OPTIMIZING TIME-LIMITED MULTI-ASPECT CLASSIFICATION

M. Couillard Dept. of Applied Mathematics, University of Western Ontario, London, Canada
J.A. Fawcett DRDC Atlantic, Dartmouth, Nova Scotia, Canada
M. Davison Dept. of Applied Mathematics, University of Western Ontario, London, Canada
V.L. Myers NATO Undersea Research Centre, La Spezia, Italy

1 INTRODUCTION

After completing an initial survey, a sidescan sonar equipped vehicle may have limited additional time to revisit some contacts in order to increase the overall classification performance. This time limitation might be a result of operational constraints or remaining battery time in the case of an automated underwater vehicle. Therefore, given a list of contacts to revisit, a path planning algorithm minimizing the total travel time between these contacts is needed. This algorithm needs to be able to determine the optimal order of visit of the contacts but also the optimal points of approach in order to maximize the classification performance.

A study conducted by Zerr, Stage and Guerrero[1] using simulated sonar images of various objects and height profiles as features showed that the highest classification performance when imaging an object twice can be achieved with an angular increment of 90 degrees between the two images. Using real multi-aspect sonar images, we extend this study and consider the classification performance of all admissible secondary aspects, as by revisiting some of the contacts at suboptimal yet still satisfactory aspects, the overall survey time can be reduced. For different types of minelike objects, we study the correct classification rates and the error rates as functions of the angular difference between aspects. We also present an algorithm based on a travelling salesman approach with Dubins curves used to generate minimum time search paths.

2 CITADEL DATABASE

The data used in this study was gathered during the October 2005 Citadel sea trial which took place at the NATO Undersea Research Centre in La Spezia, Italy. During this experiment, many mine shapes and rocks were deployed and were surveyed using a high-frequency sidescan sonar to obtain detection and classification data[2]. The main asset of this trial was the DRDC Remote Minehunting System using the Dorado semi-submersible vehicle which towed a commercially available Klein 5500 sidescan sonar. The Dorado completed flower shaped patterns around the targets to collect multi-aspect imagery at a systematic number of ranges. This special flower survey pattern produced a significant number of images at increments of 30 degrees. The objects of interest for this study are cylinder shapes, Manta shapes, Rockan shapes and rocks. Figure 1 shows examples of multi-aspects sonar images of these objects. In total, we are using 180 cylinder shape images, 180 Manta shape images, 172 Rockan shape images and 93 rock images.

3 SEGMENTATION, FEATURES AND CLASSIFICATION

Using computer-aided classification, we wish to study correct classification probabilities and global error rates as a function of the angular increments between two images of a same object. The first step of this classification task is the segmentation of the sonar images into three distinct regions: highlight or target echo (sound scattered by the target by an active sonar), shadow (regions of low acoustic energy created by an object or seabed feature blocking the sound propagation) and

background or seabed. Several methods exist for segmenting high-resolution sonar imagery, most employing some type of Markov Random Field, for instance those described by Reed, Petillot and Bell[3] or Mignotte et al.[4]. In this study, Myers' method for segmentation[5] is used. The algorithm employs a fuzzy decision function to determine whether a pixel belongs to the highlight, shadow or background regions, where a pixel's class is "flipped" depending on the classes of its neighbours. The procedure is iterative and a pixel's membership in the highlight, shadow or background classes changes over time. This kind of procedure is very similar to a Markov Random Field with one fuzzy membership function based on pixel amplitude and a second fuzzy function controlling pixel neighbour interactions. This kind of method is sometimes called fuzzy relaxation. The algorithm requires an initial image thresholded into shadow, highlight and background. With the Citadel data, this automated segmentation algorithm produced satisfactory results. The shadow threshold was chosen at the 10th percentile of the cumulative image histogram and the highlight was thresholded at the 95th percentile.

