

Smart Sensing Analysis using Advanced Process Build Methods: A Case Study based on IEER data*

Robert J. (Bob) Jannarone, David J. Cohen and John T. (Tyler) Tatum
Brainlike, Inc.

John Joseph
Naval Air Systems Command

**NDIA Joint Undersea Warfare Technology Conference,
March, 2009, San Diego, CA**

*The authors appreciate support in obtaining the data for this paper from the Johns Hopkins University Applied Physics Laboratory.

Overview

In recent years, Brainlike, Inc., has been developing smart sensing models and methods [1-10], which have resulted in monitoring and surveillance precision improvements in a variety of applications[11-24]. Brainlike technology is based on a novel, automatic and adaptive (auto-adaptive) process that continuously adjusts for changing clutter conditions. This report describes a project that was designed to show Brainlike's auto-adaptive technology can increase detection and classification effectiveness, when used in conjunction with antisubmarine warfare (ASW) sonar data[25]. The project attempted to identify auto-adaptive processing (AAP) models that can significantly reduce clutter, relative to existing ASW signal detection methods [26-27]. The report describes how AAP models were developed and compared to an existing, Improved, Enhance Echo Ranging (IEER) process, in keeping with Advanced Processing Build (APB) methods. The report outlines empirical results in general terms. Detailed results may be found elsewhere [28].

This project marks a significant step toward understanding how to remove costly and distracting clutter, in defense and other applications. Clutter in data hides threats, buries prospects, wastes time, clogs channels, saps energy, crams storage, and costs money. With clutter removed, anomalies stand out, decisions are clear, and costs are cut. Most monitoring operations would benefit from affordable clutter reduction. ASW applications have especially important and difficult challenges, since (a) shallow water submarine activity in highly cluttered environments is increasing, (b) communications channels in next generation ASW systems will become increasingly clogged, and (c) highly sophisticated detection and classification procedures already in place represent a variety of others in defense applications that have been difficult to improve upon.

The results of this project show that AAP models and related methods can significantly reduce clutter in highly cluttered environments, but they may not operate as effectively in other environments. The results point toward several explanations for these mixed results. While AAP models have been shown to remove clutter effectively, delivery of globally superior AAP models will require further analysis and model improvements.

IEER Snippet Detection and Target Classification

IEER ping data is transmitted by sonobuoys to P-3 aircraft during ASW missions [26]. Each ping is initiated by a broad band, low frequency impulse source that is deployed by the aircraft. As up to four sonobuoy hydrophones sense the direct blast, they begin transmitting echo data to the aircraft. Each sonobuoy transmits echo data from 24 directional beams, electronically combined from 40 hydrophone readings, and spaced 15 degrees apart along the horizontal plane. Each beam's data is transmitted in a time series made up of discrete, complex sample values, at a fixed number of samples per second. Each time series covers a sufficiently long time span to include echoes from targets that may be several miles away. Thus, for any given ping the amount of data transmitted includes tens of thousands of complex sample values for each among 24 beams, from each of up to four sonobuoys.

Once ping data has been received by the aircraft, the IEER process reduces sonar ping data to a sufficiently small number of snippets for immediate operator analysis. Each snippet presented to the operator includes a segment of the time series that is only a few seconds long. Operators may view a variety of displays in the time and frequency domain for the snippet's dominant beam as well as its nearest neighbor beams.

IEER processing may be grouped into initial snippet detection and secondary snippet reduction phases as follows:

- A. IEER Snippet detection. Initial detection reduces input data to a number of candidate snippets. Each candidate snippet is assigned a signal-to-noise ratio (SNR) score, which tends to have a higher value if it contains either a target or target-like clutter than if it contains background noise. Candidate snippets having an SNR score that exceeds a preset SNR cutoff value are then selected for secondary snippet classification. The cutoff value is chosen to ensure that any possible target will appear in one or more selected snippets. Each selected snippet is assigned an overall SNR score that is the highest SNR value in its dominant beam.
- B. IEER Snippet classification. Secondary classification analyzes each snippet that has been selected from initial detection. Snippet classification includes computing 13 feature scores for each snippet and producing a multivariate Gaussian (MVG) score based on them to go along with the snippet's SNR score. MVG scoring is designed to produce a higher MVG

value for a snippet if it contains a target echo than if it contains clutter. Each selected snippet, along with its MVG and SNR scores, is delivered to an operator if and only if its value exceeds a preset MVG cutoff value. The MVG cutoff value is chosen so that practically all possible targets will be delivered to the operator. (In current practice, the MVG cutoff is set to the lowest possible value, so that the operator will see all detected snippets.)

Initial and secondary processing must operate together to add clutter removal value. MVG scores that are computed during snippet reduction identify targets more precisely than SNR scores that are computed during snippet detection. Thus, detection must identify snippets first based on SNR scores, allowing classification to discriminate target snippets from clutter snippets next, based on MVG scores. Each snippet's MVG score corresponds to a log likelihood ratio (LLR), based on multivariate normality assumptions [29-31]. The MVG score is computed as a function the snippet's feature values, along with sample statistics from archival data generation (DG) datasets. DG statistics include target and clutter means and covariances, from which the MVG distance measure is computed. Each snippet's MVG score is based on two Mahalanobis distance terms: the snippet's feature vector distance from the target mean vector and its feature vector distance from the clutter mean vector.

