Ultrasound in Med. & Biol., Vol. 29, No. 2, pp. 255–269, 2003

0301-5629/03/$–see front matter

ELSEVIER

doi:10.1016/S0301-5629(02)00762-7

● *Original Contribution*

A THEORETICAL FRAMEWORK TO THREE-DIMENSIONAL ULTRASOUND RECONSTRUCTION FROM IRREGULARLY SAMPLED DATA

RAÚL SAN JOSÉ-ESTÉPAR,* MARCOS MARTÍN-FERNÁNDEZ,*
P. PABLO CABALLERO-MARTÍNEZ,* CARLOS ALBEROLA-LÓPEZ* and JUAN RUIZ-ALZOLA†
*Department of Teoría de la Señal y Comunicaciones e Ingeniería Telemática, University of Valladolid, Valladolid, Spain; and †University of Las Palmas de Gran Canaria, Las Palmas, Spain

(*Received* 15 *April* 2002; *in final form* 20 *November* 2002)

Abstract—Several techniques have been described in the literature in recent years for the reconstruction of a regular volume out of a series of ultrasound (US) slices with arbitrary orientations, typically scanned by means of US freehand systems. However, a systematic approach to such a problem is still missing. This paper focuses on proposing a theoretical framework for the 3-D US volume reconstruction problem. We introduce a statistical method for the construction and trimming of the sampling grid where the reconstruction will be carried out. The results using *in vivo* US data demonstrate that the computed reconstruction grid that encloses the region-of-interest (ROI) is smaller than those obtained from other reconstruction methods in those cases where the scanning trajectory deviates from a pure straight line. In addition, an adaptive Gaussian interpolation technique is studied and compared with well-known interpolation methods that have been applied to the reconstruction problem in the past. We find that the proposed method numerically outperforms former proposals in several control studies; subjective visual results also support this conclusion and highlight some potential deficiencies of methods previously proposed. (E-mail: caralb@tel.uva.es)

Key Words: **3-D ultrasound imaging, Irregularly sampled ultrasound data, Reconstruction grid extraction, PCA, Adaptive Gaussian interpolation.**

INTRODUCTION

Three-dimensional (3-D) ultrasonic imaging is becoming a widespread practice in clinical environments due to the potential of applications based on 3-D representation. Basically, the major drawback that physicians have to cope with when traditional single 2-D B-scans are used the need of mentally reconstructing the 3-D anatomy can be naturally overcome using 3-D ultrasound (US) imaging and appropriate further processing. But we may also enumerate other interesting benefits of 3-D echography: the spatial relationships among 2-D slices are preserved in the 3-D volume, allowing an off-line examination of scans previously recorded even by another clinician; slices that cannot be acquired because of the geometrical constraints imposed by other structures of the patient can now be readily rendered by the so-called any-plane slicing; and volume visualization and accurate volume estimation may greatly enhance the diagnosis task.

Several 3-D US techniques have been reported in the literature in the past few years. These techniques can be coarsely classified as those derived from a 2-D phased-array probe, and those that obtain a 3-D data set from 2-D B-scans acquired in rapid succession while the probe is in motion. The former technique employs a bidimensional array of piezoelectric elements and the volume is scanned by electronically steering the array elements. The main drawback of this promising technique is the limited field of view of the existing probes. The latter technique makes use of conventional 2-D US systems and a positioning system. This technique includes both the freehand and the mechanically-swept volume acquisition techniques. Freehand has received increasing attention, especially since the second half of the 1990s (Nelson and Pretorius 1997; Detmer et al. 1994), probably due to its inherent flexibility and low cost compared with the

Address correspondence to: Carlos Alberola-López, ETSI Telecomunicación, Campus Miguel Delibes, 47011 Valladolid, Spain. E-mail: caralb@tel.uva.es

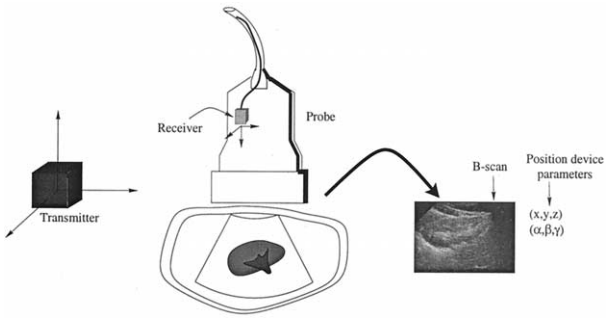

Fig. 1. Schematic description of a freehand system. The position parameters define the transformation between the receiver coordinate system and the transmitter coordinate system.

3-D US probes, and it can be now considered as a well-trusted technique.