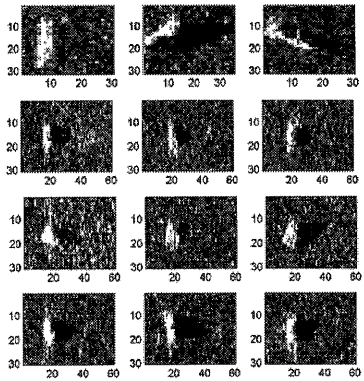

Figure 1: Multi-aspect sidescan sonar images. Line 1 to line 4: cylinder shapes, Manta shapes, Rockan shapes, rocks.

Using our segmented shadow regions, we can proceed to the second step of our computer-aided classification task by obtaining two feature sets. Our first feature set is based on the physical measurements of the objects. Each object observed on a shadow image is described by its length in the along track direction and a height profile in the across track direction. To account for potential noise in the shadow boundaries, these lengths and height profiles are filtered by considering only the region of the shadow falling inside a fitted ellipse having the same second moment as the segmented shadow. Our second feature set adds a series of shadow based features to the physical measurements. For example, the major and minor axes as well as the orientation of the fitted ellipse are used as features. We also use the shadow area and the shadow perimeter rescaled by the area and perimeter of the best fit ellipse. This rescaling is aimed at reducing the impact of range differences between some images of a given object in order to focus solely on the aspect difference.

With our feature sets, it is now possible to obtain computer-aided classification results. We use our four types of objects simultaneously in a 4-class classification problem. Both a nearest neighbour classifier and a kernel-based least square regression technique derived from support vector classifiers are used. The kernel-based regression classifier maps each training point $\vec{x}_i$, i = 1, .., n, with a Gaussian kernel function given by:

$$\kappa(\vec{x}_i, \vec{x}_j) = \exp\left(-\left\|\vec{x}_i - \vec{x}_j\right\|^2 / 2\sigma^2\right) \quad \forall\ j = 1,\ldots,n \tag{3.1}$$

Instead of using quadratic optimization to obtain the support vector weights, we use the Ridge regression[6] approximation given by:

$$\vec{\alpha} = (\kappa + \lambda I_n)^{-1}\vec{c} \tag{3.2}$$

where λ is a regularization parameter used to ensure that the inverse is not singular, κ is the complete kernel matrix and $\vec{c}$ is the class label matrix. This class label matrix is a n x C matrix of zeros, where C is the number of classes, with ones in the c^{th} column if the training point belongs to the c^{th} class. The regularization parameter is chosen to be λ = 0.01. When classifying based on the predicted labels, the maximum value among the four classes is used. For the Citadel database, the Ridge least square approximation was shown to produce classification performances comparable to the traditional quadratic optimization approach in addition to significantly reducing the computation time. Furthermore, multiple types of kernels were applied to the Citadel database and the Gaussian kernel was shown to be the one producing the highest classification performances.

To assess the impact on the classification performance of combining a secondary aspect to an original one, we first obtain single aspect classification results. We proceed exhaustively, testing each image individually. For a given object to be tested, training sets made of objects of all four classes are created. To ensure that the training will not be biased towards one of the classes, the same number of samples from each class is used in the training sets. This number is chosen as 50% of the number of samples in the class with the lowest number of observations. To ensure that our results are significant, we generate multiple random training sets for each image. The single aspect performance is obtained over all images of a given class and over all random training sets. For each class of mine shape, we obtain two classification performance measures: the hit rate and the global error rate. For a given class, the hit rate measures the percentage of images correctly classified while the global error rate combines the percentage of the objects of the other classes misclassified as the given class and the classification error of the given class.

We then compare these single aspect classification results to the classification results obtained when pairing images. Pairs of images are created based on a desired difference of aspects. The aspect of each image is computed by using the heading of the sonar according to the easting and northing coordinate system and the relative position of the target. The aspect of an object at port (starboard) side of the sonar is obtained by subtracting (adding) 90 degrees from the sonar heading. When we create a pair of images at a 0 degree increment, we combine two different images having the same aspect. A pair having an angular increment of 180 degrees is made of two images obtained from opposite sides of the object.

