# The advantage of segmentation in SAR image compression

Marco Cagnazzo, Giovanni Poggi, Luisa Verdoliva

Università di Napoli Federico II
Dipartimento di Ingegneria Elettronica e delle Telecomunicazioni
email: poggi@unina.it

***Abstract*- SAR images are severely degraded by speckle, and filtering is therefore a common practice. Filtering is especially useful before compression, to avoid spending valuable resources to represent noise; unfortunately, it also degrades important image features, like region boundaries. To overcome this problem, one can resort to a segmentation-based compression scheme, which allows one to preserve region boundaries, carry out intense denoising, and improve overall performance. In this work we assess the potential of segmentation-based compression through controlled experiments on synthetic SAR images. Numerical results seem to confirm the validity of this approach.**

## I. Introduction

Efficient and reliable compression techniques for remote sensing imagery become more and more necessary as the number and size of images to be archived and transmitted over general purpose networks grow constantly. Even the most recent standards, such as JPEG2000, cannot always guarantee a good compression performance as they neglect the peculiar characteristics of the data.

This is certainly true for SAR images, which are characterized by a wide dynamic range and are affected by a strong speckle that destroys the statistical regularities on which compression techniques rely. As a matter of fact, virtually all compression schemes proposed for SAR images (e.g., [1,2]) include a filtering phase (despeckling) which significantly improves compression performance. Indeed, it has long been observed that filtering and compression are tightly related processing steps [1]. Filtering out noise reduces the entropy of the image thus allowing for its more compact and faithful representation; likewise, compression tends to smooth out high frequencies and therefore provides a certain amount of filtering.

Unfortunately, filtering and compression not only remove noise-related high frequencies, but also cause a severe distortion of the highest-frequency components of the desired image. These correspond to region boundaries and isolated impulses, namely, some of the most meaningful features of the image, that one should try to preserve as much as possible.

The usual approach to deal with this problem is to resort to edge-preserving filters (e.g., [3,4]) which reduce their smoothing action near region boundaries. Of course, their effectiveness depends on their ability to (implicitly) identify the boundaries. Even wavelet-based despeckling [5] is based on the implicit ability to distinguish between noise components (high-frequency in all directions) and boundaries (high-frequency only in the edge-crossing direction).

It should be clear, therefore, that a more fundamental approach to image filtering and compression requires the prior identification of region boundaries or, which is the same, the segmentation of the image in homogeneous regions. Image segmentation would guarantee a number of relevant advantages:

1. important information about the region boundaries is retained in the segmentation map, which can be efficiently coded in lossless modality;
2. noise can be more easily removed in inner regions (there is no risk of damaging boundaries) with a clear improvement of image quality;
3. compression of the texture alone, without boundaries and speckle, can be much more effective, leading to higher overall performance;
4. the segmentation map is an added value for the user, and comes at no additional cost.

Segmentation, compression and filtering have many deep interactions, and intense research is under way to exploit them [6]. Indeed, they all converge toward the same broad goal, the extraction and compact representation of the most relevant features of an image, and as such they should always be carried out jointly. It must be pointed out that segmentation remains a formidable problem for which foolproof algorithms are not yet available. Nonetheless, given the intense research in the field (e.g. [7]) and the steady progress of concepts and technology, it is not unreasonable to expect that reliable image segmentation algorithms will be at hand in a few years.

This work aims at studying and quantifying the potential advantages provided by image segmentation in the filtering and compression of SAR images. To keep all variables under control, we define an abstract image model and work with synthetic images. Assuming that a perfect segmentation is available (and leaving aside the problem of *how* to obtain it) we then compare the performance of a segmentation-based compression scheme with that of a reference algorithm in a variety of operating conditions.

0-7803-7536-X/02/$17.00 (C) 2002 IEEE

In Section II, the image model is defined, and our segmentation-based coding scheme is described together with a reference conventional scheme. Section III presents and discusses the results of a number of experiments, and Section IV draws conclusions and outlines future work.

## II. IMAGE MODEL AND CODING SCHEMES

We synthesize the image as the sum of three components, a region map, a set of textures, one for each region, and additive noise (this fits SAR images if the log-intensity is taken).

An image is assumed to comprise $K$ homogeneous regions, and the segmentation map labels each pixel as belonging to one of the regions. By representing each region with its mean value, we obtain a rough approximation of the image, call it $M$, in which each region is perfectly flat, and the boundary between regions are step-like.

Each region is then characterized by a particular texture process, (obtained by passing white gaussian noise through a low-pass filter with given cut-off frequencies). The desired original image is then $X = M + T$, where $T$ is the collection of the various region textures. White gaussian noise $N$ is finally added, with its power as the only relevant parameter, to obtain the final image $Y = M + T + N$. Fig.1 shows images $X$ and $Y$ for our running example.