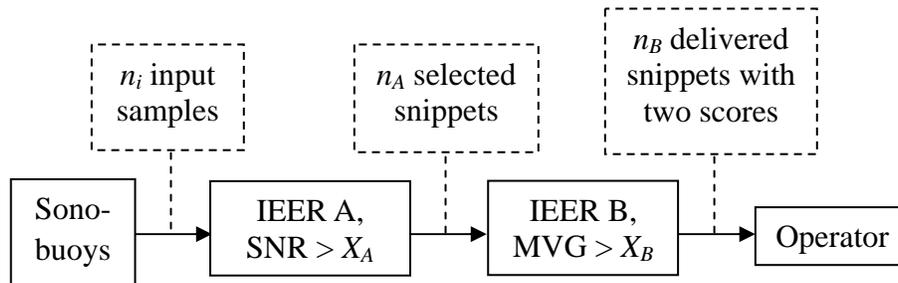


Figure 1. Two-Step IEER Processing in its Current Form

Figure 1 shows IEER processing in its current form. The number of input sample values is denoted by n_i . All samples are transmitted from sonobuoys to the aircraft. As they are being transmitted, they go through detection processing (A), which filters out nearly all samples. Remaining samples are transformed to a smaller number of candidate snippets, each of which is assigned an SNR score. Candidate snippets having an SNR value that exceeds a cutoff value denoted by X_A are selected for secondary processing. The number of selected snippets is denoted by n_A . Selected snippets then pass through classification processing (B), which produces an MVG score for each selected snippet. Selected snippets having an MVG value that exceeds a cutoff value denoted by X_B are delivered to the operator. The number of delivered snippets is denoted by n_B . (Since in current practice, X_A is set to a low value, $n_A = n_B$.)

Along with each delivered snippet, IEER processing delivers the snippet's SNR and MVG scores to the operator. Snippets and their scores are organized into a table for each ping. The table is ordered according to snippet MVG value. Operators use these tables, along with other mission indicators, to select individual snippets for further investigation.

The IEER system is designed to deliver only a small number of non-target snippets that have been caused by clutter, known as false snippets, that is. Doing so is essential for effective target recognition, for the following reasons: (a) operators examine snippets until and unless they conclude that one or more have been caused by targets; (b) operator time is limited; and (c) operator fatigue is a major concern.

In its current form, IEER processing has been shown to deliver snippet tables with a high likelihood of containing all targets and only a small number of false snippets. However, the number of false snippets must be reduced further to increase operator effectiveness and prepare for deploying more than four sonobuoys in future ASW systems.

Within any given pings, false snippets having higher MVG values than delivered target snippets cause added problems, because operators tend to examine these so-called distracting snippets first. As outlined below, the results of this project found that distracting snippets have higher frequencies than expected under MVG assumptions. These results point may partly explain why auto-adaptive processing shown improvements and may point toward further improvements.

Auto-adaptive Detection and Classification

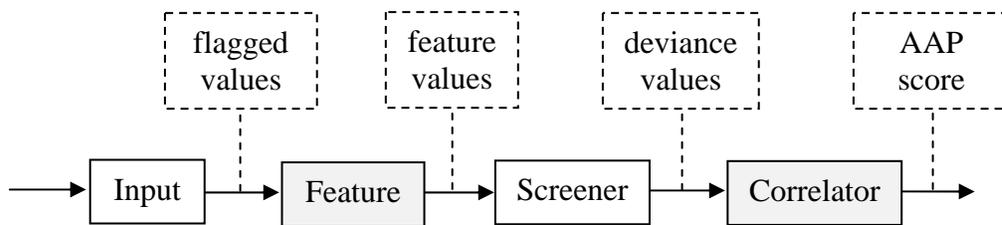


Figure 2. Auto-Adaptive Processing Stages

Brainlike AAP proceeds in a number of stages [6-8], including those shown in Figure 2 that were used in this project. The input stage flags input values having excessive magnitudes so that they will not adversely affect downstream processing. The feature stage may be used to transform input values into feature values. The screener stage converts input (or input feature) values to deviance values quickly, by comparing input values to recursively updated means. The correlator stage converts screener deviance values to more precise output scores, by comparing them to expected values based on nearest neighbor correlations.

The screener stage uses a relatively fast and simple method for auto-adaptive anomaly detection. Each time an input value is received, the screener computes the difference between its input value and a previously computed expected value. It then divides the difference by a normalizing constant to compute a deviance value. The screener then quickly and automatically adapts to the current value, by recursively updating learned metrics for computing expected values and standard errors. The screener also ensures that highly deviant input values will not adversely affect its operations [8]. Since the screener stage updates its learned metrics continuously, deviance values tend to have stationary distributions with a mean of zero and a standard deviation of 1. Stationary deviance distributions allow detection and classification rules based on them to remain robust under changing environment, sensor, and clutter conditions. Screener expected values for a given input depend only on that input's values and not values from other inputs.