There are two possible approaches when classifying a pair of images. The first approach involves creating a new object to be classified by combining the features of the two images in a single vector. Starting with a given single aspect image and a desired angular increment, we create all possible augmented pairs. We test these pairs one at a time, creating a training set made of augmented pairs from the four classes having equivalent desired angular increment. As the number of training samples can vary between all four classes, we use the same approach as in the single aspect case by generating a series of random training sets using the same number of samples from each class. In this process, we need to ensure that neither of the two images being tested in the pair is repeated in the training set. The second, faster, approach is simply to fuse the single aspect classification probabilities obtained from the classifier according to the desired angular increment between the images. We experimented with four fusion rules: the mean, the maximum, the minimum and the product rules. Because our pairs are created using a difference in aspects, our performance curves will be symmetric around 180 degrees. As in the single aspect case, we obtain the hit and global error rates for each class of mine shapes.

4 RESULTS

Our classification results as a function of the angular increments between two aspects are summarized in figures 2, 3 and 4 for the cylinder shapes, Rockan shapes and Manta shapes respectively. Each figure shows the results obtained for the two types of classifiers and our first feature set based on physical measurements. These results were obtained with the approach combining the single aspect features into one augmented vector. The percentages displayed in the legends next to the classifier types refer to the single aspect classification rates. The acronym NNC refers to the nearest neighbour classifier while the acronym KBR refers to the kernel-based least square regression classifier. The results obtained with our second feature set incorporating additional shadow based features show comparable behaviour with higher hit rates and lower error rates. The spread of each curve is also smaller in comparison with this second feature set. For each single aspect or pair being tested, 100 random training sets were generated. To verify the stability of our classification rates, we increased the number of training sets from 10 to 100. The biggest change in classification rates was found to be less than 0.2%.

Comparing the single aspect and two aspect classification rates, we see that collecting a secondary look produces a significant increase in hit rate and a significant decrease in error rate for all mine shapes. Also, the kernel-based least square regression classifier yields better classification performances than the simple nearest neighbour classifier. For all mine shapes, the angular increments yielding the maximum hit rate also yield the minimum error rate. In the cylinder shapes case, as obtained by Zerr, Stage and Guerrero[1], the optimal classification performance is reached around angular increments of 90 and 270 degrees. The worst classification performance is observed for an angular increment of 180 degrees. Interestingly, we found that the results for the Rockan shapes are distinct from those of the cylinder shapes, the optimal classification performance being obtained around a 180 degrees increment. It would seem that to maximize the classification performance when faced with Rockan shapes, one should collect sonar images on opposed sides of the object. Finally, the Manta shapes case is more ambiguous reaching an optimal classification performance at 90 and 270 degrees with a nearest neighbour classifier, but reaching this optimal performance at 150 and 210 degrees with the kernel-based classifier. Overall, it seems that an acceptable classification performance can be reached for the Manta shapes as long as the two combined aspects are at more than 90 degrees apart from each other. It is worth noting that when considering a set of acceptable secondary looks, the nearest neighbour curve obtained for the Manta shapes shows an interesting behaviour around 90 and 270 degrees. If one allows for a ±30 degrees flexibility in the secondary looks, the performance reduction at 60 or 300 degrees is significant and these secondary looks should be avoided while the performance at 120 or 240 degrees could still be considered acceptable.

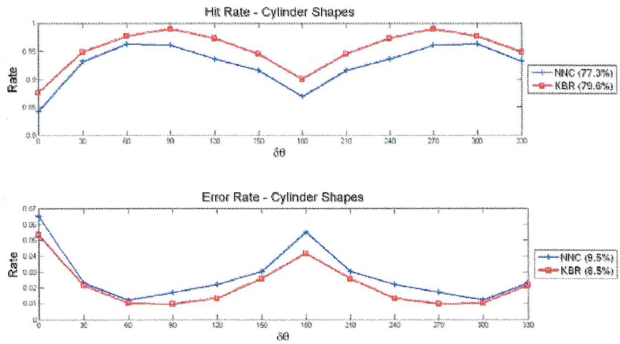

Figure 2: Cylinder shapes, hit rate and error rate as a function of the angular increments.

