

# On the Effects of Synthetic-Aperture Length on SAS Seabed Segmentation

David P. Williams and Johannes Groen  
 NATO Undersea Research Centre  
 Viale San Bartolomeo 400, 19126 La Spezia (SP), Italy  
 {williams,groen}@nurc.nato.int

## Abstract

*In this work, we quantify the relationship between synthetic-aperture length (or equivalently, along-track resolution) and seabed segmentation performance experimentally for real synthetic aperture sonar (SAS) imagery. The seabed segmentation algorithm employed uses wavelet-based features, spectral clustering, and a variational Bayesian Gaussian mixture model. It is observed that for this approach, the correct seabed segmentation rate drops approximately ten percentage points for each halving of the along-track resolution between 3 cm and 96 cm. Moreover, changing the along-track resolution has the most significant effect on rocky seabeds.*

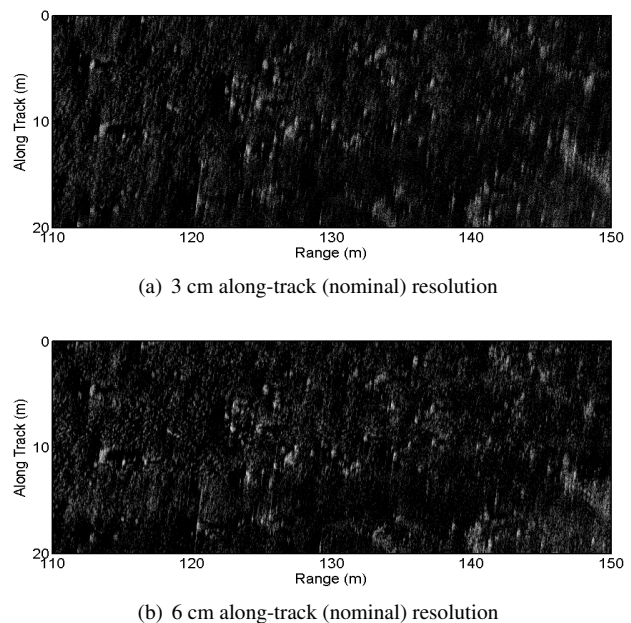
## 1. Introduction

Synthetic aperture sonar (SAS) can provide high-resolution imaging of underwater environments by simulating a (synthetic) aperture that is longer than the real physical aperture of the array. In theory, the along-track resolution of a SAS image is determined solely by the dimensions of the transducers in the array [1]. That is, high along-track resolution imaging can be achieved that is independent of wavelength and range.

In practice, the main limiting factor for successful SAS image processing is undesired motion (*i.e.*, surge, sway, heave, roll, pitch, and yaw) of the sonar platform and environmental conditions (*e.g.*, bottom type, sound-speed variations) [2]. The quality of a SAS image degrades when accumulated motion errors cannot be corrected for. Furthermore, the longer the synthetic-aperture length that is used in the processing, the more opportunities there are for motion errors to accumulate to an excessive level. Processing that uses a short synthetic-aperture length will limit the accumulated motion errors. However, the tradeoff for this increased robustness is decreased image resolution.

If a SAS image is processed with an excessively long synthetic-aperture for the amount of platform motion

present during the data collection, the quality of the resulting SAS image will degrade and appear defocused. Therefore, when significant platform motion is present, a shorter synthetic-aperture length should be employed in the processing in order to create imagery that is not characterized by motion artifacts. Although the nominal resolution of the resulting image will be lower, the quality of the image (and true achieved resolution) can be better. An example of this scenario is shown in Fig. 1.



