




Automated Side-scan Data Analysis

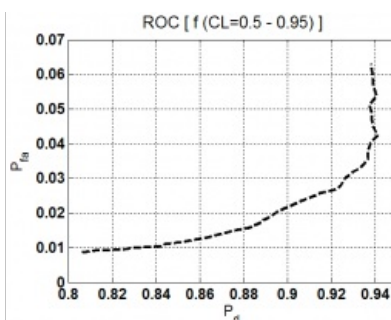



Side-scan sonar is well accepted as a tool for detection and visualisation of underwater objects and is widely used in support of safety of navigation surveys. Although side-scan sonars produce excellent images, these images are often difficult and time-consuming to review and analyse. These images typically contain spurious features and artefacts due to acoustic phenomena, such as refraction, and bottom characteristics, such as sand waves that result in a large number of false contacts. This article describes a research and development effort at Science Applications International Corporation to develop and apply advanced processing techniques in bottom target detection and image classification to support the automation of side-scan sonar data processing.



Predicted Class	True Class		
	target	clutter	uncertain
target	1418	1348	0
clutter	110	21850	0
uncertain	0	0	0
TOTAL	1528	23198	0

$$P_d = 0.92801 \quad P_{fa} = 0.058108 \quad CL = 1$$



Shallow-water hydrographic surveys typically involve the collection of concurrent bathymetric and side-scan sonar data. Recent developments in bathymetric data processing within the hydrographic community have allowed for a decrease in the labour effort to process data while at the same time providing potentially superior data products. Over the past several years, the Acoustic and Marine Systems Operation within Science

Applications International Corporation (SAIC) has been involved in an internal research and development project to reduce the labour required to process the other sonar data type side-scan sonar.

The Marine Science and Technology Division of SAIC, located in Newport (RI, USA), performs a wide range of marine survey operations, with hydrographic surveys representing a major component of these. Most hydrographic surveys that SAIC has performed since 1994 (when the National Oceanic and Atmospheric Administration [NOAA] first contracted out hydrographic surveys to private industry) have required 200% side-scan coverage, with resulting multi-beam bathymetry coverage. Although bathymetry processing is more automated than manual, due in part to the introduction of the Navigation Surface, imagery processing is almost entirely manual.

Two-step Process of Data Review

SAIC procedures for side-scan processing currently entail a two-step process of data review to mitigate missed targets. The first round of data scanning is performed by a junior hydrographer, and the subsequent verification of the data is performed by a more senior hydrographer. To reduce the labour required to process the hydrographic imagery data, and therefore improve the efficiency of imagery processing, SAIC determined that an automated process for object/target detection was needed for data processing.

The project was to develop an automated system capable of identifying natural as well as man-made targets (objects) on the seafloor with dimensions of one cubic metre and larger, regardless of strength of return and/or shape. Requirements for false detection rates were set at 5% of the total number of targets detected, with a similarly ambitious desire to identify 95% of the objects that met the object's minimum dimensional specifications.

At the prototype level, the problem was to design and implement a suite of software components to automatically detect bottom targets and then classify those detections as either legitimate or false detections to reduce the number of automatic detections a hydrographer would need to review. A generalised process flow is as follows:

1. Automatically detect bottom targets from test side-scan sonar data sets.

2. Review automatic detections to provide a ground truth data set (performed by trained hydrographers).
3. Extract parameters from each detection (for example, shape and texture).
4. Automatically measure the length, width and height of each detection.
5. Train neural net classifiers on reviewed subsets of detections to classify detections as legitimate or false alarms.
6. Apply the trained neural net to entire reviewed data sets and provide uncertainty metrics.

Although the project at its most practical level was to develop a system that would decrease the labour associated with reviewing side-scan sonar data and provide efficiency to data processing for hydrographic surveys, this technology has wide applications in areas such as harbour security, munitions detection and other such applications.

To date, the prototype software has been completed and applied to a large volume of test data from hydrographic surveys performed by SAIC for NOAA. The test data were drawn from two surveys in similar oceanographic environments using different side-scan sonars. Both areas were surveyed with a 50m range setting on the side-scan sonar and 40m line spacing, resulting in 200% side-scan coverage. Sheet F (registry number H11241) was a survey of the approaches to Little Egg and Brigantine Inlets off the coast of New Jersey. Sheet H (registry number H11455) was a survey of Holgate to Beach Haven Crest, just north of Sheet F. Sheet H (Figure 1) was surveyed with a Klein 3000 side-scan sonar and consists of 613 XTF files that span over 29 survey days. A total of 24,726 automatic detections from both sheets were comprehensively reviewed by hydrographers to provide a ground truth data set of legitimate targets and false alarms.

