



**Acoustics'08
Paris**
June 29-July 4, 2008
www.acoustics08-paris.org

automated change detection with area matching

John Dubberley, Marlin Gendron and Maura Lohrenz

Naval Research Laboratory, Bldg. 1005 Rm D-23, Stennis Space Center, MS 39529, USA
john.dubberley@nrlssc.navy.mil

When resurveying a geographic area of the seafloor during sidescan change detection operations, an automated method to match bottom objects imaged previously with objects imaged in the resurvey can increase efficiency. The geographic position of a new object relative to a historical object is a good indicator of a match. However, due to position error within the survey, there may be more than one spatially-close object in the new imagery. To complicate matters further, the reflected energy from the new object may be vastly different given a different resurvey incidence angle or the partial burial of the object. In addition, the resurveyed object image may be below the threshold set for automatic recognition and falsely eliminated. This presentation will address these problems and suggest possible methods for matching "constellations" of bottom objects using Dijkstra's Minimum Cost - Maximum Flow algorithm, control point matching, and the data association procedure.

1 Introduction

When sidescan surveys are updated with new imagery from a previously scanned area, the surveyors are interested in both the changes between the two surveys and aligning reference points. Our goal is to automate the change detection process, where possible, and aid the surveyor's decision making when automation is not applicable. The challenges to overcome include position errors internal to each survey, errors between surveys, sidescan imagery resolution, transient targets, and the movement of bottom objects.

Position errors for sidescan surveys have many sources. With the advent of widespread high quality GPS navigation instruments, errors in determining the position of the towing platform can range from one to five meters. Errors occur due to inexact sound speed estimates for tow body height and target slant range, as well as errors associated with the tow body's attitude [1]. Position errors due to cable layback and slew tend to be the largest contribution to position error and can be in the tens of meters during normal operations [2]. Some surveyors greatly reduce the cable sourced errors by using acoustic transponders; however, very few baseline surveys have been collected using transponders, and it is unlikely most new surveys will either. This implies that an automatic change detection system must be able to handle position errors of up to thirty meters.

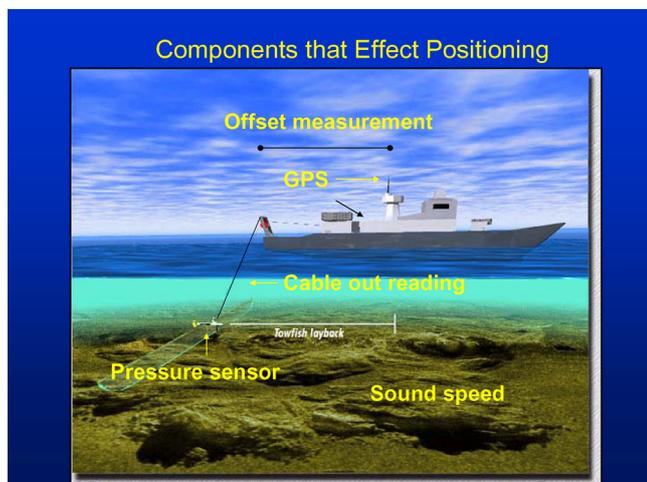


Figure 1. Illustration of some causes of position error.

Transient items must be deleted, after determining they are no longer in the survey, without eliminating objects that may have been randomly skipped by a new survey. Bottom objects move due to currents and storms, and changing burial percentage makes inter-survey comparison of objects more difficult. Resurveys also can use different sidescan sensors between surveys, creating resolution and image-matching problems.

Here we will use "Feature Matching" to best match objects individually and "Area Matching" when there is great enough ambiguity between potential matches.

2 Algorithms

2.1 Automatic Target recognition

A real-time detection algorithm, developed by the authors, ingests one scan line at a time to locate targets within sidescan imagery. Due to real-time processing considerations, the authors' algorithm relies on a patented geospatial bitmap GB technique [3]. Across-track bright and shadow positions, lengths, and intensity information are immediately gathered from the scan line and stored in two one-dimensional GBs: one each for shadows and brights. A circular lookup table is created to "window" the imagery several scan lines at a time. This lookup table is the same width as the GBs and is populated with the positions and run-lengths of shadows and brights stored in the GBs. The window is used to make the final detection decision.

Shadows and brights in a scan line are located by first adaptively obtaining a lower intensity threshold, i_{min} , such that all samples of intensity less than i_{min} are considered shadows. An upper intensity threshold, i_{max} , is set such that all samples of intensity above i_{max} are considered brights. An appropriate gamma shift converts image intensities to fit a normal distribution, such that i_{min} and i_{max} are set to the quartiles of the shifted (normal) distribution. After i_{min} and i_{max} have been determined for scan lines with maximum intensity value > 128 , the port and starboard halves of the scan lines are processed separately. Each half of the scan line can be represented by a vector, X , of length N . The following method is used to process shadows and brights for the starboard side; the port side is processed similarly.