Fig.2 (top) shows the block diagram of a conventional coder. In the following we will use the Lee filter [3] for denoising, and JPEG2000 [8] for compression.

The block diagram of the proposed segmentation-based coding scheme is shown in Fig.2 (bottom). The segmenter singles out image $M$, where each region is approximated by its mean value, and subtracts it from $Y$, leaving only texture and noise. Of course, a real segmenter would carry out this task only approximately, but in our experiments we assume an ideal behavior. The image $M$ must be encoded without loss of information (we use a contour tracer followed by chain coding and entropy coding) but for a reasonably smooth map this encoding cost is quite limited, for our running example it amounts to 0.1 bit/pixel. Denoising and compression blocks are the same as before[1].

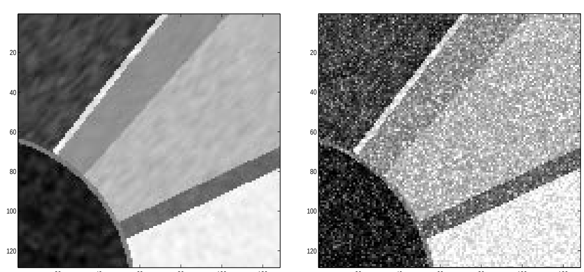

Figure 1: Original and noisy image.

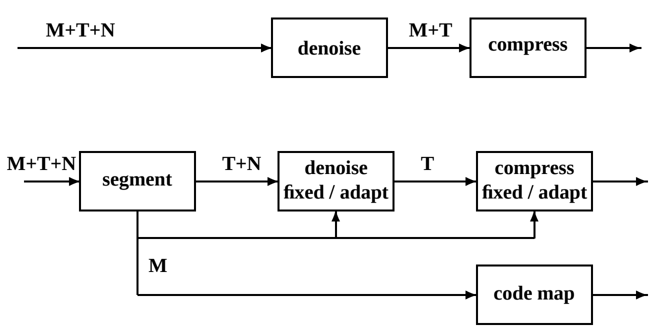


Figure 2: Reference and segmentation-based encoding schemes.

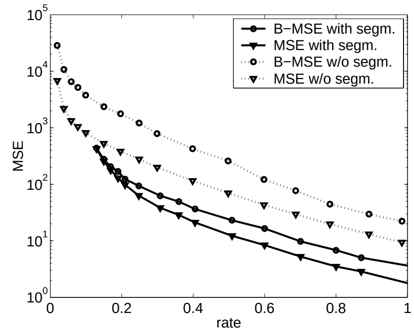


Figure 3: Global and boundary MSE for noise-free image.

## III. NUMERICAL RESULTS

In the first experiment we consider a noise-free image ($Y = X$), and compress it by JPEG2000 with no previous filtering. In Fig.3 (dotted lines) we show the mean-square error (MSE) as a function of the encoding rate (R) in bit/pixel. In addition, in order to measure edge degradation, we also report the boundary-region mean-square error (B-MSE), which is computed only on pixels that are within 3 points of an edge. It results that, even in the absence of noise, the edges are significantly degraded by the compression process. Fig.3 also reports MSE and B-MSE for the segmentation-based coder (solid lines). Despite the additional cost of segmentation, the performance gain is striking, especially for the boundary regions. As a matter of fact, MSE and B-MSE now are closer together (the increase in the latter is only due to the high frequencies associated with the change of texture from region to region) confirming that segmentation is especially valuable for boundary preservation.

Let us now consider the noisy image $Y$ (SNR=6.98 dB) and, first of all, let us study the case in which no filtering is carried out. Fig.4(a) shows that JPEG2000 has a much harder time now compressing the $Y$ image, as the MSE (always global) decreases much more slowly. What is worse, the MSE computed with respect to the *desired* image $X$ stops decreasing after a given rate, when the encoder begins devoting most of its resources to faithfully represent the added noise!

Therefore, for such noisy images, pre-filtering seems to be a mandatory step. On the other hand, even an edge-preserving

[1] It is worth noting that both filtering and compression techniques could make use the map information, as suggested in Fig.2, to adapt to the statistical behavior of each component region, with further performance improvement.

filter, like Lee's, tends to smooth out edges. This is clear from the data of Tab.1. After Lee-filtering image $Y$ the MSE decreases significantly from 11181 to 2117 with a 5×5 window, down to 1369 with a 9×9 window. However, this does not hold for the B-MSE which, after an initial reduction from 11062 to 3893, begins increasing again up to 5904, confirming that the filter, while reducing noise, is also smearing region boundaries.

