

Automated Target Detection and Grouping from Remotely Sensed Data

Brian Bonnlander, Clark Glymour and Choh Man Teng

Institute for Human and Machine Cognition
40 South Alcaniz Street, Pensacola FL 32502 USA

Summary

Targets are objects that rise above the bottom surface more than a specific amount defined by IHO survey standards. Currently target detection requires time consuming analysis by the human expert. The contracted task is to design and implement an automated target detector that can be used as a tool by human experts. We tested the implementation internally and have sent some of the test results to experts at NAVO for assessment. In addition we also designed and implemented a target grouping procedure that clusters the targets according to a proximity metric. The resulting grouping can be used to produce polygon outlines that will replace selected clusters of densely spaced targets.

Several issues and possible improvements were identified from our testing and analysis. They include alternative algorithms for target identification, computational optimization and parallelization of the implementation, application of machine learning algorithms for optimization of parameters for target identification, and a systematic testing regime.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the 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 this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 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 a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2006		2. REPORT TYPE		3. DATES COVERED 00-00-2006 to 00-00-2006	
4. TITLE AND SUBTITLE Automated Target Detection and Grouping From Remotely Sensed Data				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Human and Machine Cognition,40 South Alcaniz Street,Pensacola,FL,32502				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 29	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

1. Problem Description

IHO Standards for Hydrographic Surveys Special Publication No. 44 (International Hydrographic Bureau, 1998) specifies the minimum accuracy requirements for hydrographic surveys of various orders. One of the accuracy standards concerns “system detection capability”, which specifies the size of features that need to be reliably detected in a survey. For instance, an Order 1 survey is required to detect all cubic features that are greater than 2 meters on each side in water up to 40 meters deep, and all cubic features that are greater than 10% of the water depth on each side in depth beyond 40 meters. Similar standards are in place for surveys of different orders to accommodate areas with more or less stringent accuracy needs.

Remote sensing systems such as LIDAR and sonar provide high resolution, dense bathymetric measurements of the survey area. Each data point carries a margin of error, but the volume of data can be utilized to overcome some of this error. Target identification in this setting requires a combination of human expert judgement and precise and efficient data management and computation. In practice currently this task relies primarily on experts trained to hand pick the targets by visually inspecting a bathymetric map, a procedure that is both time consuming and prone to variations due to unarticulated differences in expert opinions. We aim to capture as much as possible the criteria the experts use in determining the acceptability of a potential target. By approximating these criteria using standardized data characteristics and objective algorithms, we aim to streamline the process of target detection and provide the experts with a reliable tool with repeatable results that will obviate much of the time consuming and tedious part of the task.

A second, related, task concerns the grouping of targets. A large feature, for instance a large coral reef, may be represented on a chart by its outline instead of a dense collection of individual small target points (Chart No. 1, NOAA and NIMA, 1997). Again this task depends on human expert judgement on whether and how the targets are to be combined. We aim to develop a procedure that will post-process the identified targets to form groupings according to their proximity to each other. This intermediate grouping will facilitate, where appropriate, the replacement of member targets in a group with an outline polygon around the group. In the following, we will focus primarily on the contracted task of target detection, but we will also report on our work on target grouping and (designed but not yet implemented) outline plotting.

2. Methodology

2.1 Operational Definition

We met with experts from the Naval Oceanographic Office (NAVO) and the University of Southern Mississippi to draw up a concrete operational target identification criterion for our initial system design. This criterion was expected to evolve based on iterative

experimentation and user feedback. It was decided that the IHO Order 1 survey requirements, as described in the previous section, was the most relevant to the coastal environment of interest to CZMIL. (This part of the requirement can be substituted with the requirements of survey standards of other orders if necessary.) Furthermore, the experts expressed a preference for identifying *all* objects satisfying the IHO height requirement regardless of the horizontal dimensions.

The target criterion articulated by the experts can be informally described as follows. For a selected area, compute the average of all depth measurements within that area, and compare this average to the depth of the shallowest point in the area. This shallowest point is considered a target if the difference in depth is (1) more than 2 meters where the average depth is not greater than 40 meters; or (2) more than 10% of the average water depth otherwise.

2.2 Algorithmic Challenges and Corresponding System Design Elements

We will highlight two of the challenges to automating the above target criterion, the first concerning the fuzzy selection of a potential target and reference area, and the second concerning the management of computational resources when dealing with a large volume of data. These challenges inform our system design.

Selection of Reference Area

Key to the success of target identification is the application of the specified target criterion to a well selected area around a potential target for computing the reference average depth. Experts determine the size, location and orientation of this area by “eyeballing” the bathymetric map. The human is superior to the machine in discerning relational structures in an image, and the expertise required for this eyeball method is difficult to capture algorithmically.

Without expert eyeballs we approximate their method by employing a reference area of varying radius. More specifically, for a given potential target, we compute several average depths with reference to a series of concentric areas approximately centered on the potential target location. This provides a more robust target detection criterion as opposed to using a single fixed reference area size. Currently a point is considered a target if it satisfies the target criterion with respect to any one of the candidate reference areas. This method can be further refined to take into account the placement of the potential target under consideration with respect to the reference area, the shape and orientation of the reference area, as well as the functional combination of results from multiple reference areas. Initial testing and expert feedback suggest that the results could be further improved by some of these modifications, which we will discuss in more detail in Sections 6 and 8.

