Ultrasound in Med. & Biol., Vol. 28, No. 8, pp. 1029–1040, 2002

0301-5629/02/$–see front matter

PII: S0301-5629(02)00548-3

● *Original Contribution*

A MULTISCALE ALGORITHM FOR THREE-DIMENSIONAL FREE-HAND ULTRASOUND

JOÃO M. SANCHES and JORGE S. MARQUES
Instituto Superior Técnico/Instituto de Sistemas e Robótica, Lisbon, Portugal

(Received 10 *October* 2001; *in final form* 8 *May* 2002)

Abstract—This paper presents a multiscale algorithm for the reconstruction of human anatomy from a set of ultrasound (US) images. Reconstruction is formulated in a Bayesian framework as an optimization problem with a large number of unknown variables. Human tissues are represented by the interpolation of coefficients associated to the nodes of a 3-D cubic grid. The convergence of the Bayesian method is usually slow and initialization dependent. In this paper, a multiscale approach is proposed to increase the convergence rate of the iterative process of volume estimation. A coarse estimate of the volume is first obtained using a cubic grid with a small number of nodes initialized with a constant value computed from the observed data. The volume estimate is then recursively improved by refining the grid step. Experimental results are provided to show that multiscale method achieves faster convergence rates compared with a single-scale approach. This is the key improvement toward real-time implementations. Experimental results of 3-D reconstruction of human anatomy are presented to assess the performance of the algorithm and comparisons with the single-scale method are presented. (E-mail: jmrs@alfa.ist.utl.pt)

Key Words: **Medical imaging, 3-D ultrasound, Interpolation, Multiscale.**

INTRODUCTION

Three-dimensional (3-D) free-hand ultrasound (US) aims to reconstruct the geometry and acoustic properties of human organs from a set of US images obtained during a clinical session. These images are associated with nonparallel planes with known position and orientation (Quistgaard 1997; Rohling and Gee 1996) (see Fig. 1). The reconstruction algorithm must be able to interpolate the data in regions that are not intersected by any inspection plane, and must also be able to reduce the multiplicative noise associated with the observed images.

Several algorithms have been proposed for 3-D ultrasound (3DUS) imagin (Rohling et al. 1999a). These algorithms usually perform volume reconstruction in two steps (Nelson et al. 1999; Carr 1996; Steen and Olstad 1994), using a voxel representation of the region-of-interest (ROI) (Chen et al. 1985). In the first step, voxels are filled with the data obtained from the inspection planes. In the second step, an average value is computed in the intersected voxels and an interpolation method is adopted to fill empty voxels.

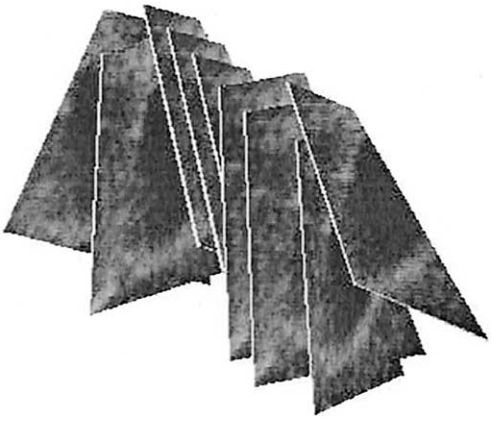

Fig. 1. US cross-sections.

3-D reconstruction of medical objects has also been addressed in a Bayesian framework by Sanches and Marques (2000b). Reconstruction is formulated as an huge optimization problem where the unknown variables are to be estimated based on a set of incomplete and noisy data (US images). The images were obtained using a free-hand commercial US equipment with a spatial locator attached to the US probe (Polhemus 1993). In this

Address correspondence to: J. M. Sanches, Instituto Superior Técnico, Av Rovisco Pais, 1049-001 Lisbon, Portugal. E-mail: jmrs@alfa.ist.utl.pt

way, the positions and orientations of the cross-sections are stored together with the images and used during the estimation process. This approach has a sound theoretical basis and compares well with other reconstruction techniques, as show by Sanches and Marques (2000b). However, the Bayesian method exhibits a slow convergence rate.

This paper tries to overcome this difficulty by using a multiscale strategy in the volume estimation.

First, a rough estimate of the volume is obtained with a small number of parameters. The volume is approximated by an interpolation of the intensity values associated with a coarse cubic grid. Then, the grid step is recursively reduced, leading to more accurate estimates of the volume of interest.

This approach has two benefits. First, it significantly reduces the computational time. The first iterations estimate a small number of coefficients (few dozens), being, therefore, much faster. Only the last iterations optimize the objective function with respect to a large number of coefficients (millions). Therefore, significant computational gain is achieved by using the multiscale approach.

Second, the use of low-resolution descriptions in the first iterations allows a faster propagation of information along the 3-D lattice, which helps to overcome the problem of missing data in specific 3-D regions.

SINGLE-SCALE ALGORITHM

This section summarizes the MAP algorithm described by Sanches and Marques (2000b) used for a single-scale (SS) description of the volume of interest.

The comparison of the method with others (Nelson et al. 1999; Carr 1996; Treece et al. 1999; Nelson and Pretorius 1997; Rohling et al. 1999b) is performed by Sanches and Marques (2000b). It is concluded that the MAP method performs better, but it is much slower.

The multiscale approach presented in the next section will overcome this drawback.

Let f be a function describing the acoustic properties of the tissue in a given ROI, $\Omega \subset R^3$. It is assumed that f is a linear combination of basis functions (interpolation functions) ϕ_i: $\Omega \rightarrow R$, i.e.;

$$f(x) = \Phi(x)^T U \tag{1}$$

where $\Phi(x) = [\phi_1(x), \phi_2(x), \ldots, \phi_N(x)]^T$ is a vector of basis functions and $U = [u_1, u_2, \ldots, u_N]^T$ is a $N \times 1$ vector of coefficients that defines the volume of interest.