The correlator stage closely resembles the screener stage in terms of how it computes robust deviance values, once it has computed expected values. The correlator differs from the screener only in the way it computes expected values. Instead of computing them for a given input as a function of only that input's values, the correlator computes them as functions of other inputs that are nearest neighbors in time and space. Expected correlator values depend on a variety of metrics that the correlator updates recursively, including but not limited to mean values, mean squares, mean cross-products, and their matrix inverses.

The correlator stage is sometimes not used, because in some, high-speed applications it operates too slowly to keep up with data in real time. For example, if an AAP correlator based on many features were used in place of initial IEER processing to assess each input sample for deviance, and if many nearest neighbors were used for estimation, special purpose hardware might be required to identify sample anomalies in real time. In other applications, the correlator adds value by discriminating unexpected anomalies from clutter more precisely than the screener. For example, an AAP correlator stage could easily identify snippet anomalies in real time, once they have been created during initial IEER processing, because the IEER process produces far fewer snippets than samples.

Figure 3 shows how the correlator may be configured for nearest neighbor, IEER anomaly detection and snippet classification. The lower right portion of the figure shows features for three beams being measured currently (Now), the center portion shows the same features for the same three beams being measured one time slice ago (Now-1), and the top left portion shows the same features for the same three beams being measured t slices ago (Now- t). The figure shows n features being measured for each beam at each time slice. The correlator may be configured so that deviance values for each of the n features being evaluated for any given beam may be based on its nearest neighbors in time and space, as shown in the figure. For example, the expected value of the center beam value, f_j , being measured currently (shown in red), may be computed as

a function of the other $n - 1$ features being measured currently for that beam, the $2n$ nearest neighbor beams being measured currently, and the $3n(t - 1)$ features that were measured recently.

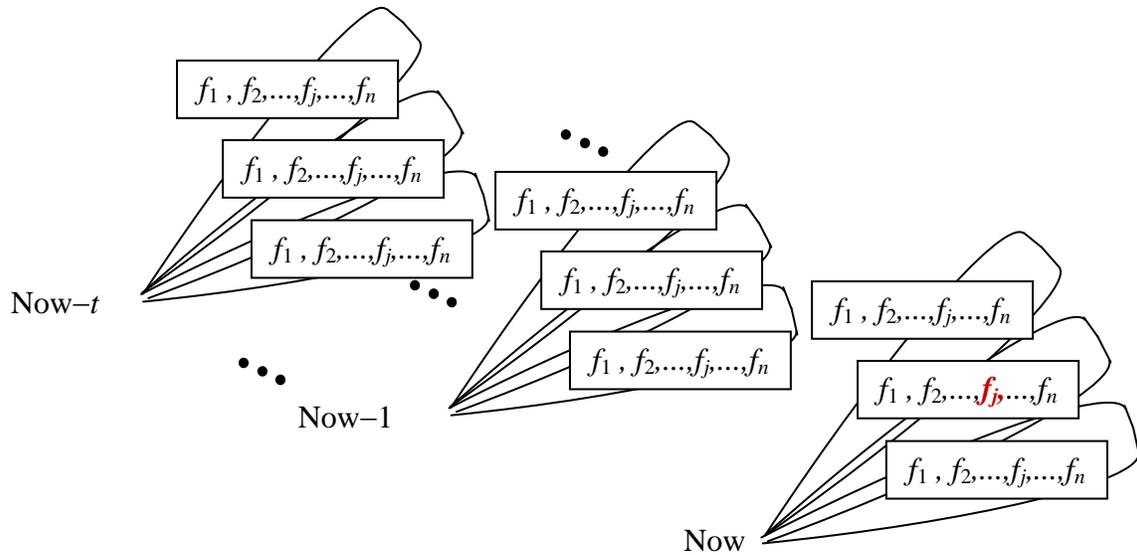


Figure 3. Auto-Adaptive, Nearest Neighbor Estimation

Computing expected values in this way produces more precision when nearest neighbor values reduce feature value uncertainty, as is well known in the signal processing and statistical literature [29-31]. Updating learned metrics for computing the expected values produces even more precision when background clutter, environmental conditions, and sensor properties are changing over time, as has been shown in a variety of studies [1-5,11-25]. This added precision may not be possible in real time, if input values are arriving too quickly or the number of independent values for estimation is too large, because (a) updating learned metrics requires time-consuming matrix inverse operations, and (b) matrix inversion time increases with the square of the number of features involved. The tradeoff between detection precision and computing practicality is one reason why the project focused on the snippet domain approach to AAP shown in Figure 3, instead of alternative, sample domain approaches. Snippet arrival rates are orders of magnitude lower than sample arrival rates.

The AAP value shown in Figure 3 is computed while the AAP process is computing feature deviance values. Its basis is a Mahalanobis distance, measured between each input vector and its most recently updated mean, and scaled by its most recently updated covariance matrix inverse.

The correlator includes a variety of capabilities that ensure fast but effective operation, including the following: (a) low-level subscript manipulation for efficient nearest-neighbor windowing, (b) direct inverse updating as well as efficient, window-based inversion; (c) efficient correction for

linear redundancies; (d) prevention of adverse outlier effects; and (e) structured learning regimes that smooth learned metric updating.