The picture is completely different when the segmentation map is available. Both MSE and B-MSE decrease consistently as the filter window grows, reaching much smaller values than in the previous case, especially near the boundaries.

| window size | w/o segm. | | with segm. | |
|---|---|---|---|---|
| | MSE | B-MSE | MSE | B-MSE |
| 5×5 | 2117 | 3893 | 1815 | 1759 |
| 7×7 | 1479 | 4439 | 982 | 937 |
| 9×9 | 1369 | 5120 | 725 | 728 |
| 11×11 | 1401 | 5578 | 622 | 634 |
| 13×13 | 1489 | 5904 | 582 | 581 |

Table.1: MSE after Lee filtering.

Fig.4(b) reports the coding performance (only global MSE) when the Lee filter is used. The MSE is evaluated both with respect to the desired image $X$ and to the noisy image $Y$. Comparing the results with those of Fig.4(a), it is clear that filtering improves performance. In fact, although the MSE with respect to the noisy original $Y$ is about unchanged, it is much smaller when the *desired* image $X$ is considered. Such original is available here only because synthetic images are considered, but a similar behavior could be expected of real images. As for the comparison between conventional and segmentation-based encoding, the difference is, once again, quite large. It is worth underlining once more that the strong noise level considered puts an insurmountable limit to the performance and after 0.25 bit/pixel (in this case) increasing further the encoding resources is a pure waste.

Finally, to gain insight about visual quality Fig.5 compares the test image encoded at 0.25 bit/pixel without (left) and with (right) segmentation. It is clear that segmentation guarantees a superior quality, especially around the region boundaries.

## IV. Conclusions and Future Work

This work suggests that segmentation has a huge potential in image coding as it allows one to separate edges from texture and to process both pieces of information in the most appropriate way. This is especially true when a strong noise component is present, as is the case for SAR images, since noise and edges often occupy the same frequencies and the former cannot be filtered without impairing the latter.

Of course, this is only an exploratory work, and a number of problems are open for further investigation, concerning the image model, the processing blocks, and the experimental setting.

Despite all the foreseeable problems, we feel that the segmentation-based approach will prove advantageous in most practical situations.

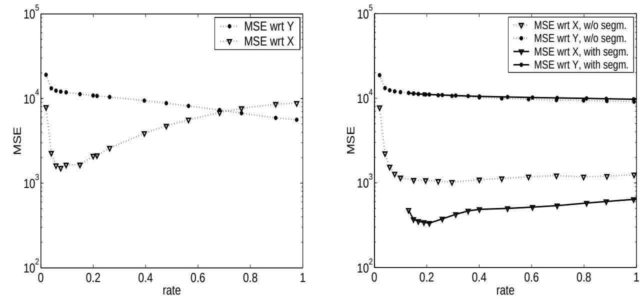


Figure 4: MSE for the compressed noisy image: (a) without filtering (b) with filtering.

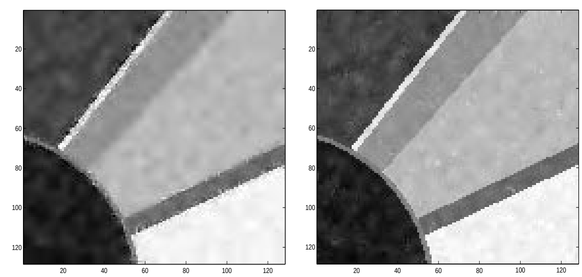

Figure 5: Decoded images at 0.25 bit/pixel.

## References

[1] D.Wei, J.E.Odegard, H.Guo, M.Lang, C.S.Burrus: "Simultaneous noise reduction and SAR image data compression using best wavelet packet basis," *Proc. IEEE ICIP-95*, vol.3, pp.200-203.

[2] Z.Zeng, I.G.Cumming: "SAR image data compression using a tree-structured wavelet transform," *IEEE Trans. Geosci. Remote Sensing*, March 2001, pp.546-552.

[3] J.S.Lee, "Digital image enhancement and noise filtering by use of local statistics," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-2, March 1980.

[4] D.T.Kuan, A.A.Sawchuk, T.C.Strand, P.Chavel, "Adaptive noise smoothing filter for images with signal-dependent noise," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-7, March 1985.

[5] D.Donoho: "De-noising by soft thresholding," *IEEE Trans. Information Theory*, vol.41, May 1995, pp.613-627.

[6] K.L.Oehler, R.M.Gray: "Combining image compression and classification using vector quantization," *IEEE Trans. Pattern Anal. Machine Intell.*, vol.17, pp.461-473, 1995.

[7] G.Poggi, A.R.P.Ragozini: "Image segmentation by tree-structured Markov random fields", *IEEE Signal Proc. Letters*, vol.6, pp.155-157, July 1999.

[8] http://maestro.ee.unsw.edu.au/ taubman/kakadusoftware/ index.html