Data Management and Data Structure

The second challenge arises from the large volume of data typically collected in a survey. Our data set covers roughly an area of 100 km² and contains over 300 million data points. Even though for our purposes each data point is represented by just three values: latitude, longitude and depth, this amounts to approximately 10GB of data in ASCII.

Ideally we would like to draw a circle (or a set of concentric circles) centered around each data point and compute the average using data points that fall within this circle. Computationally this requires, first of all, keeping at least all data points in the vicinity of the potential target in memory, and secondly, computing the distance between pairs of data points many times. We therefore instead grid the data using a relatively fine grid (approximately 1 meter by 1 meter for each cell), and consider only reference areas whose boundaries coincide with grid cell boundaries. Gridding alleviates both of the above computational difficulties. Given that we know to which grid cell each data point belongs, it takes constant time to determine whether a point is within a reference area of a potential target. Furthermore, by precomputing and storing a few statistics we can recover all the information needed for target detection without having to keep in memory all the data points involved. For instance, using efficient mathematical manipulations the average depth of an arbitrary grid-aligned reference area can be reconstructed from the average and number of points of each member cell.

Note that the “radius” of the reference area should more accurately be called the “half length of one side of the square”, since the gridded reference area is square (or more generally rectangular), but we will continue to refer to this element as the radius (of the square).

2.3 Algorithm Description

Input and output are in the form of ASCII files. The input file is comma-delimited, each row denoting one data point. A data point consists of a triple, y, x and z, indicating (in decimal degrees) longitude, latitude and (in meters) depth. The region bounded by the extremal x, y values (assumed rectangular) is divided into cells of approximately 1 meter by 1 meter; fractional cells at the high value edges of the specified bounding box are excluded. The cell sizes along the x and y dimensions are independently adjustable so that the cell may be rectangular instead of square.

The data points are read in sequentially from the input file(s) and each grid cell is populated with the following information.

- Shallowest point in the cell (technically only the depth of the shallowest point is needed, but we also record its latitude and longitude values for verification purposes). Where there are multiple shallowest points, only one is recorded.
- Number of points in the cell.
- Average depth of all points in the cell.

Note that these cell statistics are incrementally updated as data points are read and processed; the data points themselves are not kept in memory beyond each update.

Data points that are above sea level (taken to be at zero depth) are dealt with separately and not entered into the grid. This is to avoid marking topographic structures inland as targets (as they rise above the average “depth”) as well as to avoid skewing the average depth of the water portion of a cell when the cell lies on the water-land boundary. Small groups of above water points surrounded by water *may* be marked as targets; the experts have not decided on the appropriate treatment yet. For the moment these points are marked and returned separately from the list of under water targets.

The following information is computed and kept with each grid cell after the above statistics have been collected from the input data points.

- Primary target ID: uniquely numbered and at most one per grid cell, denoting whether the shallowest point in the cell is considered a target.
- Radii and reference averages: a list of radii and their corresponding average depths of reference areas with respect to which the shallowest point in the cell satisfies the target criterion.
- Target group ID: target groups are uniquely numbered and targets in the same group according to a proximity criterion are assigned the same group ID. (See the section on “Target Grouping” below.)

For each grid cell i , and for each candidate radius n , we test whether the depth of the shallowest point satisfies the IHO Order 1 survey criterion for targets with respect to the average depth of all the points within n cells of cell i . If so, we assign a new primary target ID to cell i and record the radius and reference average involved. The cell is tested for targets using all candidate radius values; each successful radius and corresponding reference average are recorded, but the primary target ID is assigned (at most) only once for each cell. All successful radii are recorded such that we may examine the distribution of targets found with each radius criterion. It is envisioned that with sufficient testing and further machine learning techniques, we would be able to identify the best radius or combination of radii to be used, in which case we do not need to keep track or try all candidate radii.

After determining whether each cell contains at least one (primary) target, the input file(s) of data points are reopened and processed a second time to extract all (secondary) targets meeting the target criterion. Secondary targets are those data points that meet the target criterion but may not be the point with the shallowest depth in its cell. A second pass is made to collect the secondary targets such that we can avoid having to keep all data points in memory at all time. During the second pass, for each incoming data point, if the cell it belongs to contains a primary target, we check the list of recorded reference averages of that cell to determine whether the data point satisfies the target criterion. Note that only cells with a primary target needs to be checked, and of those cells only the recorded reference averages need to be compared against the data point in question, since

no point can be considered a target at a certain radius unless the shallowest point in the cell can be considered a (primary) target at that radius in the first place.

The grid information can be saved to a file and re-imported after various stages (gridding, target identification, group identification) such that processing can be interrupted and resumed at intermediate points. The targets are output to a text file each with its y, x, z values together with its secondary target ID and group ID. With minimal modification supplementary information such as cell indices, effective radii and reference averages may also be output together with the target information.

