ELSEVIER

Pattern Recognition Letters 21 (2000) 917–926

Pattern Recognition Letters

www.elsevier.nl/locate/patrec

A Rayleigh reconstruction/interpolation algorithm for 3D ultrasound ☆

João M. Sanches *, Jorge S. Marques

IST/ISR, Av. Rovisco Pais, 1049-001, Lisbon, Portugal

Received 14 June 1999; received in revised form 29 May 2000

Abstract

This paper describes an algorithm for the reconstruction of 3D medical objects from ultrasound images. Reconstruction is performed by filtering and interpolating the available data using an optimal Bayesian criterion. A Rayleigh model is adopted to describe the image formation process. © 2000 Elsevier Science B.V. All rights reserved.

Keywords: Image reconstruction; Medical imaging; Multiplicative noise models

1. Introduction

Ultrasound imaging has a widespread use in medicine, e.g., in pre-natal and cardiology diagnosis. Its popularity stems from the fact that it is a low cost, non-ionizing and non-invasive technique (Quistgaard, 1997). Furthermore, it allows an interactive visualization of the anatomic structures which is not possible in other more expensive and sophisticated techniques such as CT, MRI or PET. The main weaknesses of ultrasound imaging are related to the low spatial resolution of the images and with the presence of multiplicative noise (speckle) generated by the interaction of the acoustic wave with the tissues.

Currently, most ultrasound exams are performed using the so-called B-mode scan which provides images of cross-sections of the human body along inspection planes defined by the doctor. 3D ultrasound, i.e., the reconstruction of 3D information from ultrasound data is not a widespread technology yet but it is receiving considerable attention since it provides additional information about the organs geometry and therefore allows an easier interpretation of the available data. The problem of 3D ultrasound can be formulated as follows: given a set of 2D ultrasound images corresponding to cross-sections of the human body, a 3D volume of data should be reconstructed. In most cases it is assumed that the images are obtained using conventional free-hand ultrasound equipment complemented with a spatial locator which measures the position and orientation of the ultrasound probe at each instant of time (Nelson et al., 1999). This information characterizes the inspection plane associated with each image as shown in Fig. 1.

There are two approaches for 3D reconstruction: surface based and volume based methods

☆ Partially supported by Heart 3D Project.

* Corresponding author. Tel.: +351-1-841-8195.
E-mail address: jmrs@alfa.ist.utl.pt (J.M. Sanches).

0167-8655/00/$ - see front matter
PII: S0167-8655(00)00053-2

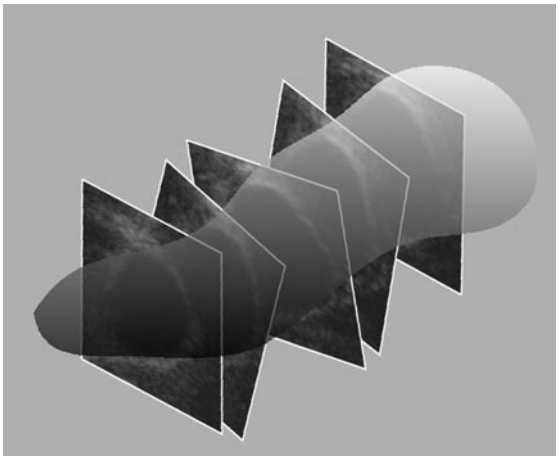

Fig. 1. Scanning geometry.

(Nelson et al., 1999). The first approach tries to estimate the organs surfaces from features detected in the ultrasound images (e.g., contours) (Raya and Udupa, 1990; Tagore, 1999; Treece et al., 1999). Therefore, only the organ's surface is reconstructed in this approach. Deformable surface models have also been used for surface estimation in medical applications (e.g., see McInerney and Terzopoulos, 1996; Delingette, 1998). The second approach is based on the estimation of the whole volume of data. A cuberille model is usually adopted (Chen et al., 1985) i.e., the 3D space is divided in cubic cells each of them characterized by an intensity to be estimated. The estimation process is usually performed in two steps (Rohling et al., 1999): bin filling and hole filling. The first is a noise reduction operation. It combines the information of the planes which intersect each voxel. The second step is an interpolation operation since it assigns a value to the voxels which were not observed during the medical exam. Several methods have been proposed for the interpolation step e.g., the voxel nearest neighbor (VNN), pixel nearest neighbor (PNN), the distance weight interpolation (DW) and radial basis function interpolation (RBF). A detailed survey of these methods can be found in (Rohling et al., 1999).

This paper addresses volume reconstruction using a set of ultrasound images and a statistical method of reconstruction (Herman and Kuba, 1999; Gooley and Barret, 1992). Unlike previous approaches, we propose a joint performance of noise reduction and image interpolation. Furthermore, instead of using different criteria, both operations optimize the same objective function. It is assumed that the function to be estimated belongs to a vector space with known basis functions different from the cubic support and constant basis functions used in the cuberille model. The unknown coefficients are obtained by Bayesian estimation methods which account for the noise reduction and image interpolation effects which were previously described. Furthermore, a non Gaussian (Rayleigh (Burckhardt, 1978)) model is adopted to describe the multiplicative type of noise observed in ultrasound data.

