

Report: Ultrasound speckle reduction using adaptive wavelet thresholding

Laura Fuentes

Université Paris Saclay

Gif-sur-Yvette, France

laura.fuentes-vicente@universite-paris-saclay.fr

1 INTRODUCTION

Medical ultrasonography serves as a crucial imaging modality in clinical diagnosis, relying on the principle of acoustic impedance. This tool employs high-frequency acoustic waves (>20kHz) emitted by a machine to interact with organs and tissues, subsequently producing echoes. The distinctive echo patterns are harnessed to compute a detailed 2-D image of the body tissue. This technique is indispensable due to its cost-effectiveness, non-invasiveness, and rapid results, providing real-time images of body tissues [1]. However, the efficiency of medical ultrasonography can be impacted by various constraints, including speckle, acoustic shadows, artifacts, etc. This study specifically addresses the issue of speckle, a granular multiplicative noise that degrades texture information and obscures details such as lines, edges, and boundaries in an image. This noise complicates segmentation tasks and feature extraction. Speckle arises from the interference of echoes from unresolvable scatterers randomly dispersed throughout the image. Despite its controversial nature, speckle also encapsulates texture information dependent on anatomical tissues. Hence, despeckling has become a crucial pre-processing step to prevent the loss of texture information. This step consists in removing speckle noise while preserving the essential features of the image.

Previous research has extensively explored the challenges associated with denoising ultrasound images, categorizing such efforts into two main approaches: single scale and multi-scale. Single scale methods involve the direct application of filters, including the Weiner filter, linear filters (e.g., median filter), and non-linear filters (e.g., statistic filters). Conversely, multi-scale techniques operate on a set of sub-images derived from an original image, which is the primary focus of this paper. This approach enables a more profound understanding of frequency amplitude variations along three directions (horizontal, vertical, and diagonal) [2]. Multi-scale applications encompass two key steps: acquiring sub-image sets and despeckling. Techniques like ridgelet, curvelet, and shearlet can be used for sub-image generation, but this paper specifically highlights the efficacy of the wavelet transform. This approach allows for simultaneous high-resolution analysis of time and frequency localized features, making it particularly adaptable to transient signals [3]. In addressing despeckling, the authors utilize a wavelet thresholding technique, involving the modification of wavelet coefficients according to a specified threshold value. However, challenges arise from these implementations, including information loss, computationally complex thresholds, and the difficulty of achieving simultaneous enhancement and despeckling. The objective is to devise a straightforward and efficient approach for determining an adaptive threshold.

The authors introduce a novel multi-scale and adaptive technique featuring a unique thresholding function, which progressively reduces wavelet coefficients to zero when their values fall below a specified threshold. This innovative approach is grounded in two key principles: firstly, leveraging the statistical properties of the image across various decomposition levels, and secondly, operating under the assumption that speckle predominantly manifests in low-valued wavelet coefficients [4].

2 METHODS

The authors introduced a multi-scale approach to tackle ultrasound image despeckling. Nevertheless, two previous assumptions are considered, gaussian nature of noise and concentration of the majority of high frequency components in diagonal sub-band. The procedure proposed can be split into 6 different steps [Appendix A.5].

Step 1: Speckle is considered to be multiplicative granular noise. In order to turn multiplicative into additive noise, the authors propose to consider the **log transformation** of the image:

$$f(x, y) = f'(x, y) \cdot \eta(x, y)$$

$$\Rightarrow \log(f(x, y)) = \log(f'(x, y)) + \log(\eta(x, y)) = L(f(x, y))$$

Multi-scale implementations involve two consecutive steps.

Step 2: The first step involves **wavelet decomposition**, enabling multi-resolution analysis by breaking down the ultrasonic image into a set of sub-images. These images capture information at different resolutions using a low-pass filter (approximate components: A) and high-pass filters (detailed components: H, V, D). The authors propose more precisely a 2-level decomposition on different types of mother wavelets [Figure 1]:




Figure 1: Two level decomposition function [4]

- First level: $[A_1, H_1, V_1, D_1] = \text{DWT}(L(f(x, y)))$
- Second level: $[A_2, H_2, V_2, D_2] = \text{DWT}(A_1)$

The second part of the multi-scale implementation involves thresholding methods to address despeckling. The process involves modifying wavelet coefficients with values below a given threshold (Step 4) and subsequently computing the inverse wavelet transform to

obtain the despeckled image (Step 5). Thresholding methods can be categorized into two groups: hard and soft. Hard techniques set to zero wavelet coefficients below the threshold, leaving other coefficients unchanged [Appendix (4)]. In contrast, soft techniques set wavelet coefficients below a given threshold to zero, and values above the threshold are reduced towards zero [Appendix (3)]. It is noteworthy that the soft technique is generally preferred over the hard approach because reducing coefficients to zero helps mitigate artifacts in the image.