2.4 Adjustable Parameters

Several parameters are user-adjustable to control the behavior of the target detector to some extent. These include, in addition to file names of various input and output files, the following parameters.

- Bounding box of the area to be processed, in terms of maximum and minimum latitude and longitude values, in decimal degrees.
- Resolution or size of each grid cell: the x and y dimensions can be specified independently for rectangular cells. The resolution is currently accepted in units of decimal degrees.
- Radius range to be considered: the minimum radius, the maximum radius, and the step increment between the minimum and the maximum, all in terms of number of grid cells. Note that the radius is defined as the number of cells away from the current cell, that is, a radius of 0 means that only the points in the current cell are included. Note also each increment of 1 increases the number of cells by 1 symmetrically in each direction of the grid; for instance, a radius of 0 encompasses 1 cell (an area of 1 cell by 1 cell), while a radius of 1 encompasses 9 cells (an area of 3 cells by 3 cells). Thus it is important to grid the data finely—not such that we may consider reference areas of minute sizes, but such that we may have better control over the successive sizes of reference areas to be considered.
- Grouping radius: targets less than “grouping radius” apart are grouped together. The grouping radius is specified in terms of number of grid cells, in any direction. (See Section on “Target Grouping” below for more details.)
- Grid loading and saving: boolean flags that determine whether a grid exists for loading and whether the current grid should be saved to a file. If not loading a grid, it is expected that one or more input files of data points will be supplied.

Optionally, the data points can be saved into files grouped by cell regions. Given the typical size of the grid (in terms of the number of cells), it would be more realistic to save regions of consecutive cells, for instance, 10 cells by 10 cells, into each file. These files are not used for target detection; they are intended to provide more quickly accessible software-independent data for post-process visualization and verification. It is

recommended that this option be turned off for deployment, due to the amount of time required for the file I/O operations.

2.5 Target Grouping and Polygon Outline Plotting

Target grouping is an intermediate step to deriving polygon outlines for groups of closely spaced targets. At the moment targets are grouped using a proximity metric similar to the one for defining areas for computing the reference average: given a grouping radius r and a target t , any target within r cells of target t is assigned to the same group as t . This is a transitive condition: all targets within r cells of target t are assigned to the same group as t , and all targets within r cells of any one of those “neighbor” targets that are themselves within r cells of target t are also assigned to the same group as t . In other words, a target may be more than r cells away from some of the targets in the same group, but it must be less than this distance away from at least one other target in the group (unless it is a singleton group). Conversely, no two targets from different groups can be less than the given grouping radius apart.

Once the targets are grouped, we can decide whether to replace all the members within a group with an outline polygon around the group. Drawing an outline polygon amounts to computing the convex hull of the member targets. We have an algorithm for computing a series of outline polygon edges by first determining member target points that are located on the convex hull. This part, outline plotting, has not yet been implemented since we need to first confer with the charting experts to obtain feedback regarding the grouping criterion and then to determine an acceptable criterion for selecting the appropriate target groups for outlining.

3. Implementation and Testing

The target detector was implemented in C++, for compatibility with the in-house developed software for related tasks at NAVO. Test data was procured from South Florida Test Facility by JALBTCX on our behalf. The data consist of LIDAR and sonar measurements, in the area around Port Everglades on the east coast of Florida. The data format is as described in the above section (“Algorithm Description”). For proof of concept we chose a subset of the given area with diverse characteristics and some known topographic and hydrographic features. The coordinates of the selected area is as follows:

min X: -80.11; max X: -80.08;
min Y: 26.08; max Y: 26.11.

This area contains some shallow water, a part of a reef, some above water offshore features, the inlet and dredged channel leading out of Port Everglades, as well as some topographic features along the coast. The square area spans approximately 11 km^2 , and the data subset contains approximately 22 million points.

The area is gridded into 3323 cells by 3323 cells, each cell measuring approximately 1 meter by 1 meter. The radii considered for target identification range from 0 to 10, with a step increment of 1 (cell). The target grouping radius is 5 (cells); this value is partially derived from the IHO survey standards which allow for a horizontal (in)accuracy up to 5m + 5% of water depth for Order 1 surveys.

One complete execution of the program, from gridding through to output of target and target group list, including saving certain diagnostic information, takes approximately 2 hours on a Power Mac G5 Quad (2.5GHz x 4) with 2 GB of memory running Mac OS 10.4.7. To start up and close down at intermediate points, loading the grid takes approximately 5 minutes and saving the grid takes approximately 3 minutes. Note that the tasks of target identification and target grouping themselves do not depend on the number of input data points, but rather the size or granularity of the grid. For a given area, the finer the grid, the more number of grid cells, and the longer it takes to look for targets in all of them. The initial gridding and the final output of secondary targets do depend on the number of input data points; however the data points do not place a strain on the available memory since they are processed sequentially and do not need to be kept together in memory all at once.


4. Results and Analysis

In the selected area, 1492 targets, clustered into 87 groups, were identified. In addition, two groups of offshore above sea level points were also isolated. These lists have been sent to NAVO experts for examination.