Detection

Simply stated, the detection algorithm is designed to find a peak followed by a shadow. The detection processing makes use of a split-window normalisation scheme commonly referred to as Constant False Alarm Rate (CFAR) detection. Conceptually, the algorithm assigns the targets contained in a side-scan image with a score that is proportional to the area of the front of the object, or width by height. Peak and shadow scores are computed individually on all side-scan channels and combined to a target score that triggers detection above a prescribed threshold.

As sand waves often have a signature similar to targets on the seafloor (a strong return on the near face of the sand wave and a shadow) and can therefore result in thousands of false detections, the detection processing includes a 2-dimensional median wave-number filter to suppress sand waves and other background interference. Post-detection processing, such as the application of linear prediction filters, is also applied to the raw data to remove unwanted textures or global trends that affect the parameters and measurements that are extracted from each detection window.

Processing and Measurements

Integration of image processing functions into the processing stream enables image filtering, segmentation, small-object detection and simple feature extraction. The parameter extraction code operates on black and white as well as greyscale images of the peak and shadow scores. A total of 172 parameters are extracted from raw and processed detection images. The resulting image parameter set captures shape, texture and statistical moments of the detected targets. An important step in parameter extraction is the determination of the image processing window within the detection window shown in Figure 3. These windows are based on the automated analysis and thresholding of the score images resulting in elliptical regions, which isolate the target and its shadow. The moment and texture parameters extracted from these windows are used for classifier training and target measurement determination (length, width and height).

Neural Network Training

The design sensitivity of the detector ensures a high probability of detection at the cost of a large number of false detections. However, this does not necessarily degrade the target identification capability when an effective classification scheme is employed. For this reason, the prototype code for the identification of targets uses multi-layer perceptron networks in association with statistical confidence metrics to manage false alarms.

The essential problem in the use of most neural network classification schemes based on error-minimisation techniques is to find a classifier that represents a model of the data as opposed to a memorisation. To achieve this balance, the ratio of training examples (detections) to network connection weights must be large ($>10:1$). To reduce the number of weights and increase the ratio, the parameters for the neural net input are selected to maximise the statistical distance between the real and false contacts (targets) and to minimise the correlation among the various image parameters.

The neural network to be trained was configured with eleven hidden layer nodes and thirty input image parameters. The training algorithm therefore attempts to minimise classification error using a total of 330 (30×11) estimated connection weights. In this case, a randomised sample of 5,000 reviewed detections is sufficient compared with the number of weights expressed as a ratio of about 15:1. Although the optimisation problem is now over-determined, constraints such as these typically increase confidence in the ability to apply a classifier to surveyed data from other geographical regions.

The results after applying the network trained on 5,000 examples to the entire set of 24,726 detections are presented in the confusion matrix shown in Figure 4. This neural network correctly classified 1,418 targets and 21,850 false alarms (clutter). However, it incorrectly classified 1,348 false alarms (clutter) as targets and 110 real targets as false alarms (clutter). Therefore, the classification results of the current single-stage neural network are a 93% probability of detection and a 6% probability of false alarm.

Uncertainty Metrics

The classification of features as targets or clutter is based on network activation values. Statistical analysis of the distribution of network activations for each class allows the identification of uncertain classifications. In the single classification example shown in Figure 5, the network activation (green circle) is plotted along with beta distributions fit to the network activations from the entire training set shown on the right. In this case, the example was correctly classified as a real target. However, in many cases, the classification is on the tails of either distribution, making the decision more difficult. In such cases, the activation distributions can be used either to define the crossover

decision point or quantify the classification in terms of a confidence level.

The activation distributions are also very useful in optimising the false alarm and detection rates. An N-stage network was developed recursively from the uncertain examples resulting from a series of training sessions using progressively larger confidence levels. This generates a receiver operating characteristic (ROC)-type analysis where detector sensitivity is analogous to confidence level. The ROC curve shown in Figure 6 was derived from a three-stage neural network where confidence levels from 0.5 to 0.95 resulted in the probabilities of detection and false alarm shown on the X and Y axes, respectively.

Summary and Future Directions

Although the prototype offers a convincing proof of concept (93% probability of detection and 6% probability of false alarm for a single-stage neural network), future success will depend largely on effective integration and fast implementations of existing software. The extraction of small targets requires the processing of side-scan data at its full resolution, which has demonstrated the need for accelerated processing speeds well in excess of those used in the prototyping phase. Consequently, continued test-bed development for fast and flexible image processing and analysis of side-scan data appears to be critical.

In the coming year, SAIC also plans to implement active learning strategies that allow the system to learn as new data are acquired. This specifically addresses classifier portability, which is the key element in the application of the automated feature-detection system to hydrographic surveys in new regions using other survey instruments.

A culmination of the current efforts in this technology area will be the development and testing of a near-real-time operational system, which will include sophisticated data-visualisation schemes to enhance the data-review process.

Email: rebecca.t.quintal@saic.com

<https://www.hydro-international.com/content/article/automated-side-scan-data-analysis>