Using the kernel-based least square regression classifier, we also compared the classification results we obtained with the approach combining the single aspect features into one augmented vector to the approach fusing the single aspect classification probabilities. For the cylinder and Mantas shapes, the augmented feature approach outperforms the fusion of the single aspect classification probabilities for all four fusion rules considered at all angular increments. In contrast, in the case of Rockan shapes, the mean fusion rule and the maximum fusion rule outperform the augmented feature approach. Also, the product rule produces comparable classification performances, while the minimum rule underperforms the augmented feature approach.

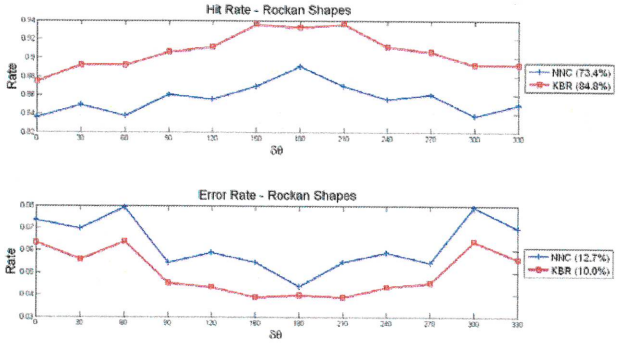

Figure 3: Rockan shapes, hit rate and error rate as a function of the angular increments.

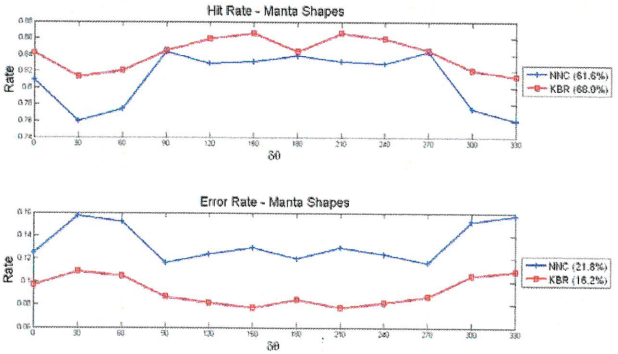

Figure 4: Manta shapes, hit rate and error rate as a function of the angular increments.

Finally, we also decided to study the classification performances as a function of the number of aspects. Results are shown in figure 5. Hit rates and error rates are obtained over all possible pairs, without considering specific angular increments. To manage the significant number of possible combinations of images for multiple aspects, we used the approach of fusing the output probabilities from the single aspect kernel-based least square regression classifier. A mean fusion rule was used. A 99% hit rate is obtained for the cylinder shapes after only three looks. Four looks are

needed to achieve the same hit rate in the case of the Rockan shapes. The Manta shapes prove to be the ones most difficult to classify, needing eight looks before reaching a 99% hit rate.

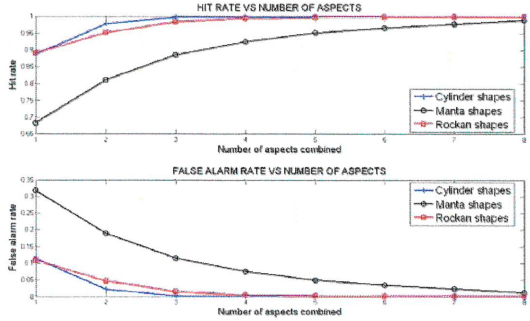

Figure 5: Hit rate and error rate as a function of the number of aspects combined.

5 OPTIMAL PATH PLANNING

After completion of an initial survey, computer-aided detection and classification algorithms generate a list of detections and associated classification probabilities. Contacts having similar probabilities of being minelike or non minelike might need to be revisited. Further classification probabilities discriminating between potential mine shapes allow us to use the secondary aspect classification performance curves shown in the previous section to identify preferred angles of approach along these contacts. The objective of optimal path planning is to combine the preferred angles of approach to each contact in order to minimize the total travelling time.