In freehand imaging, a 3-D positioning sensor provides a measure of the position and orientation of the coordinate system of a receiver (typically attached to the probe) with respect to a fixed coordinate system located at the transmitter. Figure 1 sketches the elements involved in a freehand system. Each B-scan pixel, which is initially measured in the coordinate system of the probe, can be spatially located with respect to the fixed coordinate system by means of a vector of position and a vector of orientation that jointly comprise the six degrees of freedom between two 3-D coordinate systems. A simple affine transformation allows us to convert each position in the image plane to a position in 3-D space. A critical step is the calibration of the freehand system. Calibration deals with the problem of finding the transformation between the image plane and the receiver coordinate system. The accuracy of the overall system critically depends on this calibration stage, which, if successfully carried out, provides an accurate registration of US images to 3-D space. Interesting papers have addressed this issue in detail (Detmer et al. 1994; Prager et al. 1998).

The acquired data can be seen as a scattered distribution of planes in 3-D space due to the total freedom of the physician to perform the scanning procedure. Reconstruction of a volume out of a number of planes with arbitrary orientations (hereafter referred to as irregularly sampled data) to end up with data in a regular grid is not an obvious processing task. Although a real-time visualization based on raw B-scans (without any prior reconstruction) is possible and allows online diagnosis (Prager et al. 1999) and, despite the existence of efforts of authors (Treece et al. 1999, 2000, 2001) to do signal processing directly on the raw data, a regular grid is still needed to apply further processing with off-the-shelf algorithms.

Consequently, this problem of volume reconstruction of US has been tackled by several authors (Barry et al. 1997; Rohling et al. 1997; Sanches and Marques 2000) in the past. All these studies focus on incremental interpolation of the gaps between the slices, as well as on the determination of the final intensity of a voxel when several B-scans overlap on this voxel (this latter procedure is known as compounding). Specifically, Rohling et al. (1997) focus on spatial compounding with image-based registration. The regular grid is filled using a nearest neighbor (NN) approach and the solution they propose for compounding is an average operation. Their most outstanding contribution is the registration of B-scans related to a baseline that is built from a quick pass over the region-of-interest (ROI). The registration step is intended to avoid blurring of misregistered structures due to spatial compounding. Barry et al. (1997) follow an inverse distance weighting (DW) scheme; each pixel that falls within a circular area of a fixed radius (centered at the voxel) contributes to the voxel value inversely to the distance from the voxel to the pixel. Rohling et al. (1999) provide a survey of well-known reconstruction methods like voxel nearest neighbor (VNN), pixel nearest neighbor (PNN) and DW interpolation; they also introduce a new interpolation method based on radial basis function (RBF). Finally, Meairs et al. (2000) apply a weighted ellipsoid Gaussian convolution kernel to tackle the reconstruction of irregularly-sampled data.

All the above-mentioned studies, despite their unobjectionable quality, share the same pitfall: no attention is paid to the details of the construction of the regular grid (cuberille). However, we have found that this problem is of paramount importance for further data processing and storage. It is obvious that, before applying any interpolation/compounding technique, the volume where the data are going to be resampled must be defined. The definition of the volume grid implies the selection of an orientation and extent for the volume and the selection of a voxel size. With respect to the former, on one side clinicians (ultrasonographers) often prefer that the reconstruction volume resembles as much as possible either the original sweep or some anatomical predefined plane. But, on the other side, it is needless to say that one of the most outstanding difficulties of US imaging is that a standard scanning policy does not exist; clinicians look for the best viewing direction (also called insonation angles) at will for each patient, attempting to avoid annoying effects like shadowing. Therefore, the overall scan may consist of several subscans from which a single scanning direction (say, principal direction in the examination) may not be clearly defined.

The resulting situation can be made clearer with an example: assume the physician has made a scan, consisting of two subscans, the trajectories of which draw the letter "x". Assuming that two subscans suffice, then the intersection of the segments of the "x" will comprise the

ROI. However, no slice in any of these two subscans gives a representative scanning direction. Hence, how to create a regular volume in this example is not straightforward, due to the lack of a clear scanning direction, unless an expert does provides additional information. Moreover, for the example, any reconstructed regular volume that contains all the pixels from all the B-scans will probably be very large and it will contain a lot of information-lacking voxels (*i.e.,* voxels with an intensity value imposed by the reconstruction process, and typicaly equal to zero). In this situation, we understand that, if the direct destination of the data set is a computer, to, for instance, perform a posterior registration with, perhaps, other imaging modalities, the need for a compact regular volume is high, so that common software can be used to that end and the volume can be efficiently dealt with by avoiding processing on a large number of irrelevant voxels. If, on the contrary, the direct destination of the volume is a physician, that person will have to patiently move through a large amount of slices with only a few relevant pixels in two areas of the slices (those that result from the intersection of the two tails of the "x" with the slices), until the ROI is reached. That makes the process of data analysis very cumbersome.

This paper aims at providing a theoretical framework to the problem of US volume reconstruction from freehand systems (*i.e.,* at facing the problem of the construction of the regular grid that best fits the data). This problem comprises three subproblems, namely, 1. the selection of the coordinate system of the reconstruction grid, 2. the selection of the extent of the reconstructed volume, and 3. the determination of the voxel size.