Auto-adaptive Enhancement to IEER Processing

IEER processing already includes several clutter removal methods that resemble AAP processing [25]. IEER detection identifies aberrant input values and removes their adverse affects, resembling methods used in the AAP input stage. IEER detection also continuously corrects sample values for changing baseline values, resembling methods used in the AAP screener stage. In addition, IEER detection uses spatial mean estimation to control for nearest neighbor beams, resulting in a less precise but faster alternative to the AAP correlator.

IEER MVG scores resemble AAP deviance statistics in some ways, but not others. The global AAP, Mahalanobis distance measure is similar to the MVG clutter distance term, with two exceptions: statistics for the AAP measure are continuously adjusted for changes over time, and AAP deviance values may account for nearest neighbor values in time and space. IEER classification also includes an optional provision for so-called adaptive learning. Adaptive learning replaces the clutter term in MVG scores with a term that continuously updates MVG statistic means and covariances, like the AAP distance measure.

IEER adaptive learning is not currently enabled since it has not added substantial effectiveness, perhaps because, IEER detection may not provide ample snippets for effective clutter covariance estimation. If the X_A cutoff value shown in Figure 1 were reduced, n_A could be increased, producing more snippets that would could solve the second problem.

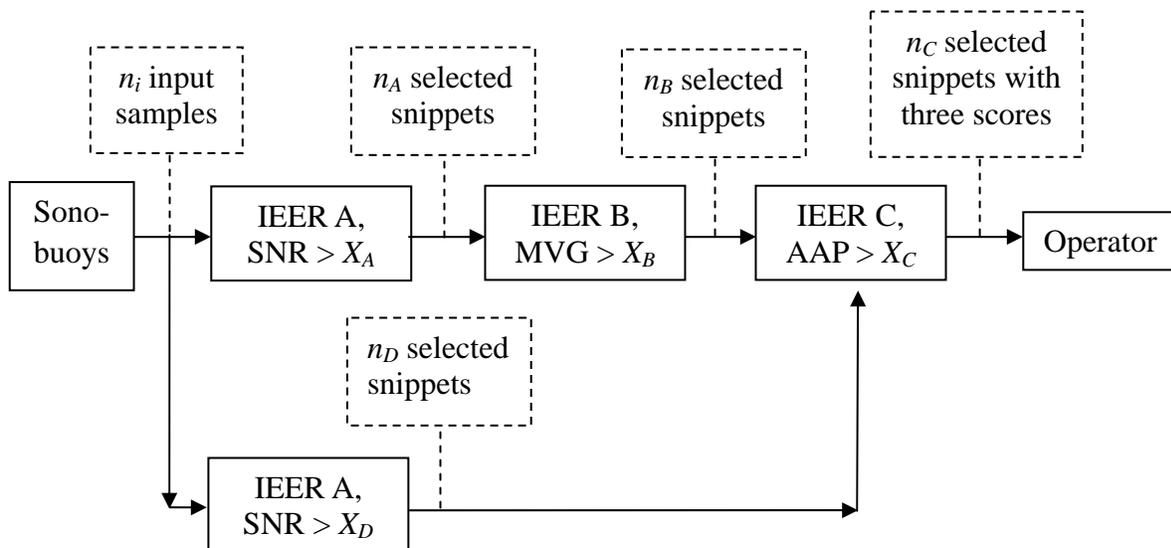


Figure 4. Combined IEER/AAP Model

Figure 4 shows the combined (IEER/AAP) model that was used in the project. The figure shows a second, initial IEER processing block, operating in parallel with the first block, and supplying its own set of SNR values to the AAP block. The second block has a different snippet selection cutoff value, denoted by X_D , than its first block counterpart, X_A . The X_D value is set lower than the X_A value, resulting in more snippets, denoted by n_D being available for AAP learning than its first block counterpart, n_A . Proceeding in this way provided ample snippets for effective AAP covariance estimation.

The Figure 4 approach is simpler to implement and integrate than alternative AAP approaches. The most complicated added component is the second initial IEER processing block, which is already operational in the current IEER system. As a result, two IEER processes that are already operational could be run in a combined system. Enhancing the IEER process in this way is much more straightforward than creating AAP replacements for either IEER detection or IEER classification. The Figure 4 model has the added advantage of delivering currently accepted IEER metrics to the operator along with a third metric, rather than replacing IEER metrics altogether.

Advanced Process Build Methods

In recent years, the U.S. Navy has developed so-called advanced processing build (APB) methods, which are designed to evaluate advanced signal processing models such as AAP, and deliver such models to the fleet quickly, in the form of new or modified products. APB methods are based in part on supplying training datasets, such as the DG datasets for this study, to advanced model developers, and withholding testing datasets, such as the TF datasets for this study, for independent testing.

Results

Results for this project were based on analyses of IEER snippets extracted from training and testing datasets provided by NAVAIR, PMA-264. Training datasets included DG02, DG03, DG05, and DG06, two of which came from shallow water and two of which came from deep water. Testing datasets included TF-20, which was highly cluttered, along with others that were much less cluttered. Further analyses were based on simulated data, designed to show how IEER snippets behave differently when baselines change than when they remain stationary.

