Classification Using Multiple Pass Fusion

Andrew D. Wilby and J. Kent Harbaugh

***Abstract*–A system architecture for the automated fusion of multiple detections by computer-aided detection and classification (CAD/CAC) algorithms from multiple passes over a scene is described. The fusion of multiple detections exploits features that are not available from a single pass.**

***Index Terms*—classification, fusion, multiple-pass**

I. Introduction

The automated detection and classification of threat objects in sidescan and synthetic aperture sonar images is an important problem for harbor security and fleet protection. The goal of automation is to relieve the burden on human operators of scanning through the large areas covered by modern sonar systems. In the typical process, a detector finds possible threat objects, a feature extractor quantifies appropriate characteristics, and a classifier discriminates potential threats from non-threat objects. The human operator then reviews the automatically generated potential threats for final classification. Each pass is processed independently; classification of each object is based on a single look from one pass.

While this procedure works well for objects that are clearly threat objects or clearly non-threat, it is often inadequate for borderline cases. In those cases, when classification is being performed by human operators, they will review looks at the location from previous passes. Often views from earlier passes, while inconclusive when considered separately, can in combination provide enough information to classify the object with confidence. We have developed an automated detection and classification architecture using looks from multiple passes to mirror this process. For each detection made, image data at the same location is automatically reviewed. Features from the current and past detections are fused along with additional multi-pass features extracted, and used by a multi-pass classifier to make a threat/non-threat decision. Furthermore, future passes by that location will automatically cue the automated detection and classification process, making use of the previously extracted features, to update the earlier classification.

In this paper, we describe in detail the multi-pass classification architecture. We will discuss our techniques for fusing detections from multiple passes and the multi-pass features extracted from the fused detections.

Manuscript received July 30, 2010

J. K. Harbaugh is with Applied Signal Technology, Torrance, CA 90502 USA phone: 310-436-7013; fax: 310-532-6787; e-mail: kent_harbaugh@appsig.com.

II. Architecture

The architecture schematic is shown in Fig. 1. A portion of a SAS or sidescan image is given to a suite of CAD/CAC algorithms. These algorithms produce a list of locations and threat confidence scores for threat objects within the image. Optionally, for each algorithm, a list of features can be returned for each detection. The interface within the architecture for these CAD/CAC algorithms is open so that the operator may include multiple third-party algorithms of their choice. The architecture does not need to know the implementation or training details of these algorithms, allowing the proprietary nature of this information to be preserved. The training and utility of each algorithm is the responsibility of their respective creators.

The detections from the multiple algorithms are then combined based on proximity. Because the detections are all based on the same image, co-registration is not an issue. The proximity fusion gives the combined set of location, scores and features (if available) to the single-pass, multi-algorithm classifier. This classifier uses the scores and features to produce a fused threat score and is described in more detail below. Since it is not expected that every CAD/CAC algorithms will find all threats, missing detections are handled gracefully in a way that does not bias the fused score. The fused detection including all the information is placed into a database.

So far, the architecture follows the typical processes of a multiple algorithm suite. In the next step, a multi-pass fusion engine compares the current detection with previous detections. Described below in more detail, this process associates together detections from different passes over the same object and generates a combined set of features for each object. This feature set is then used by the multi-pass, multi-algorithm classifier to compute an overall multi-pass score. High threat objects are presented to the operator, while both high and low threat objects are recorded in the database.