We also performed some in house analysis and validation of the software system and the identified targets and target groups. A small suite of post processing routines were implemented in matlab to visualize and analyze the results. These analyses allowed us to iteratively modify the algorithm to improve the target detection capability.

4.1 Overview Plot

Figure 1 gives an overview of the area selected for testing. Data points that are above sea level (depth 0) are colored according to their elevation values. Below sea level points are not plotted so that the targets identified may be seen more clearly. Targets are marked by their target IDs. The targets can be roughly equally divided into two broad groups, one group marking the sides of the channel coming out of Port Everglades, and the other marking some submerged offshore features in the southeast quadrant of the plot. The channel, although easy to identify by sight, is useful for verification of system performance, since the dredged bottom and the banks of the channel are relatively uncluttered and easy to inspect visually.



Note the two (barely visible) isolated blue spots in the southeast quadrant of the figure, near the bottom. These are two groups of offshore above sea level points that were identified separately from the underwater targets. The same target detection procedure can be used to detect such above water features, but NAVO experts are undecided about whether these features should be tagged as targets.

4.2 Detailed Individual Plots: Symbols and Notations

We plotted selected sub-areas, some with targets and some without, from different perspectives to better analyze the terrain and the appropriateness of the target identification criterion. Supplementary information relevant to the target detection task is also included in the plot whenever possible. First let us explain the symbols and notations used in the plots.




Figure 2(a). Center Cell with Target 1313: 3D View

Each sub-area is shown using a series of four plots, each from a different perspective:

- (a) 3D View: 3-dimensional view of latitude, longitude and depth;
- (b) Overhead View: 2-dimensional view from directly overhead;
- (c) Latitudinal Side View: 2-dimensional projection onto the latitude-depth plane; and
- (d) Longitudinal Side View: 2-dimensional projection onto the longitude-depth plane.

Figures 2(a)-(d) are a sample series of plots. In each of these plots, targets are marked by their target numbers; all data points within the plot area are colored according to their depth value (a reference color bar legend is included on the right edge of each plot); and each data point in the cell in the center of the plot is in addition marked by a red asterisk.

Figure 2(a) is a 3D view of a selected area. The vertical axis is the water depth, in meters, of the data points. The other two axes are parallel to the latitude and longitude respectively. The tick marks of these two axes are given in cell units from the center of the plot, usually from -10 to 0, and then from 0 to 10. Each number denotes the distance in terms of the number of grid cells between the tick mark location and the center of the plot. Note that 0 is repeated twice on each axis: the cell in the center of the plot is the area within the intersection of the four lines marked “0” on the two axes. Thus all distances are measured relative to the center cell, which is the cell of primary focus in

each plot. In this case, there is a target, numbered 1313, in the center cell. In addition, there are four other targets in the vicinity of the center cell: targets 952, 953, 954 and 955. (They can be seen more clearly in some of the subsequent figures plotted from various different perspectives.)

This seemingly unusual labelling is devised to facilitate the encoding of radius and reference average information used to identify targets among the points in the center cell. There are 10 concentric squares in the plot, all centered at the center cell, and each with a radius ranging from 0 to 10. (See Figure 2(b) for a clearer overhead view of the squares.) The area bounded by each of these squares of a particular radius is the area used to compute the reference average at that radius. The value of the reference average itself is denoted by the height (depth value) at which the square is plotted. Furthermore, a square is drawn in red if the shallowest point in the center cell satisfies the target criterion with respect to the reference average at this radius; otherwise the square is drawn in blue. It can be readily observed that all blue squares are above (in depth) all red squares. (See the side views Figures 2(c) and (d) for a clearer depiction.) In Figure 2(a), there is only one red square, at radius 2, indicating that target 1313 in the center cell satisfies the target criterion only when considering the reference average computed at radius 2 but not at any other radii from 0 to 10.

The same information is plotted from different perspectives in Figures 2(b)-(d). The overhead view (e.g. Figure 2(b)) is in general useful for gleaning radius information and also for orienting the distribution of targets and other data points. The latitudinal side view (e.g. Figures 2(c)) is the projection obtained by collapsing the longitude dimension; similarly for the longitudinal side view (e.g. Figure 2(d)). The two side view plots are helpful for visualizing the terrain, especially the relative depths of different points. Note that in the side views each square denoting a reference average value avg at radius r is collapsed into a line segment that extends from $-r$ to $+r$ parallel to the bottom axis at depth avg .




Figure 2(b). Center Cell with Target 1313: Overhead View




Figure 2(c). Center Cell with Target 1313: Latitudinal Side View




Figure 2(d). Center Cell with Target 1313: Longitudinal Side View

4.3 Analysis of Some Examples

Figures 2(a)-(d) were chosen to illustrate the symbols and notations used in the plots because that particular area contains relatively few data points and therefore the plots are relatively uncluttered. We will come back to this series of plots later in the next section when we discuss them in relation to the treatment of ledges and other sheer drop offs. But first let us look at two other series of denser plots illustrating some configurations of points commonly found in our test area.