Potential AAP added value was examined at both the sample detection level and the snippet classification level. At the detection level, ping sample magnitudes from the DG datasets were first examined for dataset differences and nonstationary behavior as possible causes for false detections. Ample evidence of within-ping, nonstationary behavior was found. Some time series segments showed more spiky behavior and different baseline amplitudes than others. Two

potential sources of false detections were found: (a) high amplitude data spikes, which operators had previously reported as being distracting, and (b) amplitude excursions that covered several beams, which were clearly different from beam-specific echoes from targets. One AAP model effectively removed the spikes, and another AAP Brainlike model effectively removed the multiple beam clutter. Further investigations, however, showed that both types of clutter produce negligible false snippets in other, recently developed IEER enhancements, which include effective methods for removing them.

At the classification level, an evaluation of results from current IEER processing showed that substantial differences exist over DG datasets, in frequencies of false snippets presented to operators. Differences in “distracting clutter” frequencies were especially substantial. Further results showed that clutter means and covariances among MVG features differ substantially, both between and within DG datasets. Additional results showed that observed distracting clutter (as defined previously) differed substantially over DG dataset pings far exceeded expected distracting clutter under MVG normality assumptions (as obtained from simulations — see below). These results point toward substantial added value potential for AAP models and methods.

At the classification level, analyses were conducted, comparing IEER results based on the original 13 features alone (Figure 1) to combined IEER/AAP results based on original 13 features, along with 13 AAP deviance features (Figure 4). AAP scores were obtained by reducing IEER detection cutoff values (X_D in Figure 4) so that an ample number of snippets (n_D in Figure 4) was available for AAP covariance estimation. Among those snippets, feature values corresponding to only the IEER snippets (n_A in Figure 4) were retained used to obtain sample statistics, including means and covariances for targets and non-targets, from the DG datasets. The snippet feature scores, along with the sample statistics, were used to obtain Mahalanobis distance scores and MVG scores for the DG and TF datasets.

Results [28] showed improved IEER/AAP performance relative to IEER performance alone in the TF 20 dataset. However, no such improvements were demonstrated with the less cluttered TF datasets. Further analysis showed that while IEER testing data that validated the model closely resembled IEER training data, the IEER testing data that showed no improvements did not. Most notably, AAP detection level cutoff value (X_D in Figure 4) produced a much higher proportion of clutter to target snippets in TF-20 than in the TF datasets.

Further analyses were performed to explore differences between the DG and TF-20 datasets on the one hand and other TF datasets on the other hand, in order to explain why relative IEER/AAP performance differed among them. Scatterplots of clutter Mahalanobis distance versus target Mahalanobis distance, broken down by target versus clutter snippets produced informative results. Under normality assumptions, clutter snippet distances from clutter mean vectors and

target snippet distances from target mean vectors for IEER (13–1 degrees of freedom) scoring should have means of 12 and variances of 144. Interestingly, distance means and variances for all DG datasets and TF-20 were close to their expected values relative to the others, but clutter distance means and variances for the other TF datasets were much higher.

The scatterplots also indicated that clutter distances were more positively skewed than target distances in all datasets. This result suggests that modified IEER scoring, based on a weighed instead of an unweighted difference between clutter and target distances, could improve IEER classification substantially. The result also suggests that thick-tailed, clutter feature distributions like those simulated below, could explain skewed clutter distances like these, and that AAP models could be formed that reduce clutter by removing their effects.

In order to explain IEER and IEER/AAP performance under nonstationary conditions, a variety of simulation studies was performed. Simulated datasets were based on 13 feature data and 26 feature data. Some datasets were generated in keeping stationary, multivariate normality assumptions. Others were based on non-stationary trends, along the lines of those that were observed in the DG datasets. Still others were based on generating bimodal clutter distributions, which produced distracting clutter similar to that observed in the DG datasets.

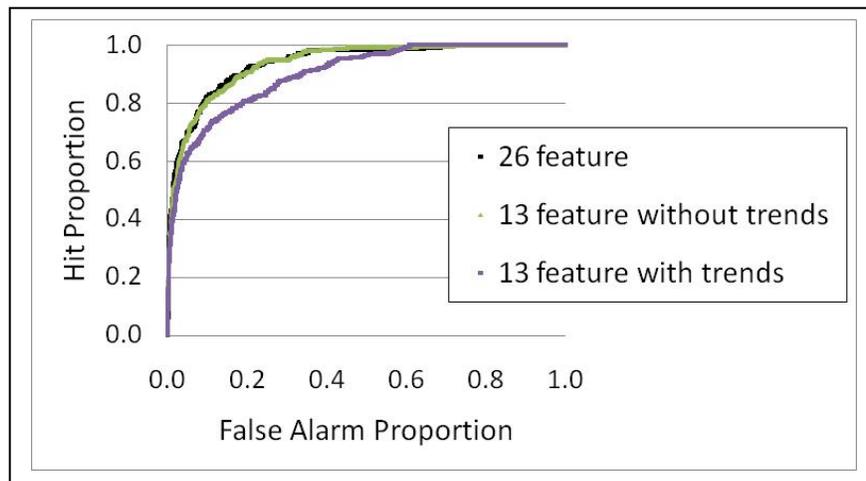


Figure 5. Hits versus false alarms based on simulated IEER IEER/AAP scoring

Figure 5 shows hits versus false alarm plots (also called receiver operating characteristic (ROC) curves [6-8]), based on two simulated datasets. One dataset was created to resemble the DG datasets, by including trends over time. The other dataset was the same as the first, except trends over time were not included. The blue, “13 feature with trends” curve was based on IEER scoring from the dataset that included trends. The black, “26 feature” curve was based on IEER/AAP scoring for the same dataset. The green, “26 feature with trends” curve was based on

IEER/AAP scoring for the dataset without trends. The first two curves resemble those obtained from the DG datasets.