The paper is organized as follows. Section 2 formulates the reconstruction problem. Section 3 describes the parameter estimation methods. Section 4 presents an experimental evaluation of the reconstruction algorithm with synthetic and medical data and Section 5 concludes the paper.

2. Problem formulation

It will be assumed that the function f describing tissue reflectivity belongs to a class of admissible functions defined in a spatial domain, $\Omega \subset R^3$, i.e., $f : \Omega \to R$. Furthermore, it is assumed that the set of admissible functions is a finite dimension vector space $\boldsymbol{F}$ with known basis functions, $(b_i : \Omega \to R)$. Each function $f \in \boldsymbol{F}$ can be expressed as a linear combination of the basis functions,

$$f(x) = B(x)^{\mathrm{T}} U, \tag{1}$$

where $B(x) = [b_1(x), b_2(x), \ldots, b_N(x)]^{\mathrm{T}}$ is a vector of basis functions and $U = [u_1, u_2, \ldots, u_N]^{\mathrm{T}}$ is a $N \times 1$ vector of coefficients. The problem addressed in this paper can be stated as follows: we want to estimate f, given a sequence of ultrasound images, with known sensor positions and orientations. This sequence is used to obtain a set of data points, $V = \{v_i\}$, with geometric and intensity information i.e., $v_i = (y_i, x_i)$, where y_i is the intensity measured by the ultrasound probe at the position $x_i \in R^3$. It is assumed that each b_i is a local

function obtained by shifting a known function $h : R^3 \rightarrow R$, i.e.,

$$b_i(x) = h(x - \mu_i), \tag{2}$$

where $\mu_i \in R^3$ is the position of the ith node of a cubic regular grid (see Fig. 2) defined in Ω. In this paper it will be assumed that h is a tri-linear interpolation function defined by:

$$h(x) = \begin{cases} \prod_{i=1}^{3}(1 - \frac{|x^i|}{\Delta}), & x \in \delta, \\ 0, & \text{otherwise}, \end{cases} \tag{3}$$

where x^i is the ith coordinate of x, Δ is the grid step and $\delta = [-\Delta, \Delta]^3$.

The grid defines a partition of Ω into cubic voxels. It can be concluded from (2) and (3) that each basis function b_i has a finite support region of eight voxels and each 3D point belongs to eight support regions. Therefore to compute $f(x_0)$ defined in (1) only eight coefficients are needed since all other basis functions are zero at $x = x_0$.

The estimation of the unknowns coefficients U, given a set of observations, V, can be formulated in a Bayesian framework. Adopting the MAP method this leads to the following optimization problem:

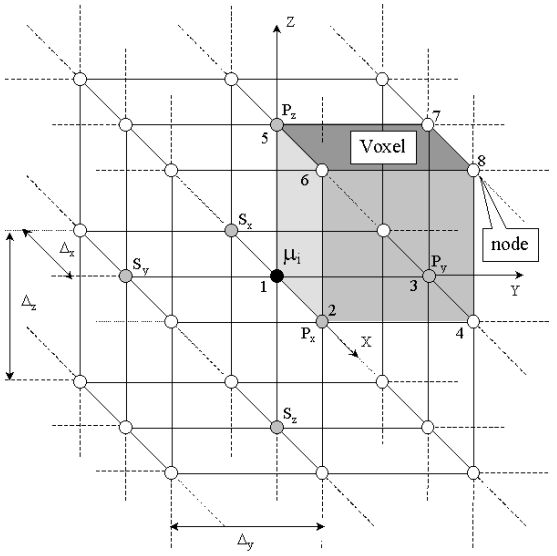

Fig. 2. 3D grid of points.

$$\hat{U} = \arg\max_{U} \ln(p(Y|U)p(U)), \tag{4}$$

where $p(V|U)$ is the sensor model and $p(Y|U)$ is the prior (Ripley, 1996). This formulation leads to a huge optimization problem that must be addressed using numeric methods.

2.1. Observation model

The observations consist of a set of 2D images obtained by intersecting the volume with plans defined by the sensor position and orientation (see Fig. 1). The volume model is to be estimated from this set of incomplete information (there are points in Ω which are not intersected by any plan). Furthermore, the intensity measurements are corrupted by multiplicative noise (speckle). This noise is produced during the image acquisition process being caused by the interaction of the acoustic wave with the tissues and the ultrasound probe surface (Abbot and Thurstone, 1979). The analysis of the physics associated with the ultrasound propagation and interaction with tissues suggests that the echo intensity measurement with the ultrasound probe has a Rayleigh distribution (Shankar, 1986; Corsini et al., 1996; Cramblitt and Parker, 1999; Keyes and Tucker, 1999; Jakeman and Pusey, 1976). The ultrasound images produced by commercial equipment are usually obtained after a nonlinear pre-processing stage which modifies the distribution of the received echo. However, the intensity variations have a non-Gaussian distribution which can be approximated by a Rayleigh distribution as before.