The basis functions, $\phi_p(x)$, are obtained by shifting a function h: $R^3 \rightarrow R$ according to:

$$\phi_p(x) = h(x - \mu_p), \tag{2}$$

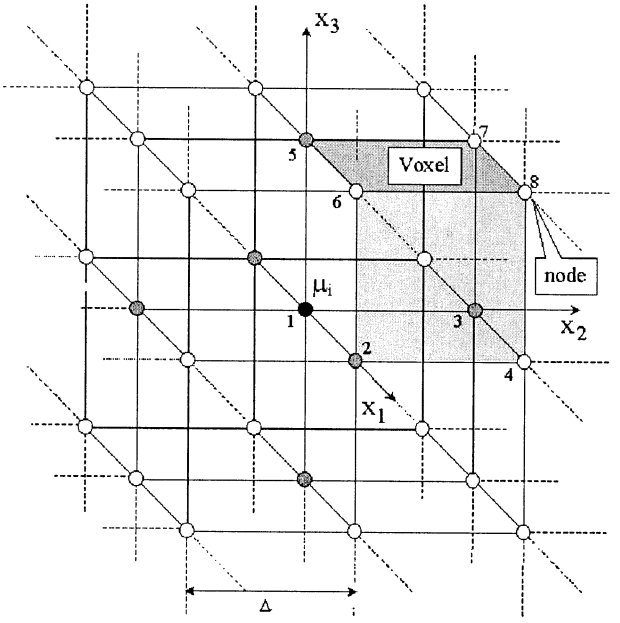

Fig. 2. 3-D grid and voxel representation.

where $\mu_p \in R^3$ are the nodes locations of a 3-D cubic grid (see Fig. 2) defined in Ω and $h(x)$ is a trilinear interpolation function defined by:

$$h(x) = \begin{cases} \prod_{k=1}^{3} \left(1 - \dfrac{|x^k|}{\Delta}\right) & x \in \delta \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

where x^k is the k th coordinate of x, $\delta = [-\Delta, \Delta]^3$ and Δ is the grid step (see Fig. 3).

A spatial locator is attached to the US probe (see Fig. 4). This accurately measures the position and orientation of the inspection plane, provided that a careful calibration procedure is performed (Prager et al. 1998).

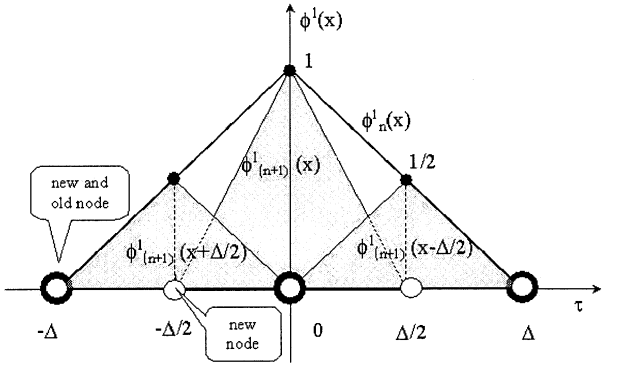

Fig. 3. 1-D basis function.

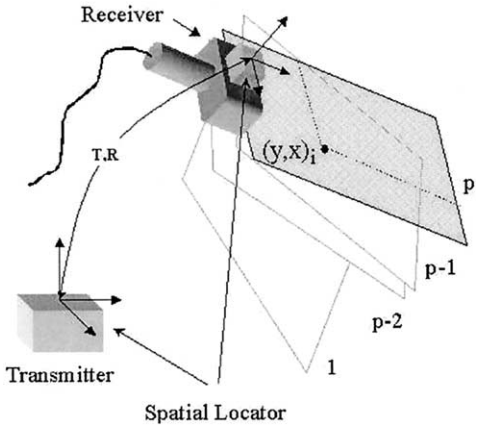

Fig. 4. Measurement of plane position and orientation.

This makes it possible to compute the 3-D positions of the pixels associated with all the observed slices.

Therefore, the available data are

$$V = \{y_i, x_i,\}$$

where y_i is the intensity of the ith pixel and x_i is its corresponding 3-D position. The 3-D coordinates x_i are assumed accurately known.[1] Only the image intensities y_i are considered to be random variables. The 3DUS problem aims at the estimation of coefficients U, given V. This can be addressed using a MAP method, as follows (Sanches and Marques 2000b):

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

where $p(V|U)$ is the sensor model and $p(U)$ is the prior density.

Suitable probability distribution must be defined for $p(V|U)$ and $p(U)$. This topic will be addressed below.

In this paper, pixel intensities are considered to be independent random variables with a Rayleigh distribution leading to the likelihood function;

$$p(V|U) = \prod_i \frac{y_i}{f(x_i)}\, e^{-\frac{y_i^2}{2f(x_i)}}, \tag{5}$$

where f is the function to be retrieved.

[1]The registering problem is not treated in this paper. This problem was published by Sanches and Marques (2000a, 2002).

These assumptions are not always true, but they provide an acceptable model for the observed data that allow us to derive analytic expressions for the estimates. Several authors consider the pixel intensities to be statistically independent (Dias and Leitão 1996), although the point spread function is sometimes greater than the interpixel distance. The Rayleigh model is one of the models used in US imaging (Burckhardt 1978; Abbot and Thurstone 1979; Wells and Halliwell 1981), and it is assumed to be appropriate for the examples treated in this paper.

Intensity errors in US images are due to constructive/destructive interference phenomenon appearing in the US images as a kind of multiplicative noise (the US wave is a coherent radiation producing effects similar to those obtained in laser systems) (Abbot and Thurstone 1979; Achim et al. 2001). The Rayleigh model arises if the number of scatters per resolution cell is large, the echo complex magnitude components, in phase and quadrature, are normally distributed and the complex phase is uniformly distributed. When the number of scatters is small or some of them are stronger than the others (which happens in strong specular reflections associated with the organ boundaries), the Rayleigh model is no longer valid. In these cases, other distributions should be used to describe the observed data (*e.g.*, the K-distribution) (Jakeman and Pusey 1976). Furthermore, the US equipment usually performs a preprocessing of the raw data to reduce the dynamic range of the RF signal for visualization purposes. This operation modifies the data distribution. However, it is possible to estimate the preprocessing function from the observed data and to compensate for the preprocess compression, obtaining an estimation of the original raw data (Sanches and Marques 2001).

In this paper, it is assumed that the original data are described by the Rayleigh model and that the observed data are decompressed.

The prior $P(U)$ plays an important role in the reconstruction process because it introduces a interpolation effect that makes it possible to recover the function coefficients even when there are no data in its vicinity. This is the basic mechanism that makes it possible to fill the volume gaps that were not intersected by any cross-section. Furthermore, the prior $P(U)$ also helps to avoid unstable behavior during the optimization process (Katsaggelos 1991).

In this paper, a Gaussian prior is used (Geman and Geman 1984):

$$p(U) = \frac{1}{Z}\, e^{-\alpha \sum_{g \in G} \sum_{i \in \delta_g} (u_g - u_i)^2} \tag{6}$$

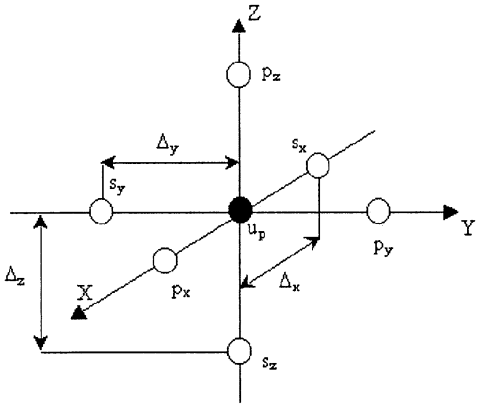

Fig. 5. Neighborhood representation.

where G denotes the grid nodes, δ_g are the neighbors of the gth node (see Fig. 5), and Z is a normalization factor. This prior models the correlation among intensity coefficients on the volume of interest, Ω. In the Bayesian context, this prior takes into account the *a priori* knowledge about the volume to estimate. The adoption of this prior is equivalent to considering f to be bandwidth-limited (*i.e.*, neighboring nodes should have similar values).