Several other simulation studies were performed, some of which addressed a concern that training datasets were too small to effectively estimate sample statistics, especially 26 by 26 covariance matrices. Results showed that standard techniques for reducing ill conditioning [32-33], when modified in keeping with heteroscedastic MVG assumptions, may improve IEER/AAP performance. Independent validation on DG data sets showed that the 26 features model performed better than 13 features alone. Other studies evaluated alternative scoring based on homoscedastic assumptions and on using clutter Mahalanobis distances alone. These studies indicated that MVG scoring, based on heteroscedastic MVG assumptions, performed better than alternatives, as expected.

Further simulation studies were performed to explain how datasets can be constructed that produce higher than expected “distracting clutter.” These studies first generated data from clutter and target distributions, and then randomly selected ping data from them. For each ping, 10 snippets were selected, with clutter snippets being generated with a probability of 0.9 and target snippets being generated with a probability of 0.1. Frequencies for distracting clutter snippets were then obtained, over pings. A variety of datasets were evaluated to reproduce distracting clutter frequencies like those that were obtained with some of the DG datasets. The datasets illustrated below came closest.

Figure 6 shows normally generated data that produced results expected under normality conditions, and Figure 7 shows corresponding distracting clutter frequencies that were obtained (the dominant frequency of zero has been capped at 20 so that other frequencies can be seen).

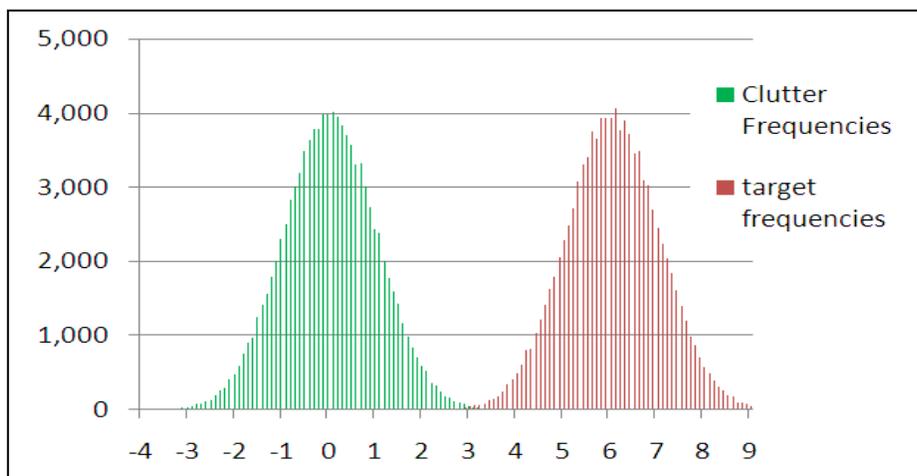


Figure 6. Normally generated target and clutter data

Since Figure 6 shows very little overlap between targets and clutter, very few clutter events would be expected to occur along with corresponding target events. Accordingly, Figure 7 shows that only one target event occurred with a corresponding clutter event, while all other target events did not have a corresponding clutter event.

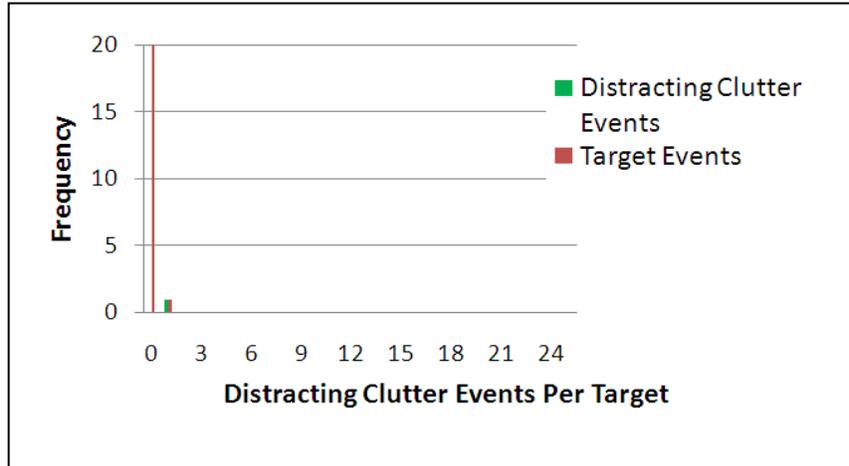


Figure 7. Normally generated target and clutter data (capped)

Figure 8 shows bimodal clutter data with much more target distribution overlap than Figure 6. The data were generated to produce the Figure 9 (capped) distracting clutter frequencies, in order to resemble those that were obtained from some of the DG datasets. Evidently, thick-tailed clutter frequencies can produce unexpected levels of distracting clutter.