Minimizing the total time needed to visit each contact corresponds to the travelling salesman problem. This problem of discrete optimization is one of the most famous examples of nondeterministic polynomial-time hard problems. A brute force approach finding the minimum travel time by enumerating all possible permutations has a computational complexity of order n!, where n is the number of contacts to be visited, and becomes quickly impractical. An exact solution technique like dynamic programming[7] reduces the complexity of the problem to order $2^n n^2$ which is still impractical for large values of n. For such large values of n, random search heuristics[7] like simulated annealing or genetic algorithm might be used to yield satisfying solutions.

In itself, the travelling salesman problem is a complex one, but our sidescan sonar problem is even more complex. Assuming that the altitude and speed of the sonar remain constant, we are trying to find the minimum travelling time between waypoints described by a 2-dimensional location, but also by a desired heading along the contacts. Furthermore, the vehicle carrying the sonar will be limited by a maximum turning radius. Dubins[8] showed that under a maximum turning radius constraint, the path minimizing the travel time between two points with fixed headings will be given by a combination of maximum turning rate segments and straight line segments. Six such combinations, shown in figure 6, exist and only one of these six possibilities will yield the minimum time solution. This increases the computational complexity of the problem beyond that of the corresponding travelling salesman problem which involves only straight lines.

Also, the difficulty of our sidescan sonar path planning is further increased by the simultaneous port/starboard side imaging capability of sidescan sonar. For a given waypoint and desired secondary aspect, two opposed headings along the contact can be followed to yield the same aspect. Each of these waypoints can be at different seabed ranges from the contact and we make the assumption than any range value between 20 and 60 meters will yield the same classification

performance. Some flexibility in the desired aspects, based on the curves shown in section 4, is also allowed in order to reduce the total travel time while preserving the classification performance. The point from where the sonar equipped vehicle is starting is known, as well as the heading of the vehicle at that point. Finally, we impose an additional straight line segment condition along the targets to ensure proper imaging. This straight line segment is chosen to be at least 10 meters.

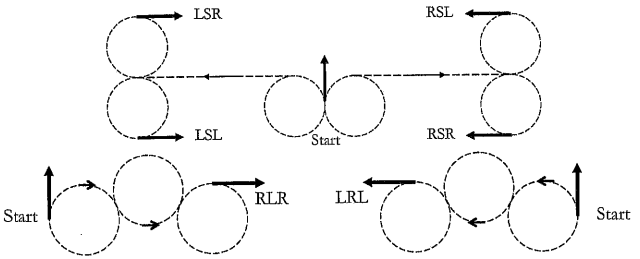

Figure 6: 6 types of Dubins curve. R: right, L: left, S: straight.

To deal with such a complex path planning problem, we constructed a three step heuristic solution technique. This simple and fast heuristic approach decomposes the problem in a series of more tractable sub-problems and produces realistic solutions. The first step is to fix the order of visit of the contacts by solving the traditional travelling salesman problem using the Euclidean distance between the contacts. Then, we move away from the contacts themselves to consider admissible waypoints along each of them. When two or more contacts are less than 120 meters apart, some waypoints might be overlapping. If these overlapping waypoints have similar desired headings, they are clustered together to reduce the number of points to visit and therefore the total travel time. The last step uses optimal Dubins curves between contacts to select the optimal waypoints of the final path solution. We proceed iteratively from the fixed starting point up to the last contact to be visited, minimizing the Dubins travelling time to the next two contacts and using this intermediate solution to select the optimal waypoint of the next contact to be visited. Figure 7 shows an example of minimum time travelling solution obtained with our heuristic algorithm. In this example, we assume the initial search followed a lawn mowing pattern with legs along the north-south direction. We also assume a minefield made of cylindrical mines, therefore the optimal secondary aspects should be at 90 or 270 degrees. Such secondary aspects will be collected by following straight line segments in the east-west direction along the targets. The solid line path illustrates the solution we obtained with no flexibility in the desired secondary looks. This path also shows the clustering of contacts by simultaneous port/starboard side imaging. The dashed line represents the path obtained by allowing ±30 degrees flexibility in the desired secondary aspects. For this specific example, allowing the use of suboptimal but yet satisfactory aspects reduces the overall survey time by more than 5%.