This model depends on a parameter α that accounts for intensity changes between neighboring nodes. Each grid node is connected to six neighbors, except boundary nodes. The α parameter controls the strength of each connection. High values of α correspond to imposing strong connections between neighboring nodes, and low values of α correspond to imposing weak connections. Therefore, choosing large values of α leads to smooth estimates, and small values of α lead to noisy estimates of the volume, with sharper transitions. The choice of α is a trade-off between noise reduction and the ability to cope with intensity transitions. Furthermore, a high value of α makes it possible rapidly to fill the gaps (*i.e.*, regions that were not intersected by any cross-section). In this paper, the α parameter was obtained by trial and error and is sometimes modified during the estimation process, starting with high values being gradually reduced.

The MAP reconstruction is the output of an optimization problem:

$$\hat{U} = arg \max_{U} \mathrm{L(U)}, \tag{7}$$

where $L(U) = \ln(P(V/U)P(U))$ is the objective function to maximize. Using eqns (5) and (6), we obtain:

$$L(U) = \Sigma_i \left[\ln\left(\frac{y_i}{f(x_i)}\right) - \frac{y_i^2}{2f(x_i)} \right] - \alpha \sum_{g \varepsilon G} \sum_{i \varepsilon \delta_g} (u_g - u_i)^2. \tag{8}$$

The optimization of eqn (8) with respect to U is a difficult problem because the number of parameters to estimate is very large (typically millions of coefficients). Furthermore, $L(U)$ is a nonconvex and nonlinear function, for which there is no close form solution (Li 1998). Therefore, numerical methods should be considered.

To solve eqn (7), the ICM algorithm proposed by Besag (1986) is used, where the joint optimization problem is converted into a sequence of 1-D optimization procedures. In each iteration, this method considers the objective function $L(U)$ as being a 1-D function depending on a single parameter, keeping all the others constant. During the iterative process, all the parameters to estimate are updated sequentially until convergence is achieved.

Let $L(u_p)$ denote the objective function as a function of u_p coefficient, keeping all the others constant. To maximize $L(u_p)$, the following stationary condition must be met:

$$\frac{\partial L(u_p)}{\partial u_p} = 0. \tag{9}$$

After straightforward manipulation, this leads to:

$$\frac{1}{2} \sum_i \frac{y_i^2 - 2f(x_i)}{f^2(x_i)} \phi_p(x_i) + 2\alpha N_v (u_p - \bar{u}_p) = 0, \tag{10}$$

where N_v is the number of neighbors of u_p, $\phi_p(x_i)$ is the basis function associated to the pth node computed at x_i and $\bar{u}_p = \frac{1}{N_v} \sum_{i \epsilon \delta_p} u_i$ is the mean intensity associated to the pth node neighbors.

This equation can be numerically solved by using the fixed point method (Press et al. 1994) leading to:

$$^{n+1}\hat{u}_p = \frac{1}{4\alpha N_v} \sum_i \frac{y_i^2 - 2f(x_i)}{f^2(x_i)} \phi_p(x_i) + \bar{u}_p, \tag{11}$$

where $f(x)$ and $\bar{u}_p$ are computed by using the estimated values computed in the previous iteration.

MULTISCALE ALGORITHM

As stated previously, the Gibbs prior has a stabilization effect in the convergence process of the algorithm, having, at the same time, a regularization effect in the

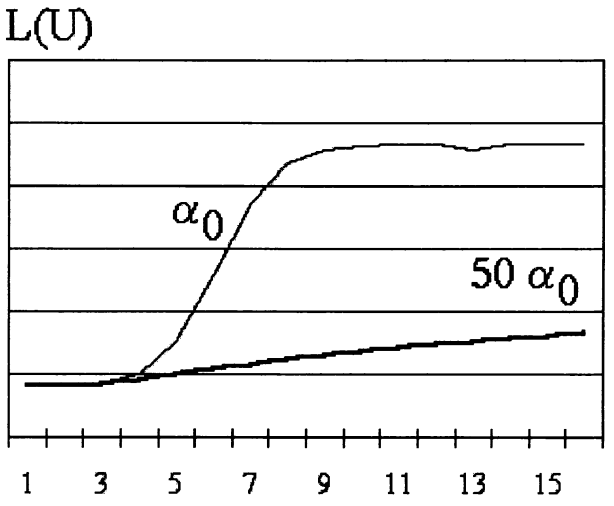

Fig. 6. Evolution of $L(U)$ along the iterative process for $\alpha = 10^{-12}$ and $\alpha = 5.10^{-11}$ using synthetic data.

final solution. This last effect can be controlled through the α parameter, which plays a key role in enforcing the convergence of the algorithm, as well as in reducing the multiplicative noise present in US images. However, the prior leads to a decrease of the convergence rate of the optimization algorithm, due to the propagation of information along the lattice nodes. This effect increases with the increase of α because an increased dependence is enforced among neighboring nodes. Figure 6 shows the evolution of the objective function during the iterative process using synthetic data, for two different values of α, α_0 and $50\alpha_0$. As shown, convergence rate for $50\alpha_0$ is clearly smaller than that for $\alpha = 10^{-12}$.

To overcome this difficulty (Herman and Kuba 1999), a multiscale (MS) approach is proposed to speed up the convergence of the sequence defined in eqn (11).

The idea is simple. First, we start with a low resolution grid (*e.g.*, eight nodes) and compute the node intensities for this grid. Then, we increase the number of nodes during the optimization process until the final resolution is achieved. The initial volume for each resolution is computed from the final estimate obtained in the grid with lower resolution.

It will be shown that this strategy improves the convergence rate of the algorithm.

To implement this strategy, it must be guaranteed that the objective function does not change during the grid refinement. This is achieved if the following conditions are met in every scale change (see Appendix B):

1. The resolution doubles from one grid resolution to the next.
2. The α parameter is proportional to the step of the grid.

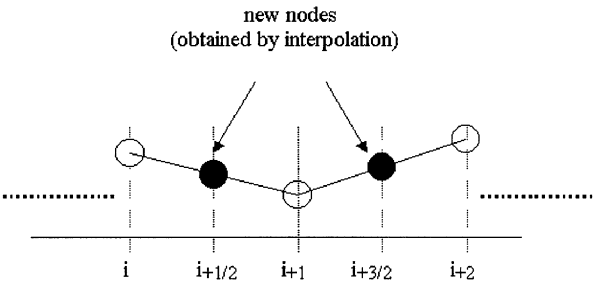

Fig. 7. Interpolation method to propagate the volume estimated in resolution grid n to the next resolution grid.