Two GBs of size $1 \times N$ are created, one for shadows and one for brights. A different gamma adjustment, γ , based on an

error approximation of the side-scan sonar parameters, is computed for position x within X , as shown in (1).

$$\gamma = e^{-\beta x/N} \tag{1}$$

β is based on the sonar parameters, such as time-varying gain. As β approaches infinity, the gamma correction approaches 0 over a greater range of X (Figure 2), and therefore has less affect on intensity thresholds.

The bright and shadow thresholds $I_{\min}(x)$ and $I_{\max}(x)$ are defined in (2) and (3). All pixels with intensity values above $I_{\max}(x)$ are considered brights, while all with intensities below $I_{\min}(x)$ are considered shadows, and the corresponding bits in the bright and shadow GBs are set, as shown in Figure 3.

$$I_{\min}(x) = i_{\min}(1 - \gamma) \tag{2}$$

$$I_{\max}(x) = i_{\max}(1 + \gamma) \tag{3}$$

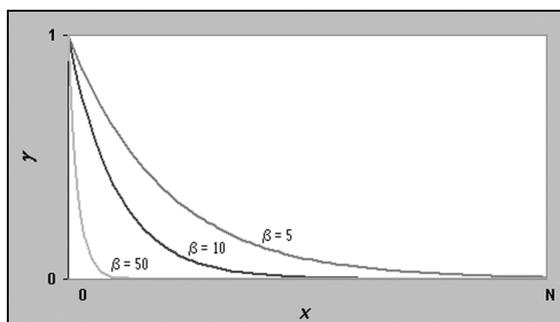


Figure 2. Gamma as a function of β .

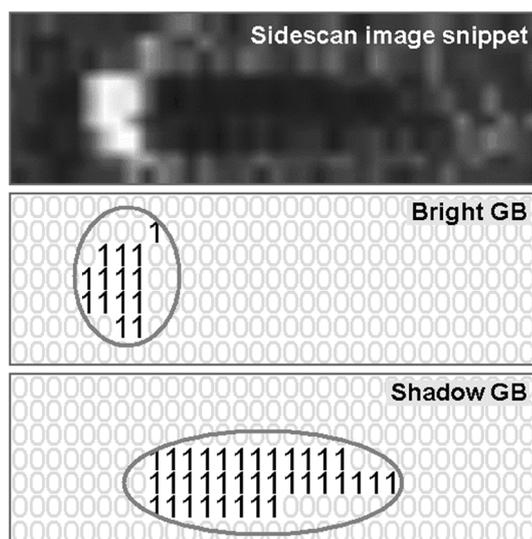


Figure 3. GBs facilitate computer-aided detection of objects in SSI.

Each row of bits in both GBs corresponds to a single scan line in the image. All pixels in the image with intensity greater than upper threshold I_{\max} are considered “brights” and the appropriate bits in the bright GB are set. Likewise,

all pixels in the image with intensity less than lower threshold I_{\min} are considered “shadows” and the appropriate bits in the shadow GB are set.

Figure 4 illustrates how the intensity thresholds vary over x for a given γ . For example, the closer a pixel is to the center of the scan, known as nadir ($x = 0$), the greater its intensity must be to be detected as a bright [4], and the lower its intensity must be to be detected as a shadow. It is interesting to note that the two threshold curves do not diverge from their respective asymptotes (i_{\min} and i_{\max}) at the same rate as they approach nadir. This is by design, because shadows are more detectable than brights in SSI [5] – [8]. In other words, a single shadow threshold value (i_{\min}) suffices for more values of x than a single bright threshold value (i_{\max}).

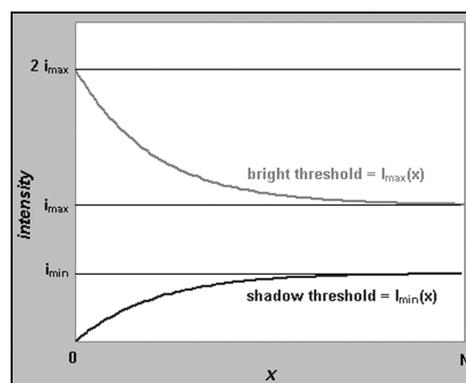


Figure 4. Intensity thresholds for brights and shadows.

Targets of various shapes and sizes are traditionally discerned in sidescan by the shadow’s dimensions, which vary as a function of both beam angle and feature size [9]. An adaptive filter is used to “complete” the target’s bright spot, reducing or enlarging it based on the shadow. The feature then can be classified more accurately using dimensions of both the shadow and completed bright spot. The filter outputs a snippet of imagery that contains only three gray-scale values: black for shadow bits, white for completed bright spot bits, and gray for all other bits (figure 5). These completed shapes are used by a feature-matching algorithm to match historical survey targets with new targets.