Step 3: The **choice of the threshold** is critical for distinguishing noise from significant information [5], and over the years, several functions have been proposed [Appendix A.3]. However, the authors introduced a global thresholding function based on diagonal sub-band at several decomposition levels to apply on detailed components ($H_i, V_i, D_i \in \{1, 2\}$).

$$\tau = 2\beta \left[\frac{|\sigma_i^2 - \sigma_D^2|}{\sigma_i} \right], \beta = \sqrt{2\text{Log}(L)} \quad (1)$$

where L represents the total number of pixels from the image, and σ_i^2, σ_D^2 refer to the noise variances of the noisy image and the detailed diagonal sub-band using [Appendix (5)].

Step 4: Regards the application of **thresholding function** to wavelet coefficients of detailed components. The paper proposes and adaptative thresholding function that scales down gradually to zero wavelet coefficients with values below threshold without altering other coefficients. This introduces a new class merging soft and hard techniques:

$$X_T(w) = \begin{cases} w \cdot e^{n_l(|w|^{-t})} & \text{if } |w| < \tau \\ w & \text{if } |w| \geq \tau \end{cases} \quad (2)$$

The resulting denoised sub-bands are:

- For level-1: H_{n1}, V_{n1}, D_{n1}
- For level-2: H_{n2}, V_{n2}, D_{n2}

Step 5: To obtain a despeckled image, the **inverse wavelet transformation** is computed. This captures the denoising modifications and returns the denoised version of the log of the image.

- At level-2: $A_{n1} = \text{IDWT}(A_2, H_{n2}, V_{n2}, D_{n2})$
- At level-1: $Lf' = \text{IDWT}(A_{n1}, H_{n1}, V_{n1}, D_{n1})$

Step 6: Since the log transformation was applied to the initial noisy image, a final step involves computing the **exponential** of the log-transformed denoised image to bring it back to the original domain. The final despeckled image is then obtained: $f' = e^{Lf'}$.

3 RESULTS

To assess the effectiveness of thresholding functions and the algorithm's adaptability to real-case scenarios, the authors undertook a comprehensive evaluation of the method's performance on both synthetic data and a real-case dataset in Matlab 17.

First, the authors analyzed 37 wavelet filters, including 15 bi-orthogonal, 9 Daubechies, 8 Symlets, and 5 Coeflets, with the aim of finding the most effective mother wavelet. The analysis revealed that Daubechies (db1) as the top performing choice, followed by Symlets, Coeflets, and Bi-orthogonal mother wavelets. For the

thresholding function referenced as (2), they selected coefficients n_l with $n_1=1$ and $n_2=0.5$ based on prior research [6]. They deployed four different metrics to evaluate and quantify each denoising algorithm's performance.

3.1 Synthetic data

The incorporation of synthetic data aimed to establish a ground truth for evaluating the denoising efficacy of the proposed algorithm. The authors deliberately introduced noise variance levels (ranging from 0.01 to 0.04) into an initially noise-free image. A comprehensive qualitative and quantitative comparison against established filtering methods (e.g., average and median filtering, SRAD) was conducted, incorporating both soft and hard thresholding techniques.

The observations indicated that as the speckle increased, the preservation of structural and textural elements in the despeckled images diminished. Notably, the proposed method consistently outperformed other techniques at every noise level [Appendix Figure 5]. Results demonstrated superior preservation of subtle features, such as edges, boundaries, and texture, as evidenced by higher SSIM scores, more pronounced features, and the lowest MSE. However, with escalating noise variance, the PSNR values decreased, implying that the despeckled images still retained some residual noise.

3.2 Real data

To assess the adaptability of the model to real-world scenarios, they authors deployed a dataset of ultrasound images from Samsung H60 Ultrasound Scanner. They observed a notable enhancement in image quality using the proposed method that consistently outperformed other methods across all considered metrics [Appendix Figure 6.a]. Improves in SSIM scores indicated a greater similarity and substantial preservation of structures in the despeckled images compared to the original ones confirmed by EKI scores. Moreover, higher PSNR values suggested effective noise reduction, leading to fewer errors and a lower MSE associated with the proposed method.