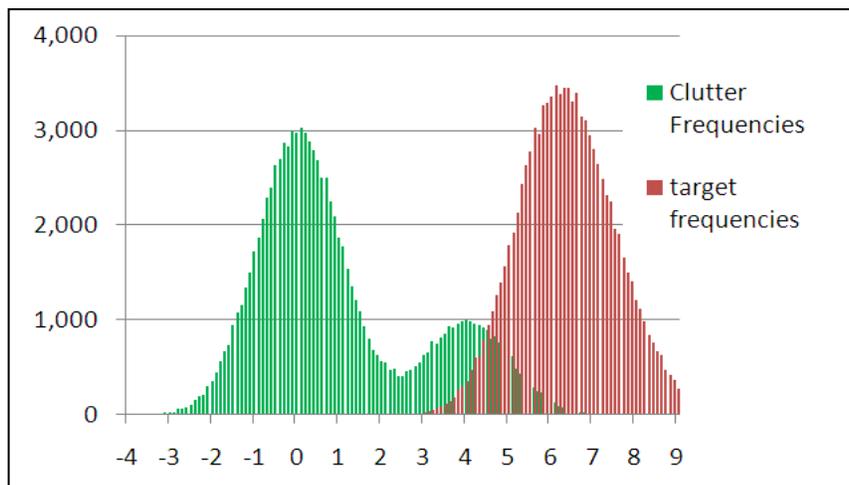


Figure 8. Normally generated target and bimodal clutter data

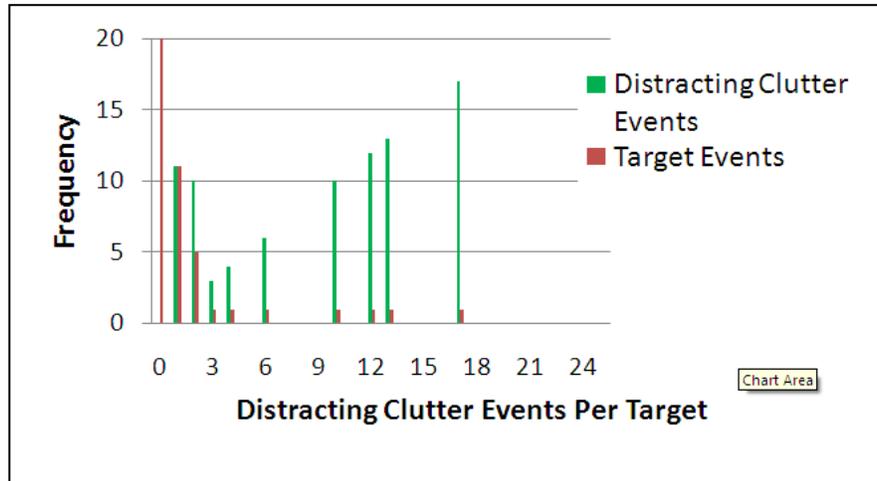


Figure 9. Normally generated target and clutter data (capped)

Conclusions

Data-based results from this project have shown that snippet detection samples, as well as target classification snippets, differ over datasets and vary over time within datasets. Data-based results along with simulation results have further shown that such trends can result in departures from MVG assumptions that will produce higher than expected false alarms. Results have further shown that one auto-adaptive processing model can remove such trends in highly cluttered datasets, but not in other less cluttered datasets.

A variety of other results have pointed toward potential AAP improvements. For example, if sufficient statistics could be identified that reflect clutter distribution departures from normality, then they could be continuously estimated by AAP models in ways that would further improve clutter removal. Also, if sufficient statistics could be identified that reflect differences in target and clutter distribution frequencies, then they could be continuously estimated and AAP learning metrics could be adjusted accordingly.

In conclusion, auto-adaptive processing models and related methods can significantly reduce clutter in highly cluttered environments, but they may not operate as effectively in other environments. Delivery of globally superior AAP models will require further analysis and model improvements, including new ways to reflect trends in clutter.