The first condition guarantees that the vectorial spaces are nested (*i.e.*, the nth vectorial space is a subset of the $(n + 1)$th space), as is shown in Appendix A. In this case, the ratio of discretization steps in two consecutive scales is:

$$r = \frac{\Delta^n}{\Delta^{n+1}} = 2, \tag{12}$$

where Δ^n and Δ^{n+1} are the steps associated to grids n and $n + 1$, respectively. With this strategy of grid refinement, the number of nodes per coordinate increases from two like 2,3,5,9,17,33,65,129, . . . , N, $2N - 1$, ... The initialization of the volume in the grid $n + 1$ is easily computed from the estimated volume in grid n by keeping unchanged the nodes that already belong to nth grid and interpolating between them to compute the value of the new ones, as shown in Fig. 7.

To met the second condition, let us consider the simplest case in which grid step is the same for the 3-D and equal to Δ^n for the highest and final resolution grid, and Δ^i for the grid step of one of the intermediate resolution grids. Let us use, for each resolution grid i, different prior parameters defined as:

$$\alpha_i = \frac{\Delta^i}{\Delta^n}\alpha, \qquad i = 1, \ldots, n. \tag{13}$$

Therefore, the MS strategy is implemented as follows:

1. The algorithm starts with a low resolution grid (*e.g.*, $2 \times 2 \times 2$).
2. The ith grid is obtained from the previous one by doubling its resolution, see eqn (12). The nodes of the new grid are preinitialized with the values of $f(x)$ computed in the previous grid, see eqn (1), at the locations of the new grid nodes (see Fig. 7) (*i.e.*, $(u_p)_0^i = f_{i-1}(\mu_p^i)$, where μ_p^i is the 3-D location of the u_p^i node.

3. For each grid, the new estimation of U is obtained using eqn (11), where the α parameter is given by eqn (13).

In this way, the objective function is invariant under grid refinements as required. Only one iteration is performed per grid scale during the first n iterations. The scale is then kept constant until a stop condition is met. The α parameter used in eqn (13) is defined by the user and it is obtained by trial and error. It was concluded, however, that the reconstruction results are not strongly influenced by the choice of α when the multiscale approach is used.

In the first iteration, each grid coefficient is initialized according to:

$$u_p^0 = \frac{2\bar{y}^2}{\pi}, \tag{14}$$

where $\bar{y}$ is the mean value of the observed data in the volume of interest Ω.

This estimate is derived from the expression of the expected value of a set of variables with Rayleigh distribution:

$$E(x) = \sqrt{\frac{\pi f}{2}}. \tag{15}$$

The MS approach also simplifies the initialization procedure because only 8 ($2 \times 2 \times 2$) nodes must be initialized. Furthermore, this initialization is always obtained from the observed data, see eqn (14), because, in the first resolution level (coarser grid), all nodes are intersected. On the contrary, in the SS method, the initialization procedure is more complex because it is necessary to adopt a strategy to fill the gaps corresponding to the nodes that were not intersected by any cross-sections.

EXPERIMENTAL RESULTS

The multiscale and single-scale algorithm are evaluated based on three figures of merit: the *a posteriori* probability $L(U)$, see eqn (7), the signal-to-noise ratio (SNR) of the reconstruction results, and the number of iterations needed to achieve convergence.

Three examples are presented. The first two examples use synthetic data (*i.e.*, cross-sections extracted from a single cube or from a set of cubes). The third example uses a set of medical data corresponding to cross-sections of a gall bladder.

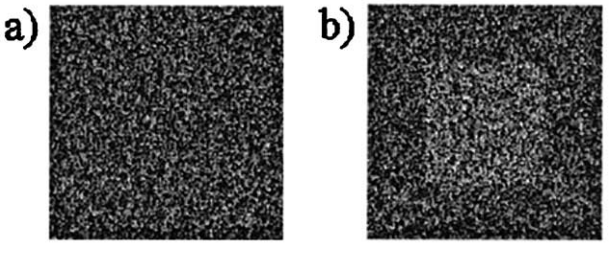

Fig. 8. Synthetic image with Rayleigh noise. (a) Background noise; (b) cube intersection.

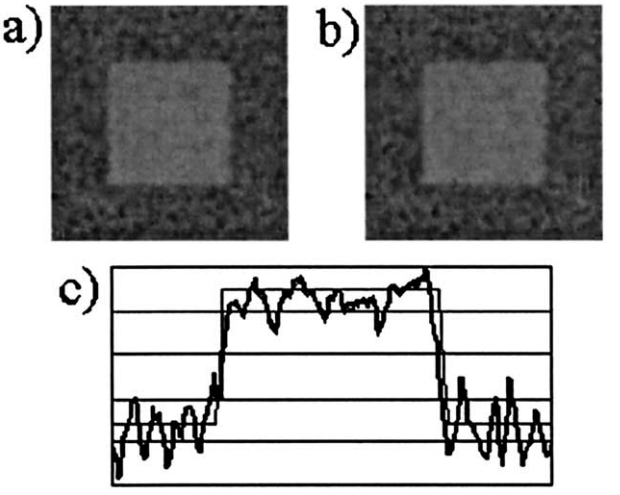

Fig. 9. Reconstruction results with $\alpha = 10^{-12}$. (a), (b) Cross-sections extracted from the estimated volume obtained with SS algorithm and MS algorithm, respectively. (c) Intensity profiles along the diagonal for the SS algorithm (thin line) and MS algorithm (thick line).

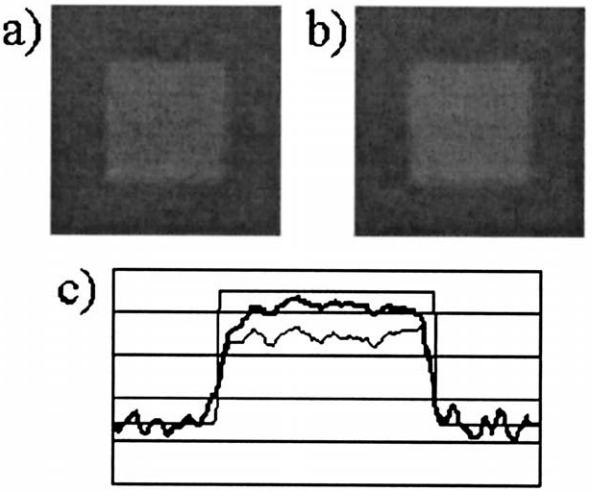

Fig. 10. Reconstruction results with $\alpha = 5.10^{-12}$. (a), (b) Cross-sections extracted from the estimated volume obtained with SS algorithm and MS algorithm, respectively. (c) Intensity profiles along the diagonal for the SS algorithm (thin line) and MS algorithm (thick line).