In this paper, pixel intensities are considered as realizations of independent random variables. Statistical independence of all elements of Y is assumed because the point spread function (PSF) of the image acquisition system is smaller than the inter-pixel distance (Dias and Silva, 1998). With this assumption the likelihood function is given by

$$p(Y|U) = \prod_i p(y_i|U). \tag{5}$$

To derive the observation model using a Rayleigh distribution for the data, we assume a large number of scatters per resolution cell with

approximately the same magnitude, no specular reflectors and a ultrasound probe which is pressure sensitive (Shankar, 1986; Abbot and Thurstone, 1979). In this case, it is assumed that the parameter of the Rayleigh distribution is equal to $f(x)$ and consequently the likelihood function is

$$p(Y|U) = \prod_i \frac{y_i}{f(x_i)} e^{-(y_i^2/2f(x_i))}. \quad (6)$$

2.2. Prior model

Three-dimensional ultrasound requires the interpolation of measured data as well as noise reduction to reduce speckle. Some assumptions must be made in order to interpolate the data between the inspection planes. In a Bayesian framework this information is included in the prior distribution Not all U vectors occur with the same probability. The available knowledge about the spatial properties of human organs and tissues can be used to define a prior distribution. In this work, a Gibbs prior is used, based on a set of quadratic potential functions. These functions are used to guarantee the smoothness of the spatial coefficients. This introduces a regularization effect which allows to recover the unknown coefficients even when there is no data in a given region. In addition, regularization also improves the convergence of the optimization algorithm. The adoption of a Gibbs distribution is equivalent to considering U as a Markov random field as stated in the Hammersley–Clifford theorem (Geman and Geman, 1984; Li, 1998). The prior distribution is given by

$$p(U) = \frac{1}{Z} e^{-\alpha \sum_{(i,j)\in\Gamma} P_i(u_j)} = \frac{1}{Z} e^{-\alpha \sum_{(i,j)\in\Gamma} (u_i - u_j)^2}, \quad (7)$$

where Γ is the set of all pairs of grid indices (i,j) such that $\|x_i - x_j\| \leqslant \Delta$ (see Fig. 3) and Z is a normalization factor. Each grid node is connected to six neighbors, except boundary nodes. The α parameter measures the strength of each connection. A high value of α correspond to strong connection between neighboring nodes (differences receive a high penalty) while low values of α correspond to weak connections. It is often convenient to assume that α varies during the optimization process, starting with a high value which is gradually reduced.

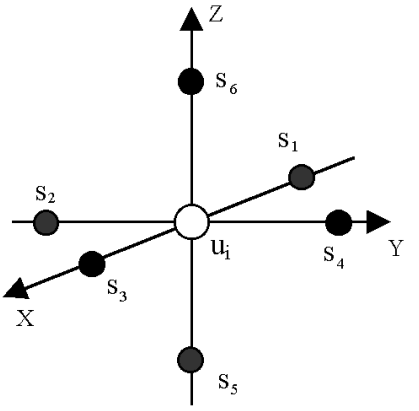

Fig. 3. Neighborhood representation.

3. Parameter estimation

This section addresses the estimation of the unknown coefficients using the Rayleigh observation model. Parameter estimation is performed using the MAP method. This leads to the maximization of the joint density $p(Y,U)$ given by

$$p(Y,U) = \prod_i \left[\frac{y_i}{f(x_i)} e^{y_i^2/2f(x_i)} \right] \frac{1}{Z} e^{-\alpha \sum_{(i,j)} P_i(u_j)}. \quad (8)$$

The maximization of $p(Y,U)$ with respect to U is a difficult problem since the number of parameters to estimate is very large (typically thousands of coefficients) and $p(Y,U)$ is a non-convex function (Li, 1998). Simpler expressions are obtained by maximizing $\ln p(Y,U)$ but the key difficulties remain the same.

Several methods are available to deal with large scale optimization problems, which fit into two broad categories (Press et al., 1994): stochastic algorithms and deterministic algorithms. Stochastic algorithms (e.g., simulated annealing) seek to obtain a global maximum but they are time consuming and require a slow annealing schedule. Deterministic algorithms do not guarantee the convergence towards a global maximum. However, they are often the most adequate choice

if a faster solution is required. The ICM algorithm, introduced by (Besag, 1986) is the one used in this paper. The ICM algorithm simplifies the optimization process by optimizing the objective function with respect to a single coefficient at a time, keeping the other coefficients constant. This is a 1D optimization problem which can be performed in a number of ways. This step is repeated for all the unknown coefficients in each iteration of the ICM algorithm.

