TEMPORAL ROBUSTNESS OF AN AUTOMATIC AURAL CLASSIFIER

Stefan Murphy and Paul C. Hines
Defence R&D Canada – Atlantic, P.O. Box 1012, Dartmouth, NS, Canada, B2Y 3Z7
stefan.murphy@drdc-rddc.gc.ca

1. INTRODUCTION

Active sonar systems are used to detect underwater manmade objects of interest (*targets*) that are too quiet to be reliably detected with passive sonar. In coastal waters, the performance of active sonar is often degraded by false alarms caused by echoes returned from geological seabed structures (*clutter*) found in these shallow regions. To reduce false alarms, a method of distinguishing target echoes from clutter echoes is required.

Research has demonstrated that perceptual signal features similar to those employed in the human auditory system can be used to automatically discriminate between target and clutter echoes, thereby improving sonar performance by reducing the number of false alarms [1]. The temporal robustness of this method is tested in this work by classifying recent echoes from 2009 using an automatic aural classifier previously trained with older (2007) echoes. Preliminary dependence on signal-to-noise ratio (SNR) is also presented.

2. EXPERIMENT

An active sonar experiment on the Malta Plateau was conducted during the Clutter'07 sea trial and repeated during the Clutter'09 sea trial. NATO Research Vessel Alliance ran a track to the southeast past Campo Vega Oilfield (ship track published in [2]). Broadband sources were used to transmit linear FM sweeps (600–3400 Hz) and a cardioid towed-array was used as the receiver. The sources and receiver were towed at a depth of 50 m. The original data set consists of over 95,000 pulse-compressed echoes returned from two underwater objects representing targets (an oil rig and a wellhead) and many geological clutter objects.

In order to avoid biasing the classification by SNR, the SNR distributions of target and clutter echoes are matched. After the SNR matching, approximately 25,000 echoes from Clutter'07 are used for *training* the aural classifier, and approximately 10,000 echoes from Clutter'09 are used for *testing* the classifier.

Many environmental factors that affect sonar echoes can change over a 2-year period; however, the aural classifier uses supervised learning and is unable to adapt to new data once trained. The two primary factors for this data set were the sound speed profile in the water column and the sea surface roughness.

Figure 1 shows that the sound speed profiles differed considerably between experiments. In 2007 the sound speed profile was downward refracting; in 2009 the profile was close to isovelocity.

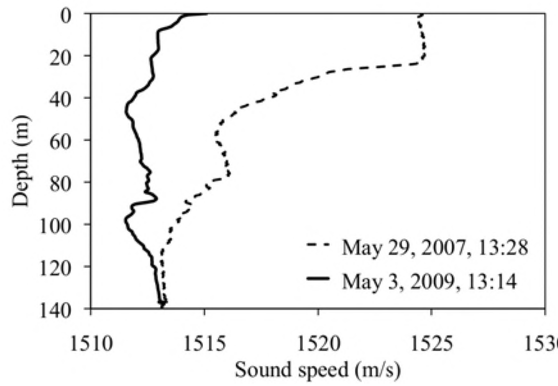

Figure 1 – Sound speed profiles for both experiments calculated from expendable bathythermograph (XBT) data.

Surface conditions also changed significantly between experiments. In 2007, the average relative wind speed of 13.0 knots resulted in Beaufort force 5–6 seas during the experiment. During the experiment in 2009, a lower average relative wind speed of 3.0 knots resulted in near flat seas. (Beaufort force 1).

3. AURAL CLASSIFIER

Details of the automatic aural classifier are published in [1]. The echoes are processed using a human auditory model that quantifies the timbre of each echo using 51 perceptual-based signal features that are not highly correlated ($r^2 < 0.81$) over the training data set.

The features are individually ranked by their ability to discriminate between target and clutter training echoes. Dimensionality is reduced by forming a subset of top ranked features, and further reduction is accomplished by principal component analysis. The number of top features and number of principal components are selected by the user.

A Gaussian classification method is used to calculate a target–clutter decision boundary in the feature space using the training echoes. The representation of the echoes was reduced to 2 dimensions by taking the first 2 principal components of the top 5 features. Figure 2 shows the minimum-error-rate decision boundary (circle) formed using Clutter'07 echoes. A scatter plot of Clutter'09 target and clutter echoes is overlaid on the decision regions. Since the data set consists of thousands of echoes, the plot is limited to a representative sample of 60 echoes.