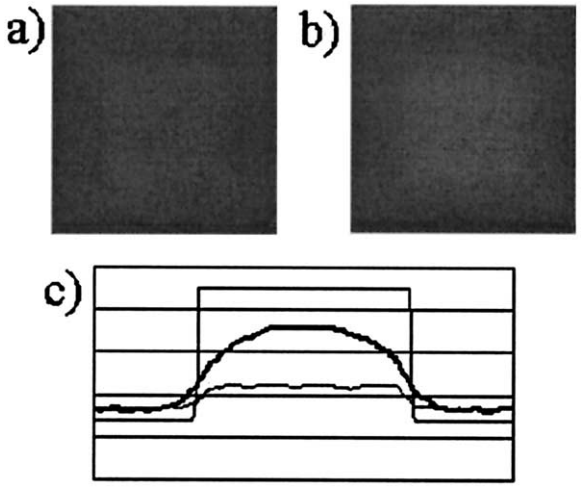

Fig. 11. Reconstruction results with $\alpha = 10^{-11}$. (a), (b) Cross-sections extracted from the estimated volume obtained with SS algorithm and MS algorithm, respectively. (c) Intensity profiles along the diagonal for the SS algorithm (thin line) and MS algorithm (thick line).

In all these experiments, a grid with $65 \times 65 \times 65$ nodes is used to reconstruct the data in the volume of interest. This means that 274,625 coefficients have to be estimated. The reconstruction process is performed during 15 iterations for the synthetic data and during 20 iterations for the medical data. From the figures representing the evolution of the objective function along the iterative process, it is possible to conclude whether convergence was or was not achieved and how far it is from the convergence.

Example 1

The first set of tests used a sequence of 50 parallel cross-sections extracted from a synthetic cube. Images with 128×128 pixels were computed by corrupting the cross-sections with Rayleigh-distributed noise. Figure 8 shows images corresponding to two different cross-sections, one extracted from the background and one intersecting the cube.

The cube was reconstructed with the MAP algorithm for three values of α ($\alpha = 5.10^{-12}, 10^{-11}, 5.10^{-11}$) after 15 iterations.

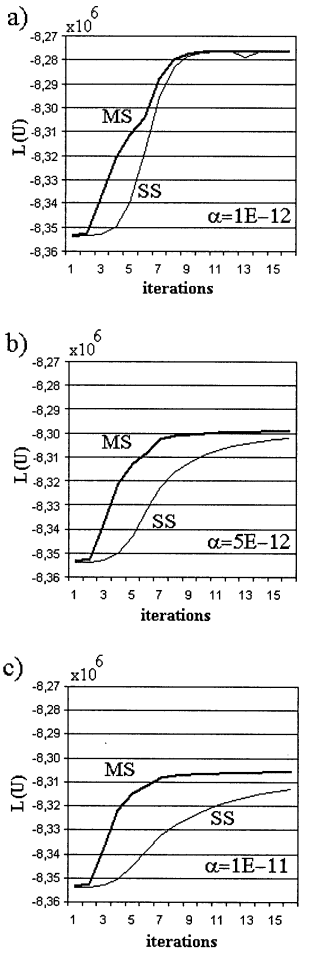

Fig. 12. Convergence of the single-scale (SS) and multiscale (MS) algorithms. (a) $\alpha = 10^{-12}$; (b) $\alpha = 5.10^{-12}$; (c) $\alpha = 10^{-11}$.

Figures 9–11 show the reconstruction results for $\alpha = 10^{-12}$, 5.10^{-12} and 10^{-11}, respectively. These figures show two cross-sections of the reconstructed volumes using the SS and MS algorithms and the intensity profiles along the main diagonal of the images.

Both methods provide similar results for small values of α ($\alpha = 10^{-12}$) but the reconstruction results become different for larger values of α.

Better results are always achieve by the MS method, which shows a smaller bias.

Table 1. Simulation results using a synthetic cube for $\alpha = 10^{-12}$, 5.10^{-12} and 10^{-11}

	$\alpha = 10^{-12}$		$\alpha = 5.10^{-12}$		$\alpha = 10^{-11}$	
	SS	MS	SS	MS	SS	MS
SNR (dB)	18.03	18.15	17.12	20.20	11.61	15.13
$L(U) \times 10^{-6}$	−8.277	−8.277	−8.302	−8.299	−8.313	−8.306

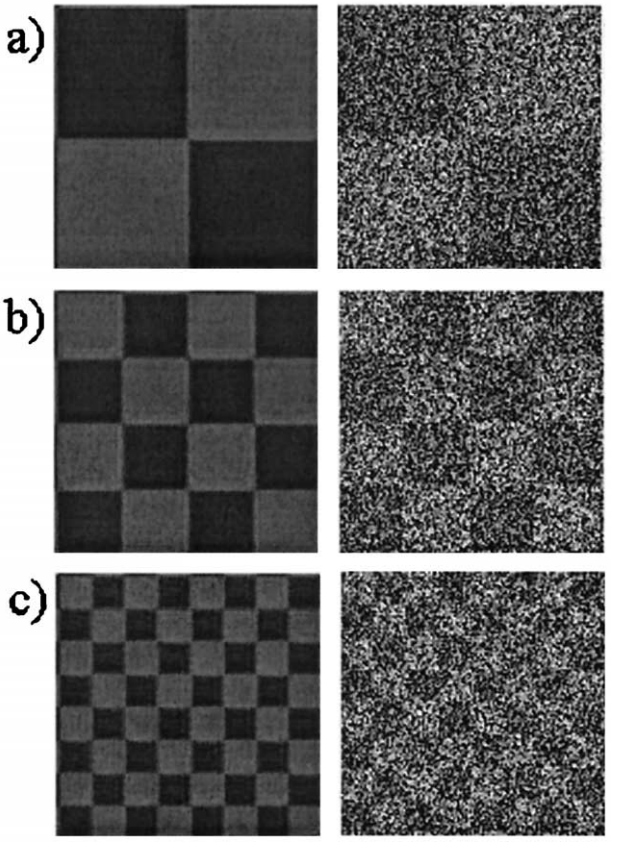

Fig. 13. Cross-sections extracted from three volumes with different number of transitions with no noise (left column) and corrupted by Rayleigh noise (right column) for the original volumes with (a) 8 cubes, (b) 32 cubes and (c) 256 cubes.