To maximize (8) with respect to a single coefficient u_p the following stationary condition must be met,

$$\frac{\partial}{\partial u_p}\ln p(Y/U)+\frac{\partial}{\partial u_p}\ln p(U)=0 \tag{9}$$

let us compute both terms of this equation. Differentiating $\ln p(Y/U)$ with respect to u_p leads to (see (6))

$$\frac{\partial}{\partial u_p}\ln p(Y/U)=\frac{1}{2}\sum_i \frac{y_i^2-2f(x_i)}{f^2(x_i)}b_p(x_i), \tag{10}$$

where $b_p(x_i)=\partial f(x_i)/\partial u_p$ is the basis function centered in μ_p. Differentiating the prior logarithm leads to (see (7))

$$\frac{\partial}{\partial u_p}\ln p(U)=-2\alpha N_v(u_p-\bar{u}_p) \tag{11}$$

with

$$\bar{u}_p=\frac{1}{N_v}\sum_{u_g\in\Delta_{\delta_p}} u_g, \tag{12}$$

where N_v is the number of control points inside the neighborhood δ_p of the pth grid node $p(N_v=6$, see Fig. 3). Replacing (10), (11) in (9) leads to:

$$\frac{1}{4\alpha N_v}\sum_i \frac{y_i^2-2f(x_i)}{f^2(x_i)}b_p(x_i)-u_p+\bar{u}_p=0. \tag{13}$$

Eq. (13) can be rewritten as follows,

$$u_p=\frac{1}{4\alpha N_v}\sum_i \frac{y_i^2-2f(x_i)}{f^2(x_i)}b_p(x_i)+\bar{u}_p \tag{14}$$

which suggests an iterative procedure to compute $\hat{u}_p$. This method is known as fixed point algorithm. Convergence is guaranteed if the right-hand side is a contraction. In practice, a single iteration is performed to update the unknown coefficient u_p in each iteration of the ICM algorithm. Eq. (14) is recursively applied to update all the coefficients until the norm of the update vector becomes smaller than a given threshold.

The algorithm is initialized using a simple interpolation method (pixel nearest neighborhood interpolation (Nelson and Pretorius, 1997). This algorithm provides an estimate of the mean intensity $\bar{y}_p$ inside each voxel. Since the mean value of a Rayleigh random variable is $\sqrt{\pi f/2}$, an initial value for u_p can be easily obtained by

$$u_p=\frac{2\bar{y}_p^2}{\pi}. \tag{15}$$

4. Experimental results

The Rayleigh reconstruction algorithm described in the previous sections was tested with synthetic and medical data. The results were evaluated according to objective criteria and by visual inspection. A comparison with two other methods was performed. The methods used in these tests were the PNN and the DW reconstruction (Rohling et al., 1999).

Three figures of merit (FOMs) were used to evaluate the estimated volume:

- signal to noise ratio

$$\mathrm{SNR}=10\log_{10}\frac{\int|f_0(x)|\,\mathrm{d}x}{\int|f(x)-f_0(x)|\,\mathrm{d}x}, \tag{16}$$

- mean absolute error

$$\mathrm{MAE}=\frac{1}{N}\sum_{i=1}^{N}|y_i-f(x_i)|, \tag{17}$$

- log-likelihood function

$$\mathrm{LL}=\frac{1}{N}\sum_{i=1}^{N}\log\left(\frac{y_i}{f(x_i)}\right)-\frac{y_i^2}{2f(x_i)}. \tag{18}$$

In these expressions sums are performed for all the data points used for testing. N denotes the number of test points and f_0 denotes the true function we wish to estimate. These FOM are computed twice for each experiment. First we compute the FOMs using all the data (integrals are

evaluated in the whole volume Ω). This provides an overall evaluation of the reconstructed volume. Then we compute the FOMs using the leave one out method as suggested in (Rohling et al., 1999) i.e., we reconstruct the volume using all the observed planes except one and compute the FOMs using the data of the image which was not used for reconstruction. The process is repeated for all other planes and average values are computed at the end. The leave one out method allows to assess the interpolation capability of the algorithms. Please note that the SNR cannot be computed in experiments with medical data since the true f_0 is unknown then.

Fig. 4 shows the results of a test performed with synthetic data. The function to reconstruct is binary: it takes a value 0.75 in the interval $[-0.5, 0.5]^3$ and 0.25 outside. A grid of $65 \times 65 \times 65$ nodes was used to approximate this function f in the interval $[-1, 1]^3$. The measurements used in this example consist of 65 images obtained by cutting f along non-parallel planes and corrupting the cross-section images with Rayleigh multiplicative noise, according to (6). Figs. 4(a) and (b) show two cross-sections of f_0. Figs. 4(c)–(f) show observed (noisy) and reconstructed images obtained with the Rayleigh algorithm. The last two images correspond to unobserved planes. Fig. 4(i)–(j) gives an overall view of the reconstructed function visualized by ray casting. The details of ray casting algorithm can be found in (Foley et al., 2000).