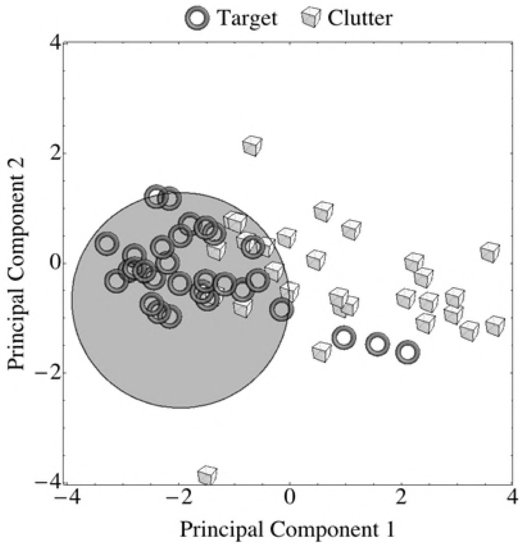

Figure 2 – Scatter plot of echoes in the reduced feature space. The gray circular target region contains 90% of the target echoes (toroids), and the surrounding white clutter region contains 64% of the clutter echoes (cubes).

4. RESULTS

4.1 Classifier performance

The minimum-error-rate operating point (shown in Figure 2) is chosen according to Bayes decision theory, with equal cost of misclassification for both target and clutter classes. In order to take all operating points into account, receiver-operating-characteristic (ROC) curves are used. ROC curves plot probability of detection versus probability of false alarm, and the summary performance metric used is the area under the ROC curve, A_{ROC}. For ideal classification, $A_{\mathrm{ROC}} = 1$, and if classification is performed by random guessing, $A_{\mathrm{ROC}} = 0.5$.

Since the number of features and principal components chosen are user-defined variables, it is possible to adjust them and monitor performance. Figure 3 shows a plot of classifier performance as grayscale intensity versus the number of top features used on the horizontal axis and the number principal components on the vertical axis. Darker color indicates higher performance. Since the number of principal components cannot exceed the number of features, data is constrained below the unit diagonal.

The peak performance ($A_{\mathrm{ROC}} = 0.903$) occurs at 29 features and 3 principal components. Having an A_{ROC} greater than 0.9 indicates a successful classifier, which demonstrates temporal robustness of the aural classifier. As a performance baseline, testing the classifier with the same Clutter'07 echoes it is trained with yields a peak A_{ROC} of 0.943.

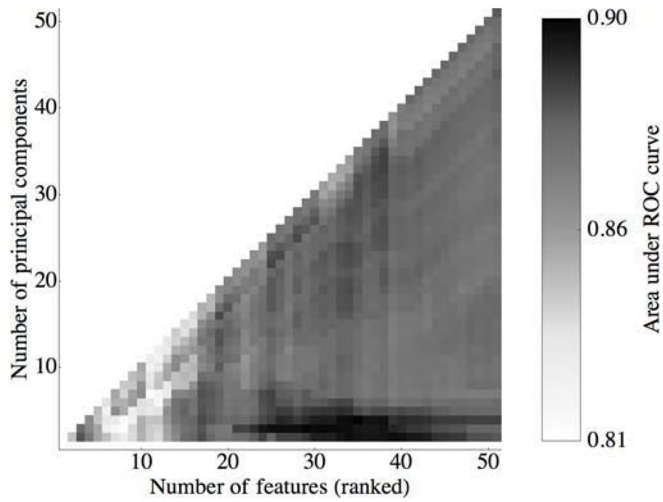

Figure 3 – Testing performance (A_{ROC}) of Clutter'09 echo classification using the aural classifier trained with Clutter'07 echoes.

5. CONCLUSIONS AND FUTURE WORK

The temporal robustness of the automatic aural classifier was demonstrated by classifying echoes using a classifier previously trained with echoes obtained 2 years earlier under different conditions.

Preliminary results show that classifier performance (A_{ROC}) increases with increasing echo SNR (dB); furthermore there is evidence that suggests the relationship is linear. This is an area of continued research.

REFERENCES

[1] Victor W. Young and Paul C. Hines, *J. Acoust. Soc. Am.* **122** (3), 1502 (2007).
[2] Paul C. Hines, Victor W. Young, and Jeff Scrutton, *Proceedings: International Symposium on Underwater Reverberation and Clutter* (ISURC, Lerici, Italy, 2008), pp. 293–301.

ACKNOWLEDGEMENTS

The authors would like to thank Doug Abraham, Charlie Gaumond, and Richard Menis for their contributions to the Clutter'09 sea trial, which was conducted by the NATO Undersea Research Centre (NURC) under Chief Scientist Peter Nielson. Funding provided in part by United States Office of Naval Research.