6 SUMMARY

After completing an initial survey, a sidescan sonar equipped vehicle may have limited additional time to revisit some contacts in order to increase the overall classification performance. Given a list of contacts to revisit, a path planning algorithm minimizing the total travel time needs to be able to find the optimal order of visit of the contacts and also the optimal points of approach along these contacts. While a specific secondary look might be optimal, knowledge of the classification performance of all secondary looks is needed as by revisiting some of the contacts at suboptimal but still satisfactory aspects, the overall survey time can be reduced.

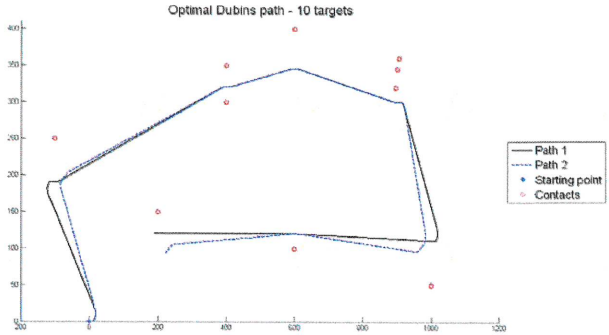

Figure 7: Minimum travelling time solutions using Dubins curves.

Using the Citadel multi-aspect sidescan sonar images of various mine shapes and rocks, we constructed secondary look classification curves to be used in conjunction with a path planning algorithm. We found that cylinder shapes are best discriminated from the other shapes by combining two aspects at 90 or 270 degrees apart. This is in contrast to the Rockan shapes case where the maximum classification performance is achieved when combining two images from opposed sides of the object. The Manta shapes classification performance seems to be high when the two combined aspects are at more than 90 degrees apart from each other. We also created a heuristic algorithm to quickly find a path minimizing the total travel time by combining a travelling salesman approach to Dubins curves between waypoints with desired headings. Our solutions take into account the simultaneous port/starboard side imaging capability of sidescan sonar by allowing multiple targets close to each other to be imaged simultaneously.

7 REFERENCES

1. B. Zerr, B. Stage and A. Guerrero, Automatic Target Classification Using Multiple Sidescan Sonar Images of Different Orientations, SACLANT CEN Memorandum SM-309. (1997)
2. V.L. Myers, B. Evans, D. Hopkin, B. Zerr, J.A. Fawcett, A. Crawford, P. Hagen and M. Flores, The CITADEL and SWIFT joint experiments with unmanned vehicles for minehunting and high-frequency sensors, Proc. 7th International Conference on Technology and the Mine Problem, MINWARA association, Monterey. (2006)
3. S. Reed, Y. Petillot and J. Bell, An automatic approach to the detection and extraction of mine features in sidescan sonar, IEEE Journal of Oceanic Engineering, Vol. 28, no. 1, 90-105. (2003)
4. M. Mignotte, C. Collet, P. Perez and P. Bouthemy, Sonar image segmentation using an unsupervised hierarchical mrf model, IEEE Transactions on Image Processing, Vol. 9, 1216-1231. (2000)
5. V.L. Myers, Sonar image segmentation using iteration and fuzzy logic, CAD/CAC 2001 Conference Proc., Halifax, NS. (2001)
6. J. Shawe-Taylor and N. Cristianini, Kernel Methods for Pattern Analysis, Cambridge University Press, 462 pages. (2004)
7. R. Rardin, Optimization in Operations Research, 1st ed Prentice Hall, 919 pages. (2000)
8. L.E. Dubins, On Curves of Minimal Length with a Constraint on Average Curvature, and with Prescribed Initial and Terminal Positions and Tangents, American Journal of Mathematics, Vol. 79, 497-516. (1957)