A comparison with the PNN and DW interpolation results can be found in Figs. 5 and 6 and in Table 1. Fig. 5 shows reconstructed images obtained by the three methods. The proposed algorithm (Rayleigh model) manages to eliminate the noisy appearance in the reconstructed images keeping an acceptable edge representation. A more detailed view can by obtained in Fig. 6 where reconstructed profiles obtained by the three methods are shown. Again, the noise reduction performed by the Rayleigh algorithm is clear.

Table 1 shows that the proposed algorithm achieves better FOMs than the two other methods, except in the case of the log-likelihood function. This drawback is explained by the regularization effect introduced by the prior. In fact, this difficulty would disappear if the a posteriori density (which includes the prior) was used instead. A comprarision with the RBF algorithm was not performed due the complexity of this method. However, the results presented in (Rohling et al., 1999) sugest that the RBF method achieves a comparable performance.

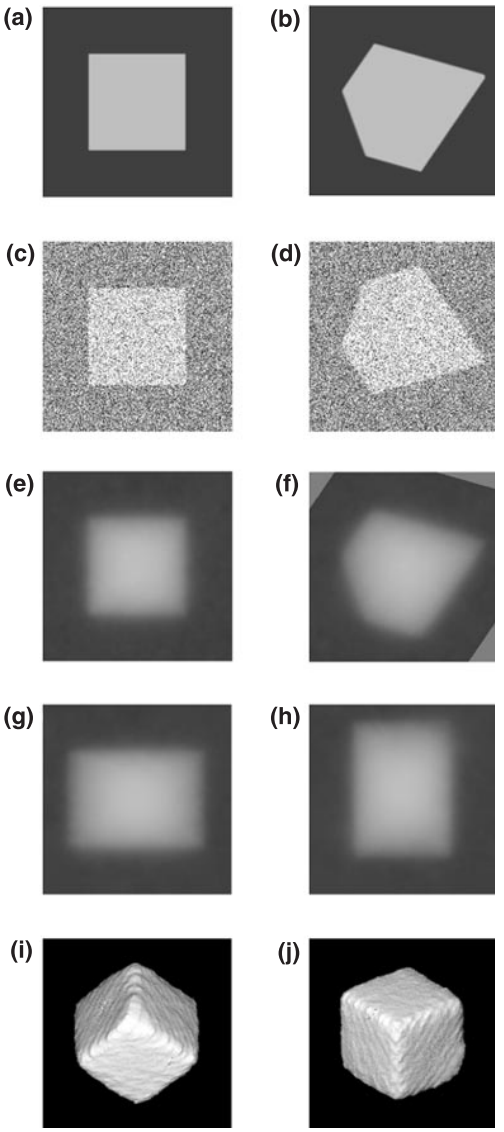

Fig. 4. (a) and (b) Ideal images without noise, (c) and (d) observed images with Rayleigh noise, (e) and (f) reconstructed images, (g) and (h) new cross-sections of the reconstructed volume, and (i) and (j) 3D reconstruction of the cube.

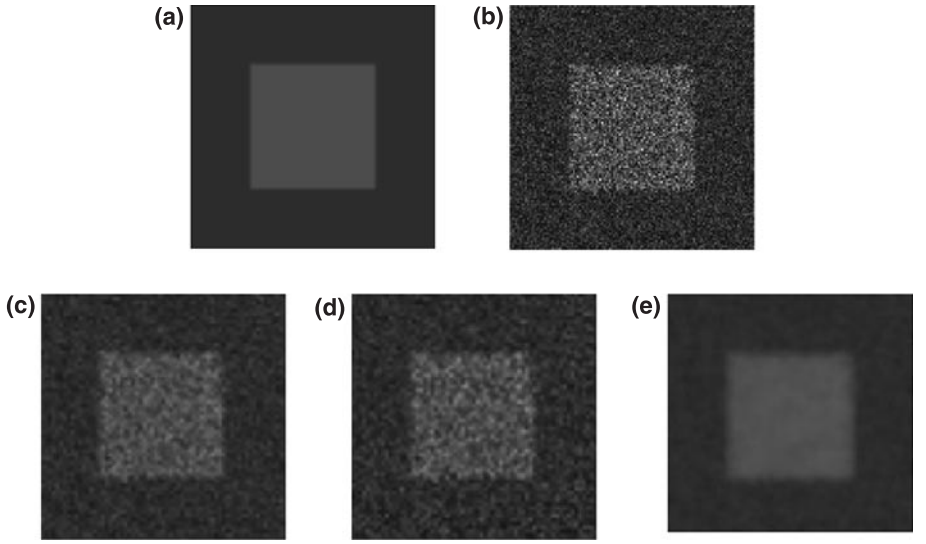

Fig. 5. Reconstructed results: (a) ideal image, (b) observed image (Rayleigh noise); reconstructed images with (c) PNN, (d) DW, and (e) Rayleigh reconstruction.