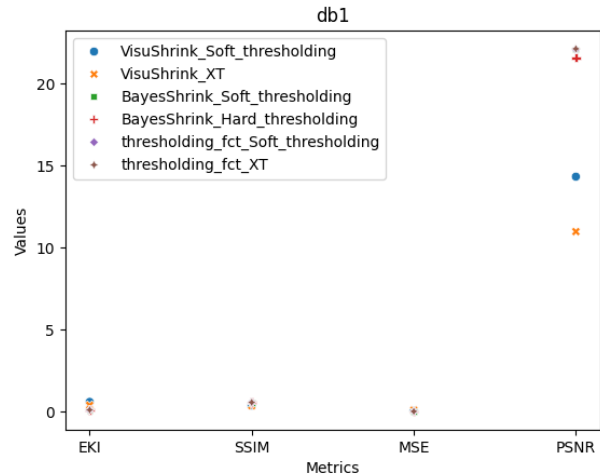
3.3 Our results

I have created a dedicated GitHub repository containing the code required to implement the proposed algorithm and replicate various figures from the associated paper.

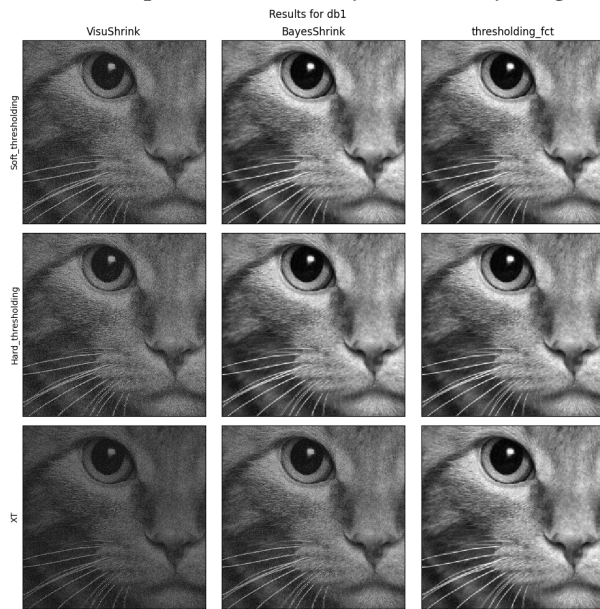
To validate the functionality of our implementation, initial tests were conducted using a synthetically noised image of a cat. To generate the synthetic noisy image, I introduced Gaussian noise with a mean (μ) of 0 and a standard deviation (σ) of 5 to a clean and noise-free image featuring a cat. These tests involved the exploration of four different mother wavelets—Daubechies1 (db1), Daubechies2 (db2), Symlets2 (sym2), and Biorthogonal (bior1.1). Furthermore, a comprehensive comparison was undertaken on a combination of various thresholding techniques outlined in [Appendix A.3], such as VisuShrink, BayesShrink [7], the proposed technique (1), and a set of thresholding functions [Appendix A.2]: Soft thresholding, Hard thresholding, and the proposed technique (2). First, we may remark that the best performing mother wavelet is as expected db1. However, noteworthy deviations in performance order were

observed, with bior1.1 achieving the second-best outcome. Secondly, I estimated each combination’s performance, and selected the two best algorithms for each metric [Figure 2.a]. Subsequently, we observe that independently from the mother wavelet, the paper’s threshold coupled to soft and XT are consistently selected featuring the better scores. Interestingly, the proposed threshold coupled with XT exhibited lower scores for EKI, similar results for SSIM and MSE. Indeed, PSNR stands out as the metric revealing the most noticeable differences. It identifies the threshold function with XT and soft thresholding, and Bayes with soft thresholding as the top-performing techniques. In addition, we can visually appreciate in [Figure 2.b) that the best results for hard thresholding and XT is obtained using the paper’s thresholding function. For hard thresholding, Bayes shrink and the paper’s threshold perform similarly. For soft thresholding, Bayes seems to perform the best. A gray effect is observed in VisuShrink technique that increases when applying the thresholding function XT. Such an effect appears as well on BayesShrink when implementing XT.