**Figure 1. SAS images of the same area of seabed that result from processing at different (nominal) along-track resolutions. The images are defocused because of significant undesired platform motion, so decreasing the synthetic-aperture length improves the image quality.**

One study [3] used synthetic data to investigate the performance of a target detection and classification algorithm as a function of image resolution (when along-track and across-track resolution were kept equal). However, the manner in which along-track resolution (or equivalently, synthetic-aperture length) affects attempts at performing *seabed segmentation* has never before been investigated. This work attempts to determine the impact of using a shorter synthetic-aperture (or equivalently, lower along-track resolution) for SAS image processing, in the context of seabed segmentation. Specifically, we seek to quantify the impact of synthetic-aperture length on seabed segmentation by using real, measured sonar data.

The remainder of this paper is organized in the following manner. Sec. 2 briefly describes the process by which lower-resolution SAS images are produced from a high-resolution SAS image. Sec. 3 briefly describes the unsupervised seabed segmentation algorithm employed in the study. Sec. 4 shows experimental results, on real SAS images, of the effects of different along-track resolutions on seabed segmentation performance. Sec. 5 contains an extensive discussion of the experimental results, and also the ways in which they can be exploited in practice. Concluding comments and directions for future work are given in Sec. 6.

## 2. Synthetic Aperture Sonar Processing

Synthetic aperture sonar (SAS) processing is the process by which acoustic returns received by hydrophones are converted into sonar imagery. For a sufficiently high pulse repetition rate, the area of seabed insonified by consecutive sonar pings will overlap. The reconstructed image of a point on the seabed will then be the coherent summation of returns over an integration angle of  $\beta$ , which is equal to the beamwidth of the transmitted signal.

SAS provides one with the opportunity to achieve constant along-track resolution independent of range. The relationship between along-track resolution,  $\delta$ , and integration angle is [4]

$$\delta = \frac{c}{4f_c \tan(\beta/2)}, \quad (1)$$

where  $c$  is the speed of sound in water, and  $f_c$  is the baseband (center) frequency of the transmitted signal. (It should be noted that the signal beamwidth  $\beta$  that is used places an upper limit on the along-track resolution theoretically achievable.)

In order to achieve constant along-track resolution independent of range, the synthetic-aperture length used must be a function of range. The synthetic-aperture length,  $L$ , needed to reconstruct the image of a point at range  $r$  is given by

$$L = 2r \tan(\beta/2). \quad (2)$$

Thus, to resolve a point at longer range will require the synthesis of a correspondingly longer aperture length.

Substituting (2) into (1) will then provide the relationship between aperture length and along-track resolution,

$$\delta = \frac{cr}{2f_c L}. \quad (3)$$

If one has reconstructed a SAS image at along-track resolution  $\delta_0$ , it is relatively simple to reconstruct the same image at a lower along-track resolution,  $\delta_i > \delta_0$ . Taking the two-dimensional Fourier transform of the reconstructed complex SAS image will result in a plane-wave decomposition. In this form, a signal amplitude is associated with each frequency,  $f$ , and wavenumber,  $k_x$ , pair. Each plane wave's angle of arrival,  $\theta$ , can then be computed via  $\sin \theta = \frac{ck_x}{2\pi f}$ .

To reconstruct a SAS image at a lower along-track resolution,  $\delta_i$ , a smaller integration angle,  $0^\circ < \beta < 90^\circ$ , must be employed (*cf.* (1)). This fact implies that only a subset of the plane waves will contribute to the reconstructed image. Specifically, only those plane waves with appropriate arrival angles,  $|\theta| \leq \beta/2$ , are retained. After filtering out the plane waves deemed to not contribute to the image when a smaller beamwidth is used, the two-dimensional inverse Fourier transform will recover the reconstructed SAS image at lower along-track resolution.

Filtering out those plane waves on the edges of the acceptable region effectively reduces the synthetic-aperture length that is used in the processing. As a result, the motion compensation requirements become less stringent, and the potential ill effects of undesired platform motion are lessened. The tradeoff for this increased robustness to platform motion is decreased along-track resolution.

In summary, if a full-resolution SAS image has already been constructed, the image reconstruction process for lower along-track resolution images can be greatly accelerated. In this work, we exploit this insight to study the relationship between synthetic-aperture length — or equivalently, along-track resolution (*cf.* (3)) — and seabed segmentation performance.