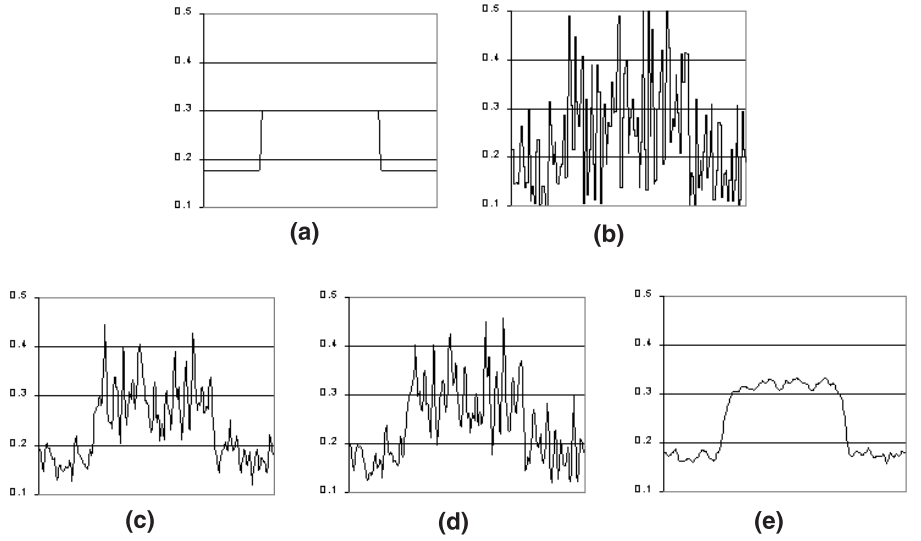

Fig. 6. Reconstructed results: (a) ideal profile; (b) measured data (Rayleigh noise); reconstructed profiles with (c) PNN, (d) DW, and (e) Rayleigh reconstruction.

Similar experiments were performed with a set of 62 ultrasound images of a gall bladder provided by the University of Cambridge in the scope of SOLUS 3D project. A grid of $129 \times 129 \times 129$ nodes was used to describe a volume of $10 \times 10 \times 10$ cm^3. The grid dimensions correspond to a trade off between computational time and accuracy. It was experimentally found that no significant improvement was achieved in these experiments by adopting a denser grid. The results are shown in Figs. 7 and 8 and in Table 2.

Table 1
Test with synthetic data (the best results are in bold)

	Global			Leave one out		
	SNR	MAE	LL	SNR	MAE	LL
PNN	9.0	59.4	−179	19.0	61.9	**−186**
DW	16.8	58.9	**−173**	16.8	62.9	−188
Rayleigh	**25.4**	**58.5**	−184	**25.4**	**61.0**	−190

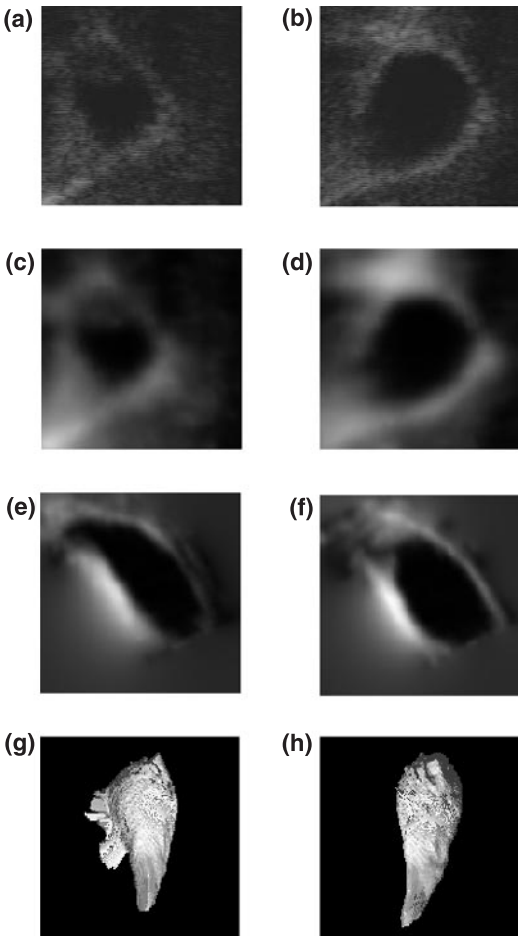

Fig. 7. (a) and (b) Ultrasound images of a gall bladder, (c) and (d) reconstructed images, (e) and (f) new cross-sections of the reconstructed volume, (g) and (h) 3D reconstruction of the Gall Bladder.