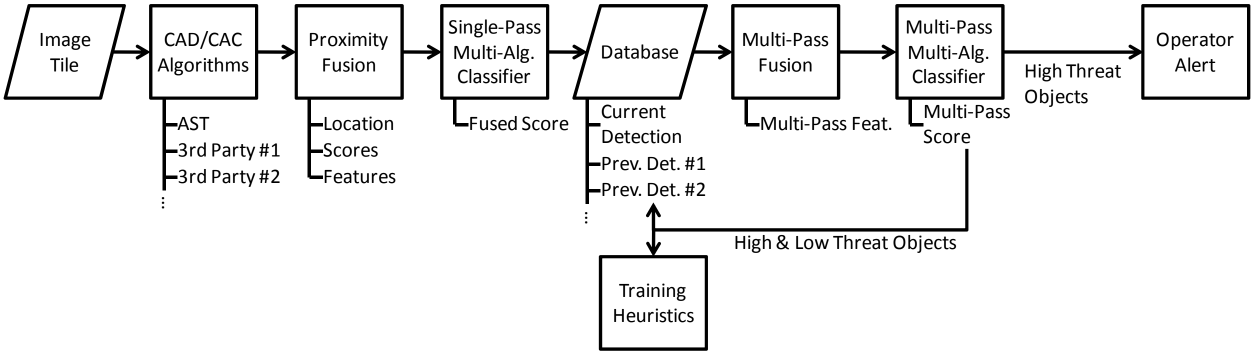

Fig. 1: Architecture for combining multiple detection and classification algorithms on single and multiple passes.

A. Multi-Pass Detection Fusion

While fusion of detections by multiple CAD/CAC algorithms within a single image are simple to co-register based on proximity, fusing detections within images from multiple passes is more difficult because of navigation and sensor variability between passes. The variability can be mitigated by procedures to accurately geo-locate the pixels of each image; however, this is not always easy to do. Although this architecture attempts to accurately geo-locate the images, to have a more robust multi-pass fusion, it does not rely entirely on proximity to associate detections from different passes.

Instead, a more generalized “parametric” proximity is used, employing the features of each detection as well as its location. The sum of the squares of the differences between two detections’ locations and other features is used as a distance metric. Each term is weighted, by default, according to the expected deviation of the location or feature. This weight can be tuned by an expert operator, for example to reduce the input of less discriminatory features; but will not be modified during regular operation. If two detections, from different passes, are within a specified parametric distance, they are considered to represent the same object. Detections beyond that distance are not associated.

This multi-pass fusion technique is intended to imitate the process used by an operator to find the same object on previous passes. While the operator will use the locations as the first indicator, an apparently co-located detection will be ignored in favor of another nearby one, if the features of the other are a better match.

Once the detections from multiple passes are co-registered, multi-pass features are added to the fused detection. These features should capture information that can only be determined by the association of detections from multiple passes. The intent is to produce a feature set that is more discriminatory than all of the features from single passes considered individually.

Several options are considered. The simplest is finding the mean and variance of each of the features over the multiple passes. This should be a better measure compared to the features individually from each pass. Potentially the statistics can be weighted according to the confidence score of the detection, so that a strong detection has more influence than a borderline one.

Another choice for a multi-pass feature is picking the “best” of detections, where best is determined by the scores. This can be done on a per-pass or per-classifier basis. For example, the best detections can be chosen by the single-pass fused score, using all the CAD/CAC algorithm features that were used to generate that score. Alternatively, since the highest score for algorithm #1 might be from a different pass than the highest score for algorithm #2, the features for each algorithm can be chosen from different passes. This choice takes advantage of the possibility that different algorithms may perform better under different circumstances, particularly different aspect angles of an object. Once the best features are determined, they can be used as is or averaged together over the top several.

One option that may be especially useful for change detection applications is an old versus new comparison of features. The difference between features from the most recent pass and the mean (or best of) features from previous passes may be useful features. The older features can be time-weighted. This difference could be exploited to find objects that are superficially static, but may in fact have been modified.

Which type of multi-pass feature performs the best is yet to be determined, and may depend on the specific targets and CAD/CAC algorithms used. In general, all of the features can be made available to the multi-pass classifier, at the cost of increased training time and space. Since the classifier will use only the most discriminatory of the features; adding additional feature will not reduce its performance. Once the classifier is trained, to improve execution speed and reduce storage requirements, only those features used by the classifier need to be computed.

B. Classifiers

Two types of classifiers are included with the architecture. The first is a trained classifier used to generate a threat score based on the available features. Both a Joint Gaussian Bayesian classifier (JBC) [1] and a Relevant Vector Machine (RVM) [2] classifier are implemented; either may be used. While the internal computational details of the two classifiers are different, their performance is comparable.