## 3. Unsupervised Seabed Segmentation

Seabed segmentation is the process by which one segments (an image of) a typically large area of seabed into different regions based on the characteristics of the seabed. For example, an area may be segmented into flat seabed, rocky seabed, and seabed characterized by sand ripples. The need to perform this seabed classification in an automated manner with no human intervention is motivated by the desire to conduct fully autonomous mine countermeasures missions with an autonomous underwater vehicle.

In this work, the “atomic” unit for seabed segmentation is assumed to be a  $2 \text{ m} \times 2 \text{ m}$  area of seabed. That is, each

2 m×2 m area of seabed corresponds to one data point. This particular size was chosen as a compromise among several factors. The larger the area chosen, the more likely that a single data point will have the unfavorable property of containing multiple types of seabed. However, if the area is too small, the distinguishing characteristics of the seabed that indicate a certain seabed type may be lost.

The unsupervised seabed segmentation algorithm employed in this work is a slightly modified form of that proposed in [5]. The method first represents each 2 m × 2 m area of seabed by a vector of 16 features derived from a wavelet decomposition [6]. Spectral clustering [7, 8] is then performed on this data, which transforms it into a lower-dimensional space via an eigendecomposition. This transformed data is then clustered using a completely unsupervised variational Bayesian Gaussian mixture model (VB-GMM) [9]. Seabed segmentation is effected in this step by assigning each data point to the mixture component that maximizes its posterior probability.

## 4. Experimental Results

### 4.1. Data Set

In April-May 2008, the NATO Undersea Research Centre (NURC) conducted the Colossus II sea trial in the Baltic Sea off the coast of Latvia. During this trial, high-resolution sonar data was collected by the MUSCLE autonomous underwater vehicle (AUV). This AUV is equipped with a 300 kHz sonar with a specified 60 kHz bandwidth that can achieve an along-track image resolution of approximately 3 cm (and an across-track image resolution of approximately 2.5 cm). The sonar data was subsequently processed into synthetic aperture sonar (SAS) imagery. For the present study, we use four images, each of which spans an area of 50 m × 110 m, from this data set. These particular images were chosen because they each contain relatively balanced amounts of flat seabed, rippled seabed, and rocky seabed.

### 4.2. Experimental Set-Up

In this study, we seek to quantify the relationship between the along-track resolution of a SAS image (or equivalently, synthetic-aperture length) and seabed segmentation performance. As noted above, four real, measured SAS images are employed to conduct this experiment.

Each of the four images undergoes the same following treatment. A given “full-resolution” SAS image, which corresponds to the image with a 3 cm along-track resolution, is re-processed (assuming a reduced beamwidth for the transmitted signal) to produce a series of images at successively lower along-track resolutions. Specifically,

five lower-resolution images are created from each full-resolution image, at along-track resolutions of 6 cm, 12 cm, 24 cm, 48 cm, and 96 cm.

Each of the six versions of a given original image then undergo the following identical processing. The wavelet features are extracted for each 2 m × 2 m area of seabed (*i.e.*, data point) of an image. Spectral clustering is then applied to this data, with  $m = 2$  eigenvectors retained so the transformed data is mapped to  $\mathbb{R}^2$ . This transformed data is then used to learn a  $k = 3$  component GMM using the variational Bayesian method. By assigning each data point to the mixture component for which its posterior probability is a maximum, segmentation of the seabed is effected.

In this study, only  $m = 2$  eigenvectors were retained because this choice permits a faster eigendecomposition and reduces the model complexity of the variational Bayesian GMM (and hence the learning time). Moreover, it was also observed that the magnitudes of the eigenvalues decrease very quickly, indicating that most of the relevant information is contained in the first few eigenvectors.