The reconstruction results presented in Fig. 7 show a close similarity with the observed images, except for the blurring effect which occurs at the transitions. Noise reduction is observed in the reconstructed images stressing the ability of the MAP algorithm in speckle elimination. The 3D reconstruction of the organ surface obtained by thresholding and ray casting is shown in Fig. 7(g) and (h).

The comparison with the PNN and DW interpolation methods is shown in Fig. 8. [1] A small visual improvement is achieved by the Rayleigh method since the reconstructed images obtained with this method are less noisy. This improvement is corroborated by the objective measures displayed in Table 2. As it can be observed all FOMs computed for the Rayleigh method are better than the ones obtained for the other two algorithms (the SNR was not computed since there is no ground truth in this experiment).

5. Conclusion

A 3D reconstruction algorithm was described based on a Rayleigh data distribution. This algorithm provides a principled approach for the interpolation and noise reduction of ultrasound images which was validated by experimental tests.

Several research directions can be explored in the future. No attempt was made in this paper to optimize the velocity of the reconstruction procedure. This issue can be addressed by using multi-scale methods (Herman and Kuba, 1999) and faster optimization algorithms. We note that the principles adopted in this paper can be easily extended in a multi-scale framework since the vector space used to define the interpolated volume

[1] The PNN and DW algorithms were implemented considering a neigborhood of 1 voxel.

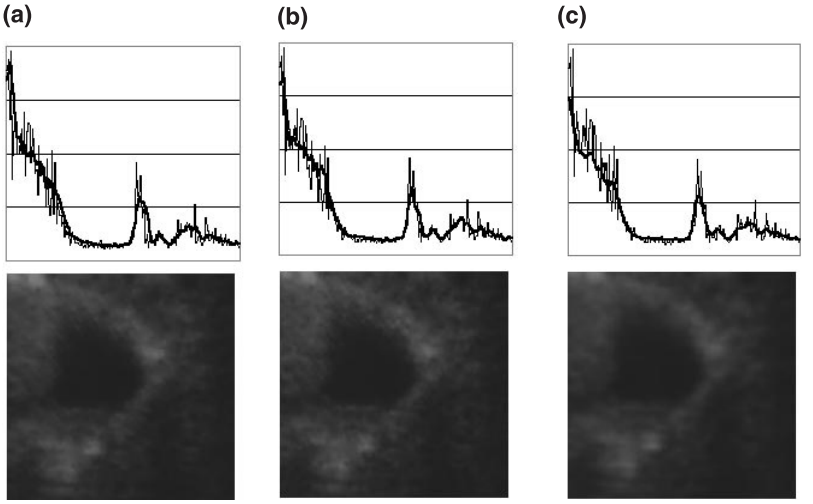

Fig. 8. Reconstructed images and profiles with: (a) PNN; (b) DW, and (c) Rayleigh reconstruction.

Table 2
Test with the medical data (the best results are in bold)

	Global		Leave one out	
	MAE	LL	MAE	LL
PNN	69.9	−195	66.7	−198
DW	69.8	−194	66.4	−197
Rayleigh	**66.4**	**−177**	**64.0**	**−178**

contains a set of nested lower resolution spaces associated to subsampled grids. Other multi-resolution spaces (e.g., wavelets (Mallat, 1989)) can also be studied.

The registration of the ultrasound images is also a key direction for future work. This is an active research area in 3D ultrasound (e.g., see Dorai et al., 1998; Rohling et al., 1998) as well as in other visualization modalities (Levoy, 1990; Nelson et al., 1999). In the case of free hand ultrasound imaging, the sensor position and orientation data is corrupted by measurement noise and the organs shapes change during the acquisition process. Both effects produce image misalignments which degrade the performance of the reconstruction algorithm. It is expected that the same principles used in this paper for volume estimation may also be applied for the estimation of alignment parameters (Sanches and Marques, 2000).

A third direction concerns the data distribution. The Rayleigh model adopted in this paper to represent speckle noise improves the results obtained under the Gaussian assumptions. However, it is important to validated this model and to determine if significant improvements can still be achieved by using more accurate data distributions.

Acknowledgements

The ultrasound images were provided by R. Prager and J. Carr from the University of Cambridge in the scope of SOLUS 3D project. We thank the helpful discussions with Andrew Gee from the same university. We also thank the anonymous reviewers for suggesting the comparison of the proposed algorithm with other interpolation methods.

References

Abbot, J., Thurstone, F., 1979. Acoustic speckle: theory and experimental analysis. Ultrasound Imaging 1, 303–324.

Besag, J., 1986. On the statistical analysis of dirty pictures. J. Roy. Statist. Soc., Ser. B 48 (3), 259–302.

Burckhardt, C., 1978. Speckle in ultrasound B-mode scans. IEEE Trans. Sonics Ultrsonics SU-25 (1), 1–6.

