing the fused ternary diagram (left), and the estimated reliability of the fused output (right). The confusion of reliability information due to the radically inconsistency of the inputs demonstrates that although the area is nominally simple, decisions to navigate in it are not.




Figure 15: Detail of the fusion in the “maybe navigable” region on the western side of the region connecting the channel to the north with the main body of the area.

Future Plans

Having spent the majority of the first year in development of the theoretical models for bathymetric reconstruction and classification fusion, future work on the project would include development of the prototype system to embody these ideas. The ob-

jectives of this work would be to expand the types of reconstructions that could be achieved, to trouble-shoot the reconstruction methods with a suitable chart portfolio in a given area, and to elicit a first estimate of the range of parameter values associated with optimistic and pessimistic reconstructions of a selection of chart configurations and metadata types.

In addition, further work in collaboration with NRL Stennis would be planned to pursue the problem of estimating weighting structures for different source data types, and methods for translating non-bathymetric chart information into the reconstruction/classification/fusion scheme.

Finally, results from these developments would allow for greater experimentation with visualization methods for this type of data. Further work on display of uncertainty contours, and in particular the secondary uncertainty associated with inconsistency of the source data reconstructions, would be expected.