The variational Bayesian approach to learning a GMM possesses the attractive property that the number of mixture components,  $k$ , to employ can be determined in a principled manner via the evidence [9]. For these experiments, however, we fix the number of mixture components to  $k = 3$  for purposes of assessment, so that uniform comparisons can be made across images.

This choice is justified by the fact that the three main seabed types of interest for purposes of seabed segmentation are (i) flat seabed, (ii) seabed characterized by sand ripples, and (iii) rocky seabed. That is, assigning an area of seabed to the correct one of these seabed types is generally the most important segmentation. Since we are interested in examining how seabed segmentation performance changes as the along-track resolution of the SAS image varies, it makes sense to consider only these three general seabed types when performing segmentation.

It should be emphasized that no labeled training data is used in the process, as the segmentation approach is completely unsupervised. Because the methods are unsupervised, no explicit correspondence between clusters and seabed types exists. However, for purposes of evaluating the segmentation results here, one can easily assign a correspondence between the three aforementioned seabed types and clusters.

### 4.3. Seabed Segmentation Results

Because the SAS images we employ in this study are real, we possess no ground-truth information regarding the true seabed type. Rather than attempting to manually ground-truth the images, we take an alternative approach. Since the objective of the study is to assess how



**Figure 2. SAS image 1 (left column) with different along-track resolutions and the associated seabed segmentation results (right column).**

**Figure 3. SAS image 2 (left column) with different along-track resolutions and the associated seabed segmentation results (right column).**

segmentation performance varies as a function of along-track resolution, we treat the segmentation results of the full-resolution images (*i.e.*, those at 3 cm along-track resolution) as ground-truth. This choice allows us to measure the change in segmentation performance as a function of along-track resolution in an objective manner.

The SAS images at each along-track resolution and their associated seabed segmentation results are shown together for images 1, 2, and 3, in Figs. 2, 3, 4, respectively. (Space constraints prevent showing the results for image 4.) For

the segmentation results, the colors white, gray, and black essentially correspond to flat seabed, rippled seabed, and rocky seabed, respectively.

The overall segmentation performance for each image is compactly summarized graphically in Fig. 5. The overall segmentation performance for each seabed type is summarized in Fig. 6.