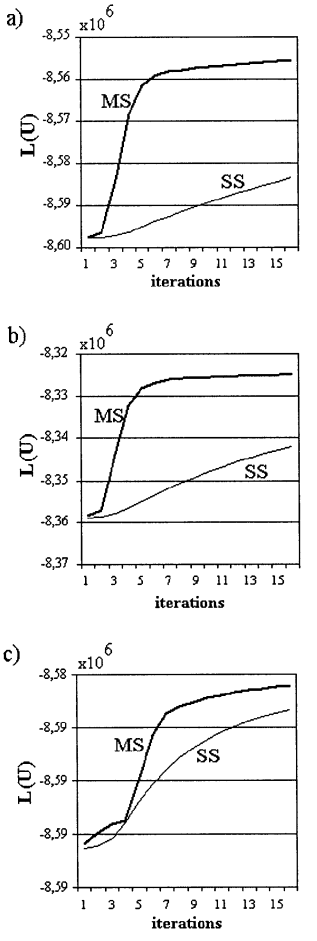

Fig. 14. Evolution of the objective function along the iterative process for the original volumes with (a) 8 cubes, (b) 64 cubes and (c) 512 cubes.

This difference can also be observed in Table 1, which shows the SNR and the $L(U)$ values for each experiment. The MS method achieves the best scores for all values of α. Figure 12 represents the evolution of the objective function along the iterative process for both methods, considering three values of α. It is visible, from these plots, that the MS approach converges faster than the SS algorithm, achieving higher values of the objective function. The differences in the convergence rates are larger when the prior is stronger. Large values of α produce a long-range smoothing, which is easily propagated using the low-resolution grid adopted in the first iterations of the MS method.

Example 2

Instead of considering an homogeneous cube as before, the volume of interest, Ω, is now filled with nonoverlapping cubes with two intensity levels. Three cases were considered, ranging from a small number of cubes (8) to a high number of cubes (512) inside Ω.

Figure 13 shows three cross-sections extracted from the three volumes (left column) as well as the corresponding images corrupted with Rayleigh noise (right column).

Each volume was then reconstructed using α =

Table 2. Simulation results for three volumes filled with homogeneous cubes

	8 cubes		64 cubes		512 cubes	
	SS	MS	SS	MS	SS	MS
SNR (dB)	10.21	12.77	10.52	11.32	10.78	11.08
$L(U) \times 10^{-6}$	−8.58	−8.56	−8.34	−8.32	−8.59	−8.58

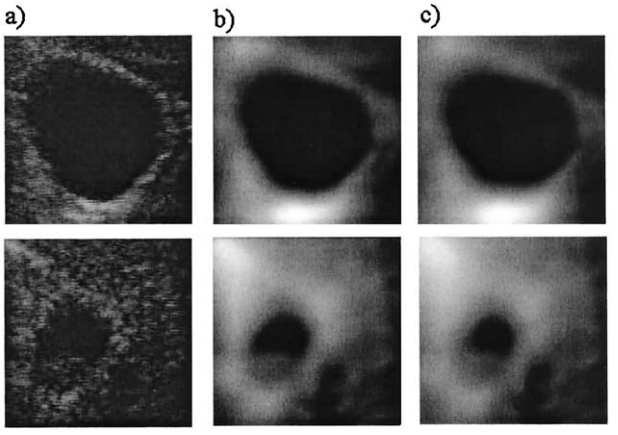

Fig. 15. (a) Two cross-sections belonging to the data set. (b),(c) the corresponding cross-sections extracted from the estimated volumes using the SS and MS approaches, respectively.

5.10^{-11} by both methods using 50 parallel cross-sections of the volume of interest with 128 × 128 pixels. Table 2 shows the SNR and L values obtained for each experiment. Better results were achieved by using the MS method, as in the previous example.

This experiment also shows that the improvement obtained with the MS approach is data-dependent. The improvement is higher when the number of transitions is small. This can be explained from the structure of the objective function L. The difference is more pronounced

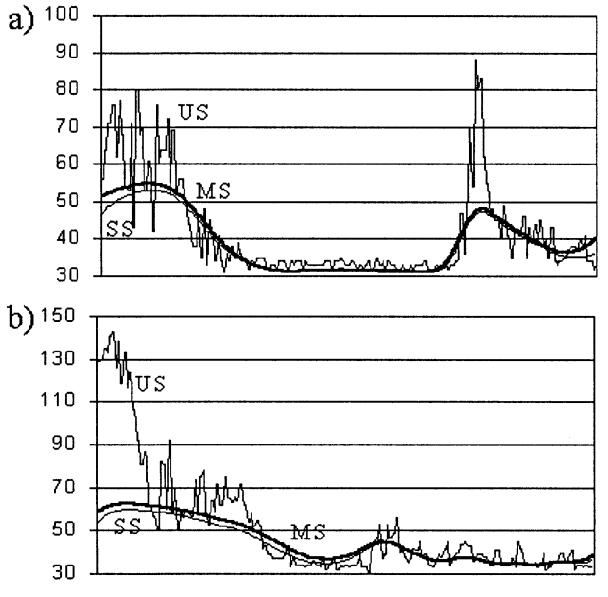

Fig. 16. Profiles extracted from the cross-sections displayed in (a) the first row and (b) second row of Fig. 15. US-ultrasound image; SS-reconstructed volume using the SS approach; MS-reconstructed volume using the MS approach.

Fig. 17. Representation of the surface of the gall bladder extracted from the estimated volume using ray-casting techniques.

(2.5 dB) when the volume exhibits large homogeneous regions that are better represented with small-resolution scales. The multiscale method makes it possible to achieve reasonable reconstruction results, in this case, after the first two iterations, using low-resolution models. This is not true when the volume of interest exhibits a large number of transitions. In this case, higher resolution models are required to approximate the function f.

Figure 14 shows the evolution of the objective function along the optimization process. Faster convergence rates are obtained by the MS approach. The convergence rate of the MS approach does not depend on the experiment, and the SS method converges slower when the volume of interest has large homogeneous regions (small number of transitions). It should also be stressed that each iteration in the MS scheme is faster than in the SS method, during the first stage of the optimization process corresponding to a computational gain.

Example 3

The experiments with real data were performed using a set of 62 images with 176 × 176 pixels corresponding to nonparallel cross-sections of a gall bladder obtained with an US probe.

Figure 15a shows two US images belonging to the data sequence and the corresponding cross-sections extracted from the estimated volumes using the SS (Fig. 15b) and MS (Fig. 15c) approaches. Both methods lead to similar reconstruction results. However, the MS approach is faster and achieves a higher value for the objective function (see Fig. 18). Figure 16 shows the main diagonal profiles extracted from the images dis-

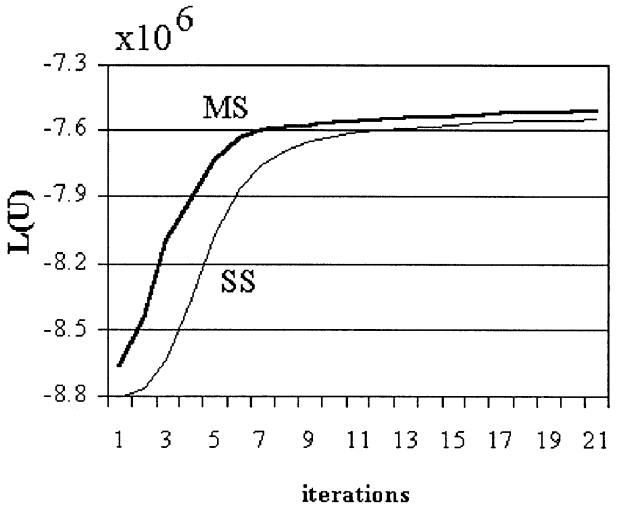

Fig. 18. Evolution of the objective function along the estimation iterative process of the gall bladder.

played in Fig. 15 and Fig. 17 shows a 3-D surface reconstruction of the gallbladder extracted from the estimated volume computed using the MS approach. This figure was obtained by applying a data segmentation algorithm and using ray-casting techniques.