Subsequently, I assessed the adaptability of the algorithm using a limited set of ultrasound images sourced from Kaggle [8]. The dataset comprised 2 normal ultrasound images, along with 1 benign and 1 malignant image. To establish a comprehensive evaluation, we averaged the results for each metric and technique across all images in the dataset. Initiating the tests, I conducted a comparative analysis of the proposed algorithm against existing denoising techniques, as illustrated in [Figure 3]. I employed a range of filters for comparative analysis, incorporating Average, Median, Weiner, and Wavelet thresholding, with the Bayes Shrinkage threshold coupled to soft thresholding. Notably, the proposed technique consistently outperformed in all metrics, aligning with our expectations [Appendix Figure 6.b). Specifically, we achieved the lowest MSE value. The algorithm demonstrated as well the highest values for SSIM and PSNR indicating a higher preservation of features and better reconstruction quality. Regarding EKI, the proposed technique consistently demonstrated superior results, with Bayes shrinkage also achieving the highest scores. I continued testing several mother wavelets for the proposed algorithm. I found that db1 and bior1.1 remained the top-performing choices, consistent with the results from synthetic trials. However, it is essential to note that vertical and diagonal detailed components were null in several instances due to a lack of variation. This limitation could directly affect the interpretability of certain results. It appears that in the paper, detailed components are represented as black squares in the schema. This representation might suggest that certain detailed components are also composed of null values. However, since the observation is not explicitly mentioned or discussed in the paper, there remains some uncertainty regarding the appropriateness of the results. As for the synthetic image, I additionally tested combinations of thresholds and thresholding functions to compare the performance of our implementation [Appendix Figure 7.a). We remark that the proposed combination remains upon the best performing combinations but to our surprise, combinations deploying VisuShrink and BayesShrink yielded same performances. As mentioned before, wavelet decomposition affect interpretability but may also limit considerably the wavelet’s thresholding capacity to modify images as most coefficients yield already null values. This phenomena could explain the




a- Filter’s performance on a synthetic noisy image



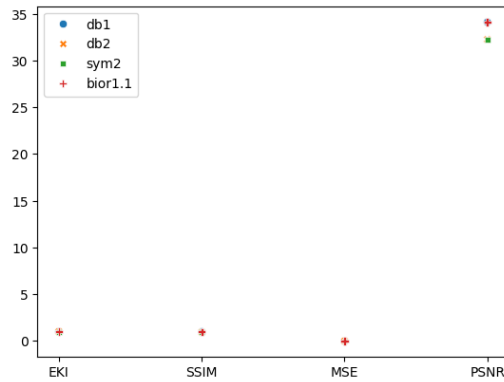
b- Synthetic noisy image denoising with 9 different filters

Figure 2: Results on synthetic data

observed similarity in results for various threshold combinations, reducing the overall variability. However, despite this constraint, the proposed technique consistently outperformed other methods. This underscores its effectiveness, even under conditions that might limit the full utilization of certain wavelet coefficients in the thresholding process. In [Figure 3.b), it is evident that most filters yield similar outputs, consistent with the metric results. A closer examination of [Figure 4] reveals that the filter does not induce drastic changes to the image. Notably, there is an improvement in feature definition in darker regions, enhanced contrast in certain areas, homogenization around similar tissue regions, and an intensification of the gray effect in specific areas.



a- Filter’s performance on ultrasound images



b- Wavelet’s performance on ultrasound images

Figure 3: Results on ultrasound images




Figure 4: Comparison of noisy ultrasound and despeckled image

4 CONCLUSION

The paper introduces a novel multi-scale approach that employs a thresholding function combining soft and hard thresholding to address speckle noise in ultrasound images. The utilization of wavelet