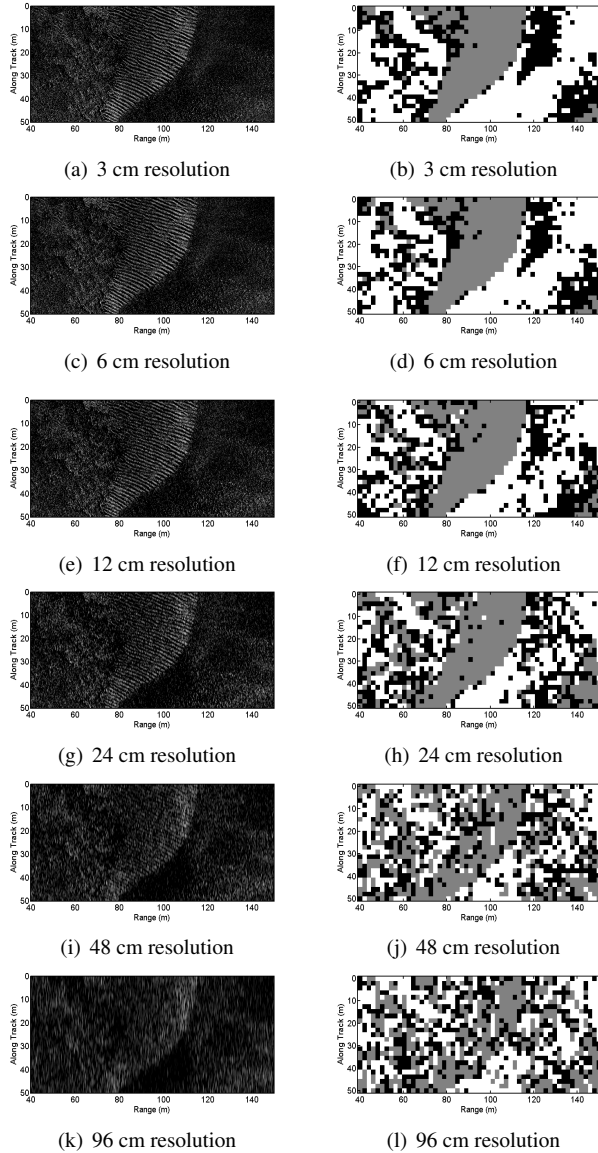


Figure 4. SAS image 3 (left column) with different along-track resolutions and the associated seabed segmentation results (right column).

## 5. Discussion

### 5.1. Effect of Along-Track Resolution on Seabed Segmentation

As one would expect, the results in Sec. 4.3 suggest that the along-track resolution of a SAS image has a strong impact on the resulting seabed segmentation. From Fig. 5, it can be observed that the correct seabed segmentation rate

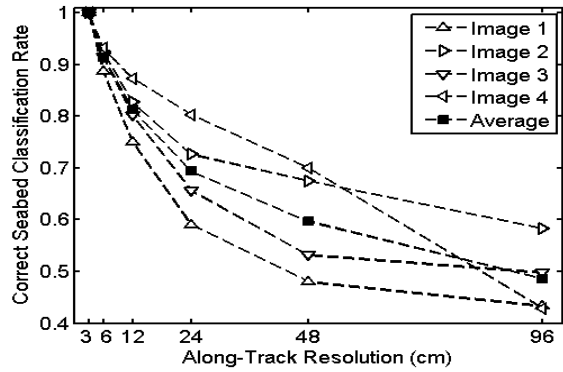


Figure 5. Correct seabed classification rate of each image as a function of along-track resolution, when the segmentation results of the 3 cm resolution cases are treated as ground truth.

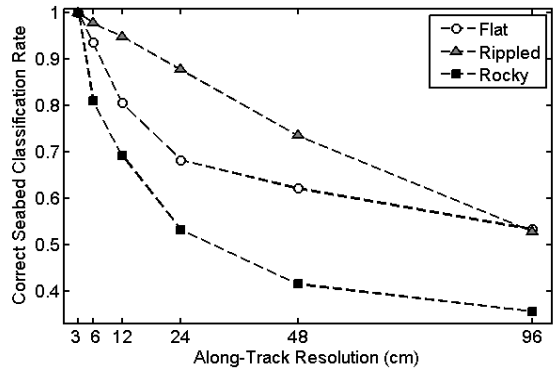


Figure 6. Average correct classification rate of each seabed type as a function of along-track resolution, when the segmentation results of the 3 cm resolution cases are treated as ground truth.

drops approximately ten percentage points for each halving of the along-track resolution. For example, the correct segmentation rate drops to approximately 90% when changing from a 3 cm along-track resolution to a 6 cm along-track resolution, and it drops again to approximately 80% at a 12 cm along-track resolution.

From Fig. 6, it can be seen that changing the along-track resolution has the biggest effect on the rocky seabed type. As along-track resolution decreases, the rocky seabed type experiences the highest misclassification rate. This phenomenon can perhaps be attributed to the fact that rocky seabed typically contains high-frequency characteristics, which would be altered if the resolution is reduced. Ad-

ditionally, (from confusion matrices not shown due to space constraints) when data points belonging to either flat or rippled seabed types are misclassified, they tend to be misclassified as rocky seabed. In general, decreasing the along-track resolution of an image affects the rippled seabed type the least, in terms of seabed segmentation.

It must be emphasized that these results are for a particular seabed segmentation approach (including a specific set of features and a particular clustering method). If a different segmentation approach were employed, the performance would of course not be identical. Nevertheless, the appearance of the SAS images at different along-track resolutions suggests that the resolution-performance dependence observed in these experiments is reasonable.