The evolution of the objective function along the optimization procedure is displayed in Fig. 18 for both methods (SS and MS).

CONCLUSIONS

This paper presents a Bayesian multiscale algorithm for 3-D reconstruction of human organs from a set of US images. 3-D data are approximated using a sequence of nested spaces with increasing resolution. Each space is spanned by a set of basis functions associated to the nodes of a cubic grid. An optimization method is proposed, making possible to: 1. estimate a linear combination of the basis functions for each resolution and 2. predict the value of the coefficients in a space with higher resolution. These two steps alternate starting from a low-resolution space until the desired resolution is achieved.

In this paper, volume reconstruction is performed using a set of cross-sections of the region of interest. The position and orientation of each cross-section is measured by a spatial locator sensor and assumed without errors. This assumption is not always valid in practice. The 3-D measurements provided by the Polhemus sensor used in this work are accurate. However, the pressure of the US probe against the human tissues causes geometric deformation that cannot be ignored. This problem has been addressed by several authors (Nelson et al. 1999; Treece et al. 2001; Rohling et al. 1998) and by us in previous work (Sanches and Marques 2002). The methods described by Sanches and Marques can be easily incorporated in to the MS algorithm described here.

Experimental tests are presented to evaluate the algorithm with US data. It is concluded that the multiscale concept leads to fast convergence rates. Convergence of the estimates using the MS method is achieved in less than 10 iterations, suggesting the possibility of using Bayesian estimation methods for interactive 3-D US imaging.

The initialization procedure is also simplified using the MS method. In fact, only the eight nodes of the coarsest grid (MS) must be initialized, instead of the million nodes of the finest grid (SS). Initialization of these eight nodes with a constant value, computed from the observed data, avoids the adoption of complex strategies to initialize the millions of nodes of the SS method, where some of them are not intersected by any cross-section.

Acknowledgements—This work was partially supported by FCT under project Heart 3D (Sapiens). The ultrasound images were kindly provided by R. Prager and A. Gee from the University of Cambridge.

REFERENCES

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

Achim A, Bezerianos A, Tsakalides P. Novel Bayesian multiscale method for speckle removal in medica ultrasound images. Trans Med Imaging 2001;20(8):772–783.

Besag J. On the statistical analysis of dirty pictures. J R Statist Soc B 1986;48(3):259–302.

Burckhardt C. Speckle in ultrasound B-mode scans. IEEE Trans Son Ultrason 1978; 25(1)

Carr J. Surface reconstruction in 3D medical imaging. Ph.D. thesis. Electrical and Electronic Engineering, University of Canterbury, Christchurch, New Zealand, 1996.

Chen LS, Herman GT, Reynolds RA, Udupa JK. Surface shading in the cuberille environement. IEEE Comput Graph Applic 1985.

Dias J, Leitão J. Wall position and thickness estimation from sequences of echocardiograms images. IEEE Trans Med Imaging 1996;15: 25–38.

Geman S, Geman D. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. IEEE Trans Pattern Anal Machine Intell 1984;6.

Herman GT, Kuba A. Discrete tomography, foundations, algorithms, and applications. Birkhauser, 1999.

Jakeman E, Pusey PN. A model for non-Rayleigh see echo. IEEE Trans Antennas Propagation 1976; AP-24(6):806–814.

Katsaggelos AK. Digital image restoration, Springer series in information sciences. Springer-Verlag, 1991.

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

Nelson TR, Pretorius DH. Interactive acquisition, analysis and visualization of sonographic volume data. Int J Imaging Syst Technol 1997;8:26–37.

Nelson T, Downey D, Pretorius D, Fenster A. Three-dimensional ultrasound. Lippincott, 1999.

Polhemus, Fasttook user's manual, revision F. 1993.

Prager RW, Rohling RN, Gee AH, Berman L. Rapid calibration for 3-D freehand ultrasound. Ultrasound Med Biol 1998;24(6):855–869.

Press WH, Vetterling WT, Teukolsky SA, Flanner BP. Numerical recipes in C. Cambridge University Press, 1994.

Quistgaard J. Signal acquisition and processing in medical diagnostics ultrasound. IEEE Signal Proc Mag 1997;14(1):67–74.
Rohling R, Gee A. Issues in 3-D free-hand medical ultrasound imaging. CUED/F-INFENG/TR 246, 1996.
Rohling RN, Gee AH, Berman L. Automatic registration of 3-D ultrasound images. Ultrasound Med Biol 1998;24(6):841–854.
Rohling RN, Gee AH, Berman L. A comparation of freehand threedimensional ultrasound reconstruction techniques. Med Image Anal 1999;4(4):339–359.
Rohling R, Gee A, Berman L, Treece G. Radial basis function interpolation for 3D freehand ultrasound. In: Proceedings of the 16th International Conference on Information Processing in Medical Imaging Visegrd, Hungary, June 1999. Springer, 1999:478–483.
Sanches J, Marques J. Alignment-by-reconstruction for 3D ultrasound imaging. Proceedings of 15th ICPR, Barcelona 2000a;3:41–44.
Sanches J, Marques J. A Rayleigh reconstruction/interpolation algorithm for 3D ultrasound. Pattern Recog Lett 2000b;21:917–926.
Sanches J, Marques J. 3D reconstruction from log-compressed Rayleigh images. Proceedings of the 6 International Conference on Image Processing, ICIP2001, Thessaloniki, Greece, October 2001.
Sanches J, Marques J. Joint image registration and volume reconstruction for 3D ultrasound. Special Issue on Ultrasonic Image Processing and Analysis. Pattern Recog Lett 2002.
Steen E, Olstad B. Volume rendering of 3D medical ultrasound data using direct feature mapping. IEEE Trans Med Imaging 1994; 13(3): 517–525.
Treece GM, Preager RW, Gee AH, Berman L. Fast surface and volume estimation from non-parallel cross-sections, for freehand 3-D ultrasound. Med Image Anal 1999;3(2):141–173.
Treece GM, Prager RW, Gee AH, Berman L. Correction of probe pressure artifacts in freehand 3D ultrasound. Proceedings of Medical Image Computing and Computer-Assisted Intervention (MICCAI 2001), Utrecht, The Netherlands, October 2001. Springer, 2001:283–290.
Wells PNT, Halliwell M. Speckle in ultrasonic imaging. Ultrasonics 1981.