Either choice requires training. Training takes two sets of feature vectors, one set for threat detections and one set for non-threat detections. A heuristic algorithm iteratively trains the classifier. At every step, each set is divided into two—one for computing classifier parameters and one for evaluating the performance of the classifier using those parameters. Performance is compared by integrating the probability of correction classification over a designated range of the probability of false classification. This is typically either the entire domain from zero to one if the operating regime is unknown, or a small range of values around the expected operating value.

An evolving subset of features is used at every step, so that features that have little discriminatory power are eliminated from consideration and those with the most are retained [3]. Both the subsets of vectors and of features are varied with each pass through the loop, so that solutions optimized to a narrow range of features are avoided. The classifier with the final parameters should perform well under most of the range of features presented by the training sets.

This trained classifier, in the single-pass multi-algorithm role, is used to generate a fused threat score based on the detections found by one or more CAD/CAC algorithms at a single location. Features used by the classifier here include both the score reported by each classifier and any features returned by the corresponding detector. The classifier, in the multi-pass multi-algorithm role, is also used to compute a multi-pass threat score, based on all of the available features and scores from the detections fused together, including the fused scores generated by the single-pass multi-algorithm classifier and multi-pass features produced by the multi-pass fusion. While the same classifier algorithm is used in both roles, obviously they use and are trained with different features.

The second type of classifier is used to partially automate the initial training of the first classifier type. Generating the sets of feature vectors needed for training can be time consuming, since appropriate threat and non-threat examples must be extracted from sensor data. To ease this process, a simple detector and classifier are used to scan the imagery. The detector finds any object that is vaguely threat-like. Next, the classifier uses a simple fuzzy-logic calculation to evaluate each detection. Any detection that has even a low possibility of being a threat is flagged to be included in the training sets. Detections that are extremely unlikely to be threats, as determined by a tunable parameter, are eliminated from consideration. An expert human operator then must decide if a detection is a threat, non-threat, or not useful for training. While human intervention is required, classifying detections is generally easier than scanning imagery for all threats and non-threats useful for training.

Once an initial training set of detections is evaluated, the CAD/CAC algorithms produce the features that will be used by the two multi-algorithm classifiers. They are trained sequentially; the single-pass classifier must be trained first, since its output of a fused threat score can be used as a feature by the multi-pass classifier.

III. Conclusions

The detection of mines in anything other than extremely benign environments is a hard problem for either a human or a computer. The MCM community routinely requires multiple looks at a field of interest in order to improve confidence in a mine-free assessment. To expect that a computer with a single look can do better than a trained operator using multiple looks is, perhaps, a little naive.

The multi-pass fusion system implemented by AST is not an automated detection/classification algorithm; it provides a framework to gather together the best-of-class CAD/CAC algorithms and to use them in a way that mimics the action of expert human operators. Working towards an efficient process to gather together information from multiple passes over an object in order to improve the quality of correct classification will, hopefully, move us one step closer to the removal of the trained operator from the minefield.

IV. References

[1] G. D. Godshalk and J. A. Smith, "Evolution of the Low Frequency Active Information Processor," Navy Journal of Underwater Acoustics, 47 (1), pp. 155–190, 1997.

[2] M. E. Tipping, "Sparse Bayesian Learning and the Relevance Vector Machine," Journal of Machine Learning Research, pp. 211–244. 2001.

[3] G. J. Dobeck, J. C. Hyland, and S. Le'Derick, "Automated Detection/Classification of Sea Mines in Sonar Imagery," in Detection and Remediation Technologies for Mines and Minelike Targets, Proceedings of SPIE, Vol. 3079, pp. 90–108, 1997.

J. Kent Harbaugh has been a Senior Research Scientist at Applied Signal Technology, Inc. since 2003. He has worked with acoustic target models, synthetic imagery, automatic target recognition algorithms and training, and change detection.

Andrew D. Wilby has been Chief Engineer with Applied Signal Technology, Inc. since 2004. Formerly with Ultra Electronics in the UK, he has been working within the sonar industry for 24 years, developing sensors and sensor processing systems for minehunting, ASW and commercial applications.