## 5.2. Utility of Experimental Results

The observed experimental results can subsequently be exploited in several situations, two of which we describe in the following. Processing sonar data to achieve high-resolution SAS images is a computationally demanding task. But by processing the sonar data at a lower along-track resolution, computation time can be reduced. This fact can be vitally important if the data processing is to be performed onboard a sonar-equipped vehicle that has limited processing capability.

For example, an AUV may be tasked to locate a safe route (*e.g.*, that avoids rocky areas where mine detection is unfeasible) from one location to another, for subsequent assets to utilize. To achieve this goal in a timely manner, the AUV must adapt its route based on the data that it collects and processes onboard during the mission. By knowing the anticipated degradation in seabed segmentation performance for processing at reduced along-track resolutions, an appropriate resolution can be selected at which to process the data that balances different conflicting constraints (*e.g.*, mission time, processing capability, segmentation accuracy).

These results can also be exploited when significant undesired platform motion (*e.g.*, surge, sway, heave, pitch, roll, and yaw) occurs during data collection. The motion may be such that the SAS processing cannot be performed at the highest along-track resolution (using the longest synthetic-aperture) because the resulting image would contain motion artifacts (*cf.* Fig. 1). So instead, the processing must be done at a lower nominal resolution (shorter synthetic-aperture) that results in adequate image quality, albeit at a less than optimal resolution. The results obtained in this work establish the drop in seabed segmentation performance that is likely to be experienced by processing at the reduced resolution. Therefore, one can use these results to help determine whether it is more worthwhile to re-collect the data at sea (hopefully without plat-

form motion) so that better segmentation performance can be achieved by processing the new data at a higher resolution, or if the present (admittedly sub-optimal) segmentation performance is sufficient to not warrant undertaking a second expensive data collection experiment.

## 6. Conclusion

In this work, the relationship between synthetic-aperture length (or equivalently, along-track resolution) and seabed segmentation performance was quantified experimentally for real, measured SAS imagery. It was observed that the correct seabed segmentation rate drops approximately ten percentage points for each halving of the along-track resolution between 3 cm and 96 cm. Moreover, changing the along-track resolution has the most significant effect on rocky seabeds.

It was also explained how the relationship between along-track resolution and seabed segmentation performance obtained in this work can be exploited in practical, operational settings. Namely, the insight provided by the results permits informed decisions to be made about the computation needed onboard an AUV, and about the potential benefits of re-collecting data that is marred by significant undesired platform motion.

## References

- [1] M. Hayes and P. Gough. Broad-band synthetic aperture sonar. *IEEE J. Oceanic Engineering*, 17(1):80–94, 1992.
- [2] A. Bellettini and M. Pinto. Theoretical accuracy of synthetic aperture sonar micronavigation using a displaced phase-center antenna. *IEEE J. Oceanic Engineering*, 27(4):780–789, 2002.
- [3] R. Manning. Small object classification performance of high-resolution imaging sonars as a function of image resolution. In *Proc. MTS/IEEE OCEANS*, pages 2156–2163, 2002.
- [4] J. Groen, R. Hansen, H. Callow, J. Sabel, and T. Sæbø. Shadow enhancement in synthetic aperture sonar using fixed focusing. *IEEE J. Oceanic Engineering*, 34(2):1–16, 2009.
- [5] D. Williams. Unsupervised seabed segmentation of synthetic aperture sonar imagery via wavelet features and spectral clustering. *Proc. IEEE International Conference on Image Processing*, 2009.
- [6] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, 1999.
- [7] M. Meilă and J. Shi. Learning segmentation by random walks. In *NIPS 13*, pages 873–879, 2000.
- [8] A. Ng, M. Jordan, and Y. Weiss. On spectral clustering: Analysis and an algorithm. In *NIPS 14*, pages 849–856, 2001.
- [9] M. Beal and Z. Ghahramani. The variational Bayesian EM algorithm for incomplete data: Application to scoring graphical model structures. *Bayesian Statistics*, 7:453–464, 2003.