transform allows for the dissection of a discrete-time signal into various scale components and orientations, revealing both coarse and fine details. The modification of these detailed components, assuming that noise primarily resides in lower frequencies, serves as an effective pre-processing step to enhance image quality in computer vision. This approach outperforms single-scale methods by considering an increasing amount of information during image modification. Specifically, gradually scaling coefficients under a given threshold to zero effectively removes speckle while preserving more information. Nevertheless, the despeckling techniques show limitations when applied to detailed components with null values. Implementations in such cases are hindered due to lack of possible modifications underscoring a scenario where the proposed approach may not achieve optimal results. Experimentations demonstrated on both synthetic and real datasets, showcase enhanced clarity with preserved edges and boundaries. The filtered images exhibit reduced speckle, improved edge preservation, and a heightened level of detail compared to the original images. Performance measures affirm the superior preservation of structures. Despite the successful outcomes, the implementation process encountered challenges. A notable contradiction emerged between the global threshold presented in the paper’s methodology and the suggestion, in the step-by-step presentation, to compute the threshold individually for each detailed component $(H_i, V_i, D_i, \forall i \in \{1, 2\})$. Additionally, the paper did not specify how to compute σ_i corresponding to the noise variance of the noisy image. To address these issues, I delved into the literature, explored potential meanings, and selected the most effective approach based on a set of images. Another challenge involved computation time. In fact, some images took longer than expected to be despeckled, leading to a significant reduction in the dataset size. In conclusion, the findings presented in this article underscore and provide valuable insights that contribute to ultrasound image despeckle, paving the way for continued advancements in medical imaging.

REFERENCES

- [1] *Ultrasound image: definition*. URL: <https://www.nibib.nih.gov/science-education/science-topics/ultrasound>.
- [2] Savita Gupta, RC Chauhan, and SC Saxena. “Robust non-homomorphic approach for speckle reduction in medical ultrasound images”. In: *Medical and Biological Engineering and Computing* 43 (2005), pp. 189–195.
- [3] Rajesh Ganesan, Tapas K Das, and Vivekanand Venkataraman. “Wavelet-based multiscale statistical process monitoring: A literature review”. In: *IIE transactions* 36.9 (2004), pp. 787–806.
- [4] Anterpreet Kaur Bedi and Ramesh Kumar Sunkaria. “Ultrasound speckle reduction using adaptive wavelet thresholding”. In: *Multidimensional Systems and Signal Processing* (2022), pp. 1–26.
- [5] S Grace Chang, Bin Yu, and Martin Vetterli. “Adaptive wavelet thresholding for image denoising and compression”. In: *IEEE transactions on image processing* 9.9 (2000), pp. 1532–1546.
- [6] G Andria et al. “Linear filtering of 2-D wavelet coefficients for denoising ultrasound medical images”. In: *Measurement* 45.7 (2012), pp. 1792–1800.
- [7] Stefan Van der Walt et al. “scikit-image: image processing in Python”. In: *PeerJ* 2 (2014), e453.
- [8] *Kaggle ultrasound dataset*. URL: <https://www.kaggle.com/datasets/anaselmasyr/datasetbusiwithgt>.
- [9] David L Donoho and Iain M Johnstone. “Ideal spatial adaptation by wavelet shrinkage”. In: *biometrika* 81.3 (1994), pp. 425–455.
- [10] David L Donoho and Iain M Johnstone. “Adapting to unknown smoothness via wavelet shrinkage”. In: *Journal of the american statistical association* 90.432 (1995), pp. 1200–1224.
- [11] Umme Sara, Morium Akter, and Mohammad Shorif Uddin. “Image quality assessment through FSIM, SSIM, MSE and PSNR—a comparative study”. In: *Journal of Computer and Communications* 7.3 (2019), pp. 8–18.

A APPENDIX

A.1 Quick-links

- Ultrasound speckle reduction using adaptive wavelet thresholding [4]
- GitHub repository

A.2 Thresholding functions

A.2.1 Soft thresholding.

$$\hat{w} = \begin{cases} 0 & \text{if } |w| < t \\ \text{sign}(w)(|w| - t) & \text{if } |w| \geq t \end{cases} \quad (3)$$

A.2.2 Hard thresholding.

$$\hat{w} = \begin{cases} 0 & \text{if } |w| < t \\ w & \text{if } |w| \geq t \end{cases} \quad (4)$$

A.2.3 Exponential thresholding.

$$X_T(w) = \begin{cases} w \cdot e^{n_l(|w| - T_{k_l})} & \text{if } |w| < T_{k_l} \\ w & \text{if } |w| \geq T_{k_l} \end{cases}$$

A.3 Threshold definitions

A.3.1 VisuShrink. [9] [10]

Minimizes the maximum possible error over a fixed number of pixels in an image. This method offers a global threshold that is defined as follows:

$$T_{univ} = \sigma_{im} \sqrt{2 \log(L)}$$

Where L represents the size of the image (width×height), and σ_{im} the standard deviation of the noisy image.