Example 1

Figure 3 depicts an area with fairly uniform depth except for a small strip of points whose depth measurements are well above the rest of the points in the surrounding area. All the points in that strip of shallower points were marked as targets. From the clear difference between the shallower points and the uniformly deeper bottom, we can expect that the shallower points would be considered targets with respect to a wide range of radii. This is in fact the case; the points satisfy the target criterion at all radii considered (from 0 to 10).




Figure 3(a). Center Cell with Targets 612, 615, 616: 3D View




Figure 3(b). Center Cell with Targets 612, 615, 616: Overhead View




Figure 3(c). Center Cell with Targets 612, 615, 616: Latitudinal Side View




Figure 3(d). Center Cell with Targets 612, 615, 616: Longitudinal Side View

Three targets were identified in the center cell. Target 612 was recorded as the primary target of the cell, and targets 615 and 616 were secondary targets identified during a second pass over the input data points. The target numbers are not necessarily consecutive simply due to the order of presentation of the points in the input files, but this does not distract from the functional role of the targets.

The red colored points down among the blue colored points are points (in this case non-targets) belonging to the center cell. Recall that all data points in the center cell are marked with a red asterisk in the plots. The red asterisks may seem distracting in these plots, but they are helpful in locating the points in the center cell in certain plots, especially when the shallowest point is not a target. An example is Figure 6(b). Without the red asterisk marking, it would not be obvious to find the center point we need to focus on in that plot.

Example 2

The data points in the next series of plots are not as uniform, but they nonetheless exhibit a surface trend. In Figure 4, the bottom slopes down in the general direction from west to east (away from the shore). Again there is a small set of points much shallower than the surrounding bottom. These shallower points are aligned roughly in the north-south direction, although their depth values indicate that they are not necessarily closely strung together but rather there exist several “strings” of shallow points.

The data points in Figure 4 cannot be as intuitively categorized as those in Figure 3. Due to the more varied terrain, only the shallowest of the shallow points satisfy the target criterion. The radii at which the targets in the center cell successfully qualify range from 4 to 10. By examining the plots, especially Figures 4(a) and (c), we can see that the center cell is located near the “inflexion point” of the slope, that is, near the center cell the slope gradient starts getting steeper towards the deeper water. Thus, at smaller radii (< 4), the smaller reference areas may not include enough of the deeper points, whereas as the radius increases and therefore the reference area expands, the additional points on the deeper side of the reference area outweigh the additional points included on the shallower side.

Three targets are identified in the center cell of Figure 4. Just as in Figure 3, one of them is a primary target of the cell and the other two are secondary targets. Even though the secondary targets are not as shallow as the primary target (Figure 4(d)), they nonetheless are shallow enough to satisfy the IHO standards with respect to the same reference averages that qualify the primary target.




Figure 4(a). Center Cell with Targets 168, 169, 170: 3D View




Figure 4(b). Center Cell with Targets 168, 169, 170: Overhead View




Figure 4(c). Center Cell with Targets 168, 169, 170: Latitudinal Side View




Figure 4(d). Center Cell with Targets 168, 169, 170: Longitudinal Side View

5. An Extended Analysis

The sample plots examined in the previous section all to some extent conform to our expectation, or at least there is a reasonable explanation in each case supporting the target detection mechanism. We will now turn to a more extended series of plots, all near each other in the area in or along the channel out of Port Everglades. As mentioned before, the channel provides a useful ground for testing and validation, since we know approximately the contour and structure of the channel.

In the following we will only show two views of each series of plots, the overhead view and the longitudinal side view. These two views are the most informative in the area under study, since the channel runs almost straight out in the west-east direction.

5.1 A Target

First of all let us re-examine Figure 2. This is an area on the south side of the channel. From Figure 2(d) we can see that target 1313 is approximately half way on the slope that constitutes the south edge of the channel. The four other targets visible in the plots in Figure 2 are all better classified as being on the edge of the bank rather than in the channel. (Note that all these points are below sea level. The dredged channel is at a depth of 12-13 meters whereas the undredged surrounding area including the channel banks is at a depth of 5-6 meters.)

Target 1313 satisfies the target criterion only at radius 2. By examining Figure 2(d) we see that the radius 2 reference area maximizes the ratio of deeper (channel bottom) to shallower (channel bank) points. For reference areas with radii larger than 2, the denser points on the bank lead to a significant increase in the average depth. For radii less than 2, there are not enough data points in the reference areas to produce a sufficiently large difference between the target point in the center cell and the computed average. For example we can see in Figure 2(b) that the reference area at radius 0 includes only 1 data point—target 1313 itself, and therefore the average at this radius must be identical to the depth of the target point.

5.2 Another Target

Figure 5 shows an area that is slightly shifted from that shown in Figure 2. In Figure 5 the plots are centered at the cell with target 952, which can also be seen on the channel bank in Figure 2. In this case, target 952 is classified as a target at radii 2, 5 and 8. The same kind of forces are at work between the group of deeper points in the channel proper and the group of shallower points on the channel bank. This gives rise to a flip-flopping behavior of the target detector as the target criterion is alternately satisfied and unsatisfied with increasing radius.