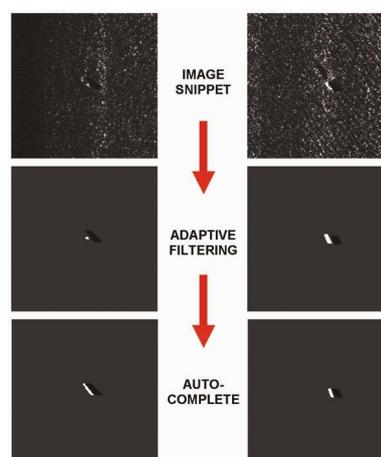


Figure 5. A target’s bright spot is “completed”, based on its shadow.

In the series on the left, the bright spot is enlarged to match the shadow. On the right, the bright spot is reduced.

2.2 Feature Matching

The actual matching process starts when the ACDC system has detected a contact. Most surveys will overlap their sidescan swaths so each bathymetry point is imaged twice from opposite headings. Even if there are swath position errors, this overlap helps ensure all points will be seen. Imaging potential objects on opposite headings makes it more likely to identify objects that may be partially buried.

The highest probability weight is assigned to the position match between the object and the previously surveyed object, since in most cases, there will be only one previous survey object in the area. Next, the objects' sizes are calculated and compared. Adequate latitude (150%) is allowed here so that images with differing views will not automatically rule out a potential object match. If the match is sufficient, it is marked and the contact is amended to include the new image. If the contact has only one possible match, and the match is not strong enough to automatically flag (or not weak enough to automatically ignore), the contact is sent to the surveyor to match. Otherwise, the image is forwarded to the area matching routine.

2.3 Area Matching

Area matching is the process of examining the positions of surrounding objects and using their geometric relationships to identify objects as those previously seen. Complications to this process include: some objects in both the previous and new surveys may be transients, some objects will have moved, and some objects may have not been identified in one or the other survey.

Various algorithms were considered to help solve the area matching problem. Neural Networks are commonly applied to pattern matching problems. They are flexible and easily applied. Their weakness lies in their training patterns, which are slow to evolve and challenging to tune to not over- or under-match. For these reasons we rejected Neural Networks.

We first tried Dijkstra's Minimum Cost - Maximum Flow algorithm for the area-matching problem. This method treats each point and its potential matches as nodes on a directed graph, with the object being evaluated as the root node. Probabilities of match are assigned to each node edge. The highest probability node geometry is selected and compared to the node geometry in the previous survey. Dykstra's algorithm is a method to efficiently solve these nodes. After preliminary work, we found that the graphs very rarely needed tree graphs with more than one connected node. All such graphs would reduce into the feature matching case. The algorithm was set aside because, in most cases, it did not improve the match.

Rigid Data Association (RDA) has been used to improve navigation for Autonomous Underwater Vehicles (AUVs) [10]. Typical AUV navigation systems get a GPS fix at the surface and use an Inertial Navigation System to estimate movement through the water and a Doppler Velocity Log to track movement over the sea floor, when applicable. These systems can bring swath-to-swath position errors down to 5

meters or less by assuming there is a correction transformation for each swath with constant position, heading, and velocity errors. This rigid transformation can be calculated by finding the position error between two or more objects imaged in both swaths. Each pair of potential matches creates a transform. These transforms are searched to find the one that causes the least overall match error. Using this transform, the position of all objects can be adjusted and then the closest object from the previous leg is chosen as the match for each object.

Rigid Data Association is ideal for area matching when using only a "lawn mower" pattern survey with low position errors. However, when change detecting between different surveys, when a lawn mower pattern is not used, and when non-constant current fields are present, RDAs applicability breaks down. In future work we plan to use it when it can be automatically recognized as applicable.

The final method we examined for area matching during change detection is Control Point Matching (CPM) [11]. This method simplifies the general matching of constellations to matching triangles. A window is placed around the newly sensed point, and all triangles formed by that point in the window are calculated in both the new and old surveys. The centroid of each triangle is calculated, and distances between the centroids of new and former survey triangles are calculated. If the centroid distance is less than the maximum error ellipse axis of the compared triangles, they are labeled a preliminary match. Next, the sphericity of the affine transformation between possible match triangles is calculated to estimate their similarity [12]. This results in a number on the unit interval. The triangle with the largest sphericity is the match, as long as the sphericity is greater than one half.