A.3.2 BayesShrink. [5]

This method is based on the assumption that: for every sub-band, wavelet coefficients can be described by a Generalized Gaussian Distribution. The threshold value minimizes the Bayesian risk:

$$T_{Bayes}(\sigma_x) = \frac{\sigma_l^2}{\sigma_x}$$

Where σ_x defines the standard deviation of the respective sub-band and

$$\sigma_l = \left[\frac{\text{median}(|D_l|)}{0.6745} \right]^2 \quad (5)$$

A.4 Metrics Formulas

In order to address denoising evaluation, four different metrics are used to assess the algorithm's performance [7] [11]. In the first three metrics, we will call the initial noisy image f and the denoised image g .

A.4.1 Structural Similarity Index (SSIM). Analyzes the preservation of structures in the filtered image while depicting the similarity persistence. This function takes as input the original image acquired and filtered image, and is based on luminance, contrast and structural terms parameters.

$$l(f, g) = \frac{2\mu_f\mu_g + c_1}{\mu_f^2 + \mu_g^2 + c_1}$$

$$c(f, g) = \frac{2\sigma_f\sigma_g + c_2}{\sigma_f^2 + \sigma_g^2 + c_2}$$

$$s(f, g) = \frac{\sigma_{fg} + c_3}{\sigma_f\sigma_g + c_3}$$

$$c_1 = (0.01 \times L)^2, c_2 = (0.03 \times L)^2 \text{ and } c_3 = c_2/L$$

A.4.2 Edge Keeping Index (EKI). Studies the amount of edges preserved in the filtered image when compared to the original noisy image. High EKI values indicate an efficient preservation of edges.

$$EKI = \frac{D(\Delta f - \overline{\Delta f}, \Delta g - \overline{\Delta g})}{\sqrt{D(\Delta f - \overline{\Delta f}, \Delta g - \overline{\Delta g})}}$$

$$\text{Where } D(f, g) = \sum_{x=1}^m \sum_{y=1}^n f(x, y)g(x, y)$$

A.4.3 Mean Squared Error (MSE). Measures the amount of variation between the despeckled image and the reference image. The lower the MSE value, the lower the error.

$$MSE = \frac{1}{mn} \sum_{x=1}^m \sum_{y=1}^n [f(x, y) - g(x, y)]^2$$

A.4.4 Peak Signal-to-Noise Ratio (PSNR). Quantifies the amount of noise suppressed in the filtered image. The higher the value the better the quality of the output image.

$$PSNR(dB) = 10 \log_{10} \left(\frac{\text{peakval}^2}{MSE} \right)$$

Where, peakval can be selected by the researcher or chosen from a range of values defined by the image datatype.

A.5 Step-by-step algorithm

Algorithm 1 Image despeckle

- Step 1:** Log transformation of the image
- Step 2:** Wavelet decomposition
- Step 3:** Choice of threshold
- Step 4:** Thresholding function implementation
- Step 5:** Inverse wavelet transformation
- Step 6:** Exponential transformation of the denoised image

A.6 Figures




Figure 5: Comparison of average performance measures for test images [4]


Performance measures	Average filtering	Median filtering	SRAD	Shearlet with SRAD	Shearlet with MSRAD	NLM filtering	Wavelet shrinkage method	Proposed method
EKI	0.342303	0.326886	0.41669	0.76203	0.459838	0.64107	0.6386	0.838143
SSIM	0.698353	0.798411	0.663866	0.74334	0.65550	0.99929	0.9097	0.999517
PSNR (dB)	30.62995	27.13495	25.1017	23.1511	31.32264	34.56615	32.5861	36.93762
MSE	0.012599	0.01234	0.110242	0.085602	0.178275	0.005700	0.2282	0.001829

a-Paper results [4]


	paper	bayes filter	weiner filter	average filter	median filter
EKI	0.998462	0.998085	0.806110	0.455949	-0.256040
SSIM	0.992341	0.953422	0.964555	0.879736	0.970562
PSNR	34.107628	32.176712	27.271601	32.207617	29.322205
MSE	0.000498	0.000702	0.001906	0.000613	0.001243

b- Our results

Figure 6: Comparison of the performance of the filtering methods



a- Wavelet's performance on ultrasound images



b- Implementation of different filters on ultrasound images

Figure 7: Results on ultrasound images