Figure 5(a). Center Cell with Target 952: Overhead View




Figure 5(b). Center Cell with Target 952: Longitudinal Side View

Issues Observed from the Analysis

This series of plots illustrates several issues. First, in sparse areas a small number of data points can easily skew the reference average, and therefore the satisfiability of the target criterion. Second, the decision as to whether a data point is a target depends critically on the choice of the reference area with which to compute the comparison average. Here the factor in question is the radius of the reference area. By choosing a different radius and therefore a reference area of a different size, we can obtain a different classification of the same data point. The following plots will further illustrate additional factors that can influence the behavior of the target detection procedure via the choice of the reference area. Third, in view of the above issues, it may be helpful to devise and include with each target classification a confidence measure that takes into account the sufficiency and suitability of the data and reference area(s) considered. Let us analyze another series of plots and then we will summarize the issues in Section 6 and suggest some possible improvements to the target detector to address these issues in Section 8.

5.3 A Non-Target

The last set of plots we consider here is still in the same area. This time the plots are focused on a center cell without any targets. Refer to Figure 6(a) to see the relative position of the center cell to, for example, target 1313 which is in the center cell of the plots in Figure 2.

The non-target point in question here, in the center cell in the Figure 6 plots, is the green colored point almost directly above target 1313 in Figure 2(d). Recall that with the longitudinal side view, all the points along the same longitude are collapsed into the same column, even if their latitude values differ substantially. In Figure 2(d) even though it appears that the green non-target is in the same cell as target 1313, the overhead view in Figure 6(a) shows more clearly that they are 4 cells apart along the latitude dimension. (In Figure 6 this non-target point is now marked with a red asterisk instead of a green dot.)

Figure 6(a) also shows that this non-target in the center cell and target 1313 are about the same distance away from the shallower points on the channel bank to the south. Figure 6(b) shows that the two points are also of similar depths; the non-target is even slightly shallower than target 1313. In other words, these two points have very similar characteristics, yet only one of them, the deeper of the two, is classified as a target according to the target detection procedure.

By comparing Figure 6(a) and Figure 2(b), some subtle differences between the two situations can be observed. The points closest to the non-target are slightly shallower than the ones closest to target 1313. These closest and shallower points weigh in at every radius considered, which results in shallower reference averages at all radii compared to the corresponding averages for target 1313. The average depth at radius 1 for the non-target is even shallower than the non-target point itself.




Figure 6(a). Center Cell with No Target: Overhead View




Figure 6(b). Center Cell with No Target: Longitudinal Side View

6. Issues

In Section 5.2 we mentioned some issues regarding the target detection procedure. Here we provide a more detailed discussion of these and other issues noted from our analysis.

Sparse Data Area and Confidence Measure

The last series of plots again re-emphasizes the importance of choosing a suitable reference area. A few data points, with slightly different depths and located slightly differently, could influence whether a point qualifies as a target. Obviously in dense areas with many data points, as in Figures 3 and 4, this problem is greatly diminished. However, it would still be helpful to develop a measure for the confidence with which the target classification is made, based on the number, location and other characteristics of the data points in the area.

Non-Uniform Classification and Robustness of Varying Radius

A related issue that needs to be explored concerns the robustness of the varying radius method. Points that are clearly shallower than the surrounding area, for example those in Figures 3 and 4, typically satisfy the target criterion homogeneously at a wide range of radii. Target 1313 satisfies the target criterion only at radius 2, and target 952 satisfies the criterion at radii 2, 5 and 8. We are not sure at the moment whether these points should be classified as targets or non-targets (awaiting expert assessment). However, the singular or uneven classification may serve to indicate that the targets so found may possess some questionable characteristics, or in the extreme case, the uneven classification may be used to declare these points as non-targets. This suggests that radius inhomogeneity might serve as an indicator of confidence in assigning target status to a point.

Flexible Placement and Shape of Reference Area

Another issue concerns the placement of the reference area relative to the candidate data point under consideration, or more precisely, the candidate cell in which the data point under consideration is located. From Figures 2(d) and 6(b), it is clear that if we shift the reference area such that it defines a square that extends from the edge towards the middle of the channel, the reference area will include (in addition to the candidate point) mostly deeper points that are at the dredged bottom and few of the shallower points atop the channel bank. This will allow us to classify as targets both target 1313 and the point in the center cell of Figure 6, and it will allow us to do so without having to be too precise about the radius chosen, since we can expect that both points would be classified as targets over a fairly wide range of radii.

Conversely, if these points should not be considered targets, both of them will be classified as non-targets if we extend the reference area from the candidate cell towards

the channel bank (as opposed to towards mid-channel as in the previous case). The reference area so constructed will include mostly the shallower points on the bank and the resulting reference average will not qualify the two points as targets.

As mentioned before we are consulting with the experts regarding whether these points should be classified as targets. However, given the similarity of their characteristics, we believe they should be treated uniformly, either both as targets or both as non-targets.