References

1. R.J. Jannarone, D.J. Cohen, J.T. Tatum, D. Statter, J. Joseph, and M. Wardlaw, "Recent Smart Sensing Developments: Undersea Warfare Surveillance Prospects," Paper presented at the NDIA Joint Undersea Warfare Technology Conference, March, 2009, San Diego, CA.
2. R.L. Smith and R.J. Jannarone, "An Auto-Adaptive Statistical Procedure for Tracking Structural Health Monitoring Data," *Proc. SPIE*, 5391, 166-176, 2004.
3. R.J. Jannarone, *Concurrent Learning and Information Processing: A Neurocomputing System that Learns during Monitoring, Forecasting, and Control*. Chapman & Hall, New York, 1997.
4. R.J. Jannarone, "Concurrent Information Processing, I: An Applications Overview," *Applied Computing Review*, Vol. I, pp. 1-6, 1993.
5. R.J. Jannarone, K.F. Yu, and Y. Takefuji, "Conjunctoids: Statistical Learning Modules for Binary Events," *Neural Networks*, Vol. 1, pp. 325-337, 1988.
6. R. J. Jannarone, J. T., Tatum, et. al. "Efficient Processing in an Auto-Adaptive Network," patent pending, Brainlike, Inc..
7. R. J. Jannarone, J. T., Tatum, et. al. "Auto-Adaptive Network," patent pending, Brainlike, Inc.
8. Brainlike Studio, Version 2 Manual, Brainlike, Inc., 2008.
9. "Brainlike Sensing: enabling technology for the wireless revolution." Brainlike, Inc., Technical Report, November, 2008.
10. Brainlike Products and Services. <http://www.brainlike.com/products.htm> .
11. R.J. Jannarone, C.M. Traweek, T. Wettergren, and J. T. Tatum, "Damn the Torpedoes or Let's Wait and See? Decisive Inference for Making the Right Choice," *Military Operations Research Journal*, in review.
12. R. Jannarone, J.T. Tatum, J. Joseph, N. Naluai, M. Wardlaw, and A. Slaterbeck, "Auto-Adaptive Systems for Undersea Warfare: Added Value Analysis for Selected Applications," paper presented at the 2007, NDIA Joint Undersea Warfare Technology Conference, San Diego.
13. R. Jannarone, "Classification and Detection Improvement through Auto-Adaptive Clutter Estimation," Office of Naval Research Maritime Sensor Program Review, Newport RI, August, 2006.
14. R.J. Jannarone and J.T. Tatum, "A Novel Process for Littoral Target Recognition: Preliminary Analysis Results," Proceedings of the 25th Oceans Conference, Washington, DC, 2005.
15. "Clutter Removal and Substantially Improved Submarine and Mine Detection Through Affordable "Brainlike" Methods," Final Report under Office of Naval Research Phase I STTR Contract #N00014-05-M-0287, April, 2006.
16. "Shallow Water Attack Prevention," Brainlike Inc., December, 2003, www.Brainlike.com.
17. "Smart Sensor Engineering: a case study based on accelerometer data," Brainlike, Inc., December, 2008, http://www.brainlike.com/pdfs/Smart_Sensor_Engineering_Brainlike_Case_Study.pdf .

- “Brainlike Homeland Security Relevance,” Brainlike Surveillance, Inc., white paper, December, 2003, www.Brainlike.com.
18. “Equipment Breakdown Prevention,” Brainlike, Inc., December, 2003, www.Brainlike.com.
 19. “Equipment Health Monitoring,” Brainlike, Inc., December, 2003, www.Brainlike.com.
 20. “Brainlike Monitoring Improvement Illustration,” Brainlike, Inc., December, 2004. www.Brainlike.com.
 21. “Return on Investment Analysis,” Brainlike, Inc., December, 2003. <http://www.brainlike.com/savings.htm>.
 22. “Auto-Adaptive Submarine Detection,” Brainlike, Inc., August, 2003. (http://www.brainlike.com/pdfs/shallow_water_attack.pdf).
 23. “Auto-Adaptive Mine Detection (http://www.brainlike.com/pdfs/target_recognition.pdf).
 24. Auto-Adaptive Meter Monitoring,” Brainlike Surveillance, Inc., August, 2003. (http://www.brainlike.com/pdfs/monitoring_improvement.pdf).
 25. "Clutter Removal and Substantially Improved Submarine Detection, through Affordable 'Brainlike' Methods," Navy STTR, Contract #N68335-07-C-0096, December, 2006.
 26. S. Shelikoff, “Analyzer Subunit, and Improved Echo Ranging Functional Overview,” CDRL No. B008, Naval Air Warfare Center, Patuxent River, MD, Revision 5, L-3 Communications Titan, 24 August, 2007.
 27. F.B. Shin, D.H. Kil, & R. F. Wayland. Active Impulse Echo Discrimination in Shallow Water by Mapping Target Physics-Derived Features to Classifiers. IEEE Journal of Oceanic Engineering, Vol. 22, No. 1, 1997.
 28. R.J.. Jannarone, D.J. Cohen, J.T. Tatum, and J. Joseph, “Smart Sensing Analysis using Advanced Process Build Methods: A Case Study based on IEER data,” Paper presented at the NDIA Joint Undersea Warfare Technology Conference, March, 2009, San Diego, CA.
 29. H. V. Poor. *An Introduction to Signal Detection and Estimation* 2nd Edn, Springer, New York, 1994.
 30. E.L. Lehmann, *Theory of Point Estimation*, Wiley, New York, 1983.
 31. E.L. Lehmann, *Testing Statistical Hypotheses*, 2nd Edn., Wiley, New York, 1986.
 32. Effron, & C. Morris. Stein’s Estimation Rule and its Competitors – An Empirical Bayes Approach. J. Amer. Statist. Assoc. Vol. 68. 117-130. 1973.
 33. W. James, & C. Stein. Estimation with Quadratic Loss. Proc. Fourth Berkeley Symp. Math. Statist. Prob. Vol. 1. 361-379. Univ. of California Press, California, 1960.