APPENDIX A

Refinement

In this Appendix, we will show that doubling the resolution of the 3-D grid leads to a sequence of nested vector spaces. Any function defined in a low-resolution space can be easily expressed as a linear combination of the basis function associated with a high-dimension space.

Consider a new 3-D grid formed by a set of nodes at locations $\mu_\lambda = (x_0 + i\Delta_x/2, y_0 + j\Delta_y/2, z_0 + k\Delta_z/2)$, where (x_0, y_0, z_0) is the coordi nate of the bottom left corner of the (ROI).

To prove that the new set of basis functions span a vectorial space that contains the older, we just have to show that any basis function defined in the nth grid can be expressed as a linear combination of the basis functions defined in the $(n + 1)$th grid.

Let:

$$\phi_n^1(x) = h(x/\Delta_n^1) \tag{A.1}$$

be the first component of the basis function defined in the nth grid.

Thus

$$\phi_n(u) = \phi_n^1(u^1)\phi_n^2(u^2)\phi_n^3(u^3), \tag{A.2}$$

where $u = (u^1, u^2, u^3)$ and

$$h(x) = \begin{cases} \Pi_{i=1}^3(1 - |x^i|) & x \in \delta \\ 0 & \text{otherwise} \end{cases}. \tag{A.3}$$

We shall assume, without lack of generality, that $\phi_n{}^1(x)$, represented in Fig. 3, is centered at the origin, because any other basis function can be obtained by simply shifting it.

It is easy to show that (see Fig. 3):

$$\phi_n^1(x) = \sum_{i=-1}^{1} \phi_n^1\left(i\,\frac{\Delta_n^1}{2}\right) h\left(\frac{x - i\Delta_n^1/2}{\Delta_n^1/2}\right) \tag{A.4}$$

where

$$h\left(\frac{x - i\Delta_n^1/2}{\Delta_n^1/2}\right) = \phi_{(n+1)}^1(x - i\Delta_{(n+1)}^1) \tag{A.5}$$

is the first component of the basis function defined in the $(n + 1)$th grid and $\phi_n^1(i\Delta_n^1/2)$ is the value of the first component of the nth grid basis function calculated at the nodes of the $(n + 1)$th grid. Replacing eqn (A.4) in eqn (A.2) leads to:

$$\begin{aligned}\phi_n(u) = &\sum_{i=-1}^{1} (\phi_n^1(\mu_i^1)\phi_{(n+1)}^1(u^1 - \mu_i^1)) \\ &\sum_{j=-1}^{1} (\phi_n^2(\mu_j^3)\phi_{(n+1)}^2(u^2 - \mu_j^2)) \\ &\sum_{k=-1}^{1} (\phi_n^3(\mu_k^3)\phi_{(n+1)}^3(u^3 - \mu_k^3)),\end{aligned} \tag{A.6}$$

where $\mu_g = (\mu_i^1, \mu_j^2, \mu_k^3)$ is the position of the gth node in the $(n + 1)$th grid. After rearranging the terms, the basis function defined in the nth grid can be written as follows:

$$\phi_n(u) = \sum_{g=(i,j,k)} \phi_n(\mu_g)\phi_{(n+1)}(u - \mu_g) \tag{A.7}$$

Equation (A.7) shows that a basis function defined in the nth grid can be obtained as a linear combination of the basis functions associated to the $(n + 1)$ th grid, where the coefficients are the values of the n th grid basis function calculated in the positions of the nodes of the $(n+1)$ th grid. This means that the space defined in the n th grid is nested in the space defined in the $(n+1)$ th grid.

APPENDIX B

Invariance of the objective function under grid refinement

Let us consider a general Gibbs distribution:

$$P(f) = \frac{1}{Z}\, e^{-\alpha E(f)}, \tag{B.1}$$

with energy defined by:

$$E = \int_{R\subset\Omega} |\nabla f(x)|^2 d\upsilon, \tag{B.2}$$

where $f(x) = \sum_p u_p b_p(x)$.

Let us discretize this integral using a 3-D grid with step $(\Delta^1, \Delta^2, \Delta^3.)$ The approximated value is:

$$E \approx \sum_{p\in G} |\nabla f_d(x_p)|^2 \Delta^1\Delta^2\Delta^3, \tag{B.3}$$

where G is the set of all index of the grid and x_p is the 3-D position of the p th node.

The gradient can be approximated by first order backward differences,

$$\overline{\nabla} f_d(x_p) = \begin{pmatrix} (u_p - s_1)/\Delta^1 \\ (u_p - s_2)/\Delta^2 \\ (u_p - s_3)/\Delta^3 \end{pmatrix}, \tag{B.4}$$

where s_1, s_2, s_3 are the neighbors of u_p (see Fig. 5). Substituting eqn (B.4) in eqn (B.3) leads to:

$$E = \sum_p C_x(u_p - s_1)^2 + C_y(u_p - s_2)^2 + C_z(u_p - s_3)^2 \tag{B.5}$$

where $C_1 = \dfrac{\Delta_2\Delta_3}{\Delta_1}$, $C_2 = \dfrac{\Delta_1\Delta_3}{\Delta_2}$ and $C_3 = \dfrac{\Delta_1\Delta_2}{\Delta_3}$.

When $\Delta_1 = \Delta_2 = \Delta_3 = \Delta$, eqn (B.5) takes the next form:

$$E = \Delta \sum_p \sum_i (u_p - s_i)^2 \tag{B.6}$$

and substituting in eqn (B.1) leads to:

$$P(U) = \frac{1}{Z} e^{-\alpha\Delta} \sum_p \sum_i (u_p - s_i)^2. \tag{B.7}$$

Therefore, instead of using eqn (6) to define the prior to perform the MAP estimation, we should use the eqn (B.7), which is a better approximation of eqn (B.5), which is independent of the discretization grid.

Furthermore, this approximation depends on the error by approximating $f(x)$ by $f^n{}_d(x)$ and the gradient $\nabla f(x)$ by eqn (B.4). However, in this case,

1. Because we are dealing with functions belonging to a finite dimension vectorial space, see eqn (1) with linear derivatives with respect to the coefficients, and
2. the vectorial space associated to nth grid is nested in the vectorial space associated to the $(n + 1)$th grid (as proved in Appendix A), then

$$E_d^n(U^n) = E_d^{n+1}(U^{(n+1)}) \tag{B.8}$$

where $E_d^n(U^n)$ and $E^d_{(n=1}$ $(U^{(n+1)})$ are the approximated discrete energies computed in n th and $(n + 1)$ th grids, respectively.

Thus, our discrete energy, $E_d^n(U^n)$, that is to be maximized is kept constant under a grid refinement.