More generally, instead of requiring the reference area to be centered at the candidate cell, we may allow the reference area to be shifted asymmetrically in various ways as long as it still includes the candidate cell under consideration. We may also allow the reference area to take a shape other than a square, although for efficiency reasons rectangular areas (including squares) are preferred. Assessing target status with shifts of the cell containing the candidate point to each of the four corners of the reference area is essentially taking a quadrant of a larger, but still symmetrical, reference area. Shifting to one corner or another instead introduces a directionality to the reference area with respect to the candidate point. At issue is finding effective criteria for choosing a direction, or for using a combination of directional results in target selection.

Ledges, Drop offs, Channels and Other Features

This analysis leads to yet another issue. Our definition of a target includes any object or point that satisfies the IHO criterion on *any one* side of the object. Using this definition, points along a ledge or other steep drop offs, for instance a channel bank, are considered legitimate targets. We are working with NAVO experts to refine this definition. Should they require that an object be classified as a target only if it protrudes on *all* sides instead of just one, the “shifting box” strategy discussed above, combined with a varying radius, would be a viable method for identifying targets according to this revised definition.

7. Target Grouping

We also implemented a procedure for target grouping, which takes as input the list of targets identified by the target detection procedure and assigns the same group ID to targets that are less than a specified distance apart. (See Section 2.5 for a description of the grouping procedure.) Figure 7 is an overhead view of the target groups in the test area. Targets within the same group are plotted using the same color. There are a total of 87 groups for 1492 targets. Due to the size limitation of the figure, the color difference between adjacent groups may not be clearly visible.

As previously discussed, the groupings can be further processed to obtain polygon outlines of selected groups of densely spaced targets. We have in place an algorithm for this step but it has not yet been implemented.




Figure 7. Target Groups. Targets in the same group are of the same color.

8. Potential Further Work

The issues described in Section 6 identify potential further research of practical importance. We will not re-elaborate the issues here but will instead focus on possible improvements and directions of future research.

8.1 Programmatic Variations

Desymmetrize the Radius and Reference Area

Currently all reference areas are centered at the candidate grid cell under consideration. The only variable is the size of the reference area, determined by a symmetrically applied radius. As discussed in Section 6, it may be helpful to consider other placement and shape of the reference area. The more fine control of the radius and reference area will allow us to, if desired, exclude ledges and identify as targets only those points that protrude on all sides.

Shifting the location of the candidate grid cell within the reference area creates a dimensional asymmetry, as noted above. A corner location of the grid cell, for example, creates a reference area approximating a quadrant of a reference area with a radius twice that of the radius from the center of the reference area. We might require that a point be counted as a target if placing it at any corner or the center of the reference area meets the IHO target criterion. Alternatively, we might place the candidate cell at the midpoint of each of the four edges of the corresponding reference areas. The variation scales linearly in the number of different placements.

Medians Instead of Averages

Averages can be skewed by a few extremal points—a hole, for example—while medians generally provide a statistic more robust to outliers. Medians are, however, computationally demanding. Merely to compute the median of each grid cell (with a single data pass) would require storing every data point in memory. These points have to remain in memory through much of the processing, since medians of cells cannot be aggregated without consideration of all point depths within each cell. (Contrast this with the computation of averages. The average over several cells can be composed, with supplementary information, from the averages of the individual cells, without having to go back and examine the original data points.) If, however, target selection with medians should prove more satisfactory than with averages, the increase in performance might justify the rather steep increase in computational costs. The problem is natural for parallelization, which could allow considerable time (but not memory) savings.

Parameters for Grid Cell Size, Radius Range and Radius Step Increment

For testing we have used a grid cell size of approximately 1 square meter, and a radius range of 0 to 10, with a step increment of 1 (that is: 0, 1, 2, ..., 10). Larger grid cells will reduce computational time and memory requirements, but we will have less control over the radii that may be chosen, and more important, the step increment that defines the minimum difference between one reference area and the next possible larger reference area. Using concentric reference areas centered at the candidate cell, as we did, a 1m x 1m grid cell size will allow for reference areas of size 1m x 1m, 3m x 3m, and so on, whereas a 2m x 2m cell size will inherently only allow for areas of size 2m x 2m, 6m x 6m, and so on. It is not possible to define an area of, for example, 5m x 5m with 2m x 2m grid cells.

These parameters, grid cell size, radius range and radius step increment, are all adjustable in the program, but we need to determine the optimal values for these parameters for different situations. One approach is to apply machine learning techniques, which will be discussed in Section 8.3.

Expert Assessment and Switch for Different Applications

Many of the above variants introduce extra relational criteria for “target,” with consequent costs in memory, run time, or both. Consultation with experts and analysis of

comparative results should indicate whether and which of these extra computational costs are worthwhile.

It seems possible that there is no firm consensus, only general guidelines, among experts on all of these issues about target criteria, and even that each expert is not able to completely articulate his/her own selection criteria. For example, considerations drawn from the past experience of an expert but external to the database of x , y , z values sometimes justifiably influence decisions about whether a point should be marked as a target. In addition, mine experts and ecology experts will have different requirements for targets. In that case, a possible procedure is to implement switches for several target criteria that can be activated depending on the application. Metadata recording the criteria used and other supplementary information can also be automatically recorded.