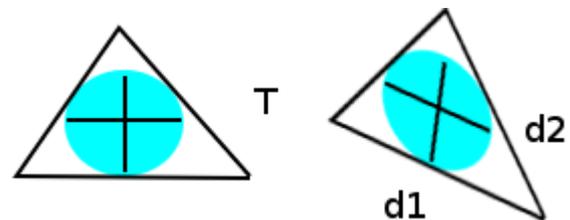


Figure 6.

Inscribed circle in left triangle related to the right triangle by affine transform T . Ellipse axes $d1$ and $d2$ form the transform of the circle.

On relatively difficult testing sets, Control Point Matching has been shown to correctly match 90% of the points and mismatch 2/3 of the remaining points [11]. Given the flexibility and general applicability of Control Point Matching, we chose this for our area matching algorithm in this project.

3 Preliminary Results

To demonstrate that our feature-matching algorithm works as expected, we used an initial test data set of 120 contact points randomly distributed in a one square kilometre area. Two survey data sets were computed from the initial data set by adding a random distance vector from a Gaussian distribution with a 7-meter 95% confidence radius. Three sets of points were too close together for feature matching

to apply. Two other points were not matched, one more than expected but within the tolerances of such a small sample size.

Since there is an operator available to match by sight and context, it is better to have a lower number of automatic matches if there is a corresponding decrease in mismatches. Control Point Matching must be tuned down by increasing the sphericity match threshold until the number of false matches decreases. In our preliminary test experiments, we reduced the number of false matches to below one percent of total contacts, with a drop in matches to 70% of total contacts. We believe with further work we can increase this match percentage with a similar or lower level of false matches.

4 Conclusion

Automatic intersurvey Change Detection is actively desired by surveying organizations. Any improvement over manually deconflicting data points will result in significant manpower savings. Given that the expected average contact density will allow feature matching to complete the match 90% of the time, Control Point Matching tuned down to eliminate false positives should allow us to exceed our goal of 80% autodetect on surveys of sufficient data quality. It is believed that 80% autodetection should halve the manpower required for deconfliction. Savings of fifteen minutes per contact should also realize deconflicting between survey swaths.

Acknowledgments

Work done in this paper was sponsored by the Office of Naval Research. The authors would also like to acknowledge NP6 James Hammack and Jan Saunders at the Naval Oceanographic Office.

References

- [1] S Anstee "The Effects of towfish motion on sidescan sonar images", Commonwealth of Australia DSTO 44 (1994).
- [2] R. Trahan Jr., D. S. Yeadon, M. Thiele Jr. "Optimal estimation of layback distance for marine towed cables", *IEEE Journal of Oceanic Engineering*, 16, 217-222 (1991).
- [3] M. Gendron, P. Wischow, M. Trenchard, M. Lohrenz, L. Riedlinger, and M. Mehaffey, *Moving Map Composer*, Naval Research Laboratory, U.S. Patent No. 6,218,965, (2001).
- [4] P. Blondel and B Murton, *Handbook of Seafloor Sonar Imagery*, New York, John Wiley & Sons, 23-46, (1997).
- [5] M. C. Mignotte, C. Collet, P. Perez and P. Bouthemy, "Markov Random Field and Fuzzy Logic Modeling in Sonar Imagery: Application to the Classification of Underwater Floor," *Computer Vision and Image Understanding*, 79, 4-24, (2000.)
- [6] M. C. Mignotte, C. Collet, P. Perez and P. Bouthemy, "Sonar Image Segmentation Using an Unsupervised Hierarchical MRF Model," *IEEE Transactions on Image Processing*, 9, 1216-1231, (2000).
- [7] S. Reed, E. Dura, D. M. Lane, Y. J. Petillot, "Extraction and Classification of Objects from Sidescan Sonar Bell," *IEEE Workshop on Nonlinear and Non-Gaussian Signal Processing*, 1-2, (2002).
- [8] S. Reed, Y. Petillot, J. Bell, "An Automatic Approach to the Detection and Extraction of Mine Features in Sidescan Sonar," *IEEE Journal of Oceanic Engineering*,. 28. 90-105, (2003).
- [9] J. P. Fish, H. A. Carr. *Sound Underwater Images: A guide to the Generation and Interpretation of Sonar*. American Underwater Search and Survey, Ltd. (1990).
- [10] E. Coiras, F. Baralli, B. Evans "Rigid Data Association for Shallow Water Surveys" , *IET, Radar, Sonar, & Navigation* , 1, 354-361 (2007).
- [11] D. Skea, I. Barrodale, R. Kuwahara, and R. Poeckert, "A control point matching algorithm", *Pattern Recognition*, 26 269-276 (1993).
- [12] N. Ansari and E. Delp, "Partial shape recognition: A landmark-based approach," *IEEE Trans. Pattern Anal. Machine Intell*, 12 470-483, 1990..