Chen, L.S., Herman, G.T., Reynolds, R.A., Udupa, J.K., 1985. Surface shading in the cuberille environement. IEEE Computer Graphics Appl., 33–43.

Corsini, G., Mossa, A., Verrazzani, L., 1996. Signal-to-noise ratio and autocorrelation function of the image intensity in coherent systems: sub-Rayleigh and super-Rayleigh conditions. IEEE Trans. Image Process. 5 (1), 132–141.

Cramblitt, R.M., Parker, K.J., 1999. Generation of non-Rayleigh spekle distribution using marked regularity models. IEEE Trans. Ultrasound, Ferroelectrics and Frequency Control 46 (4), 867–874.

Delingette, H., 1998. Initialization of deformable models from 3D data. In: Proc. ICCV'98 January. Bombay, India, pp. 311–316.

Dias, J., Silva, T., 1998. Restoration of coherent images using a compound markov random field. In: Proc. 10th Portuguese Conf. Pattern Recognition, IST, Lisbon, pp. 109–114.

Dorai, C., Wang, G., Jain, A.K., Mercer, C., 1998. Registration and integration of multiple object views for 3D model construction. IEEE Trans. PAMI 20 (1), 83–89.

Foley, J.D. et al., Computer Graphics : Principles and Practice, Addison-Wesley, Reading, MA.

Geman, S., Geman, D., 1984. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. IEEE Trans. Pattern Anal. Machine Intell. PAMI-6 (6), 721–741.

Gooley, T.A., Barret, H.H., 1992. Evaluation of statistical methods of image reconstruction trough ROC analysis. IEEE Trans. Med. Imaging 11 (2), 276–283.

Herman, G.T., Kuba, A., 1999. Discrete Tomography, Foundations, Algorithms, and Applications. Birkhauser, Basel.

Jakeman, E., Pusey, P.N., 1976. A model for non-rayleigh see echo. IEEE Trans. Antennas Propagation AP-24 (6), 806–814.

Keyes, T.K., Tucker, W.T., 1999. The K-distribution for modeling the envelope amplitude of a backscattered signal. IEEE Trans Ultrasound Ferroelectrics Frequency Control 46 (4), 883–887.

Levoy, M., 1990. A hybrid ray tracer rendering polygon and volume data. IEEE Comput. Graphics Appl. March, 33–40.

Li, S.Z., 1998. Close-form solution and parameter selection for convex minimization-based edge-preserving smoothing. IEEE Trans. PAMI PAMI-20 (9), 916–932.

Mallat, S.G., 1989. A theory for multiresolution signal decomposition the wavelet representation. IEEE Trans Pattern Anal. Machine Intell. 11 (7), 1989.

McInerney, T., Terzopoulos, D., 1996. Deformable models in medical image analysis: a survey. Med. Image Anal. 1 (2), 91–108.

Nelson, T., Downey, D., Pretorius, D., Fenster, A.. Three-Dimensional Ultrasound, Lippincott.

Nelson, T.R., Pretorius, D.H., 1997. Interactive acquisition, analysis and visualization of sonographic volume data. Int. J. Imaging Systems Technol. 8, 26–37.

Press, W.H., et al., 1994. Numerical Recipes in C. Cambridge University Press, Cambridge.

Quistgaard, J., 1997. Signal acquisition and processing in medical diagnostics ultrasound. IEEE Signal Process. Mag. 14 (1), 67–74.

Raya, S.P., Udupa, J.K., 1990. Shape-based interpolation of multidimensional objects. IEEE Trans. Med. Imaging 9 (1), 3–42.

Ripley, B., 1996. Pattern Recognition and Neural Networks. Cambridge University Press, Cambridge.

Rohling, R.N., Gee, A.H., Berman, L., 1998. Automatic registration of 3D ultrasound images. Ultrasound Med. Biol. 24 (6), 841–854.

Rohling, R.N., Gee, A.H., Berman, L., 1999. A comparation of freehand three-dimensional ultrasound reconstruction techniques. Med. Image Anal. 4 (4), 339–359.

Sanches, J., Marques, J.S., 2000. Alignment-by-reconstruction for 3D ultrasound imaging, ICPR 2000, Barelona.

Shankar, P., 1986. Speckle redution in ultrasound B-scans using weighted averaging in spatial compounding. IEEE Trans. Ultrasonics, Ferroelectrics, Frequency Control UFFC-33 (6), November.

Tagare, H.D., 1999. Shape-based nonrigid correspondence with application to heart motion analysis. IEEE Trans. Med. Imaging 18 (7), 570–579.

Treece, G.M., Preager, R.W., Gee, A.H., Berman, L., 1999. Fast surface and volume estimation from non-parallel cross-sections, for frehand 3D ultrasound. Med. Image Anal. 3 (2), 141–173.