8.2 Refining and Redefining Target Criteria

The target specification mandated by IHO survey standards leaves much room for interpretation. The task of target identification relies heavily on human expertise. The automated target detection procedure we have implemented corresponds roughly to the procedure described to us by the experts, but our analysis and their feedback indicated that a refinement of the target criteria would more precisely match their treatment of different hydrographic and also topographic features. These include for instance above sea level objects, ledges, mesas, and donut-shaped features such as underwater volcanoes.

Currently all data points above sea level are processed separately and are not included as input to the main target detection procedure. However, comparatively small above water regions that are surrounded by water might arguably be considered targets. The experts are undecided about the treatment of such features, but if needed, a pre-process might check, for points above water, whether their neighboring cells according to some user specified radius, or radial sector, are above or below sea level, and include such points as input into our procedure accordingly. The same target detection procedure can be applied to pick out both targets above and below sea level.

The present algorithm uses depth 0 as the sea level to identify land data points. This criterion may be unsatisfactory where there are large tidal effects covering and uncovering rocks for example. Investigation requires consultation with experts and assembly of a body of easy and hard cases for land/water demarcation.

Here are some other features that deserve closer examination. Internal points in submerged islands—underwater mesas—will not count as targets by the present criteria, and perhaps they should be. Points on ledges—long running sequences of contiguous points whose neighboring points in one direction are approximately equi-shallow, and whose neighboring points in the opposite direction are much deeper—whether natural or from dredging, might be considered non-targets. Algorithmic modification may be required to ascertain that these features are treated uniformly according to expert judgement.

8.3 Applying Machine Learning Techniques to Refine Target Classification Criterion

The most important consideration is to avoid false negatives, since missed targets may pose a threat to navigation and other activities. In contrast, false positives, though not desirable, may be more acceptable (if not too excessive), since the human experts could examine and eliminate them at their discretion. Accordingly, it would be useful to have an expert compile a small list of hard and critical cases—cases meeting two criteria: (1) it was difficult to recognize the target; and (2) if not identified, the target would be a genuine threat.

Given a larger set of targets and non-targets identified by an expert, automated machine learning procedures, probabilistic decision trees for example, can search for the combination of grid sizes, radii, and placement shifts of reference areas that optimizes fit to expert judgement. This will allow us to investigate issues such as whether the procedure would be more robust if we required a point to simultaneously satisfy the IHO target criterion at multiple radii. With such a database available, results of alternative criteria can also be compared through informative true and false positive trade-offs indicated by Receiver Operating Characteristic (ROC) curves.

8.4 Confidence Measure

In sparse areas, the designation of targets can be easily swayed by a few strategically located data points. There will also be borderline cases in which experts do not necessarily agree or in which an expert is uncertain of the appropriate designation. We believe a confidence measure would be a helpful inclusion into the overall design, to indicate how well a data point conform to the prototypical target. At its simplest the confidence measure can be the inverse of the number of points in the reference area. More sophisticated measures will also take into account the relationship between the points and other survey characteristics.

In cases where uncertain or non-uniform expert judgement is correlated with variation in target status according to the specific radius selected for the reference area, it may be possible to incorporate the radius, or directional dependence, or both, into a measure of uncertainty to be reported as metadata.

8.5 Groupings and Outlines

We have implemented a procedure for automated target grouping, but we have not implemented the algorithm for drawing polygon outlines around these groupings, which amounts to computing the convex hull given a set of data points. At issue is whether such features would be of value sufficient to warrant their implementation. A further

research issue concerns the choice of appropriate distance parameter for grouping, and the introduction of an optional switch for that parameter, together with corresponding metadata for the output.

8.6 Optimization

We have endeavored to implement a reasonably efficient algorithm, but we do not claim that the program is fully optimized. The current implementation allows for rapid prototyping to facilitate the development and testing of algorithmic modifications. Once the algorithm has stabilized, we can focus on optimizing the performance. Possible time and memory improvements include, for instance, data buffering, partitioning the reference area to isolate components that do not need to be recomputed, reordering the execution sequence and, of course, minimizing the computation and storage of diagnostic information. For instance, currently all radii are attempted and results recorded even though a point is classified as a target if it satisfies the target criterion at any one radius. We could cut short the computation as soon as we find a first radius that allows the point to be classified as a target, but we purposely compute and record results for all radii such that we may analyze the performance using machine learning techniques such as those described in Section 8.3.

In addition, if very large areas are to be analyzed all at once, depending on the algorithm, parallelization by geographical sector may be valuable, and routines need to be optimized to synchronize results obtained from parallel processes and for regions where sectors processed in parallel overlap.

References

International Hydrographic Bureau, 1998. IHO Standards for Hydrographic Surveys Special Publication No. 44, 4th Edition.

Department of Commerce (National Oceanic and Atmospheric Administration) and Department of Defense (National Imagery and Mapping Agency). Chart No. 1: Nautical Chart Symbols, Abbreviations and Terms, Tenth Edition, November 1997.