Principal component analysis (PCA) has been used in this study to deal with the first of the aforementioned subproblems. Considering the 3-D positions of the pixels as samples of a population, we look for the coordinate system that achieves the largest data variance in each direction while being uncorrelated with the others. The size of the reconstructed volume is pruned using the eigenvalues information provided by PCA. Furthermore, an adaptive Gaussian convolution kernel for dealing with irregularly-sampled data is introduced and compared with some well-known reconstruction techniques. We will demonstrate the application of this reconstruction method to *in vivo* data.

STATE OF THE ART

Notation and terminology

The position of a B-scan pixel is computed as a set of homogeneous coordinate system transformations. Following the notation provided by Prager et al. (1998), the overall transformation can be written as:

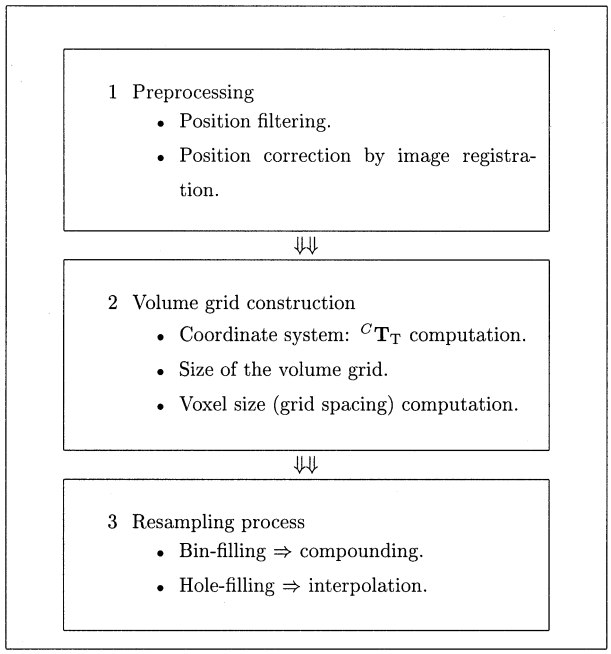

Fig. 2. Reconstruction process.

$$ {}^{C}\boldsymbol{x} = {}^{C}T_{T}\,{}^{T}T_{R}\,{}^{R}T_{P}\,{}^{P}\boldsymbol{x}. \tag{1} $$

The position of every pixel in every B-scan, ${}^{P}\boldsymbol{x}$, is transformed to a 3-D position into the C coordinate system, ${}^{C}\boldsymbol{x}$. P is the coordinate system attached to every B-scan; R and T are the coordinate systems of the position sensor receiver and transmitter respectively and, finally, C is the coordinate system of the reconstructed volume. The transformation between P and R stems from the calibration process. The second transformation between R and T is given by the position sensor's records. The last transformation is a custom transformation that allows us to choose the best representation of the reconstructed volume. The sub- and superscripts P, R, T and C will be used in each variable/transformation to denote the coordinate system that the variable/transformation is referred to.

Reconstruction: a survey

Figure 2 sketches the main steps of the reconstruction. A critical issue of every reconstruction technique is the accuracy of the position measurements. This accuracy would ideally depend only on the finite resolution of the position sensor but, in reality, electromagnetic contamination in the environment where the measurements are taken may distort the transmitter's magnetic fields (Leotta et al. 1997). The main drawback of electromagnetic position devices is the susceptibility to interference from metals and electronic devices. These circumstances

can be alleviated with a proper setup of the inspection room; however, it is almost impossible to completely get rid of them. Assuming that the probe position varies slowly between consecutive slices, smoothing the measures given by the position sensor contributes to the suppression of high-frequency artefacts and, consequently, the system accuracy can be enhanced.

The second step in the reconstruction process is the selection of a coordinate system for the reconstructed volume (*i.e.*, the definition of ${}^{C}T_{T}$) as well as the extent of the volume and the grid spacing (voxel size). Although every reconstruction process has to deal with this issue somehow, to our knowledge, it has not been explicitly treated in the literature. Barry et al. (1997) propose the use of a Key US frame to define the axes of the reconstruction volume. This Key frame is chosen by a user and, typically, turns out to be one that is centrally located and depicts a complete cross-section of the tissue of interest. The volume axes are parallel to those of the Key US frame and the origin of the volume is at the center of the Key frame. Clearly this approach works correctly when the B-scans are quasiparallel; however, under other circumstances, for example, when the volume consists of several sweeps in different directions, the selection of a representative Key frame may be difficult to do, or such a representative frame may even not exist; in addition, it will usually not be the best solution in terms of minimizing the amount of data needed to represent the scanned volume. As far as the grid spacing is concerned, several authors pose the problem as a trade-off between resolution and size of the reconstructed volume (Barry et al. 1997; Rohling et al. 1998; Leotta and Martin 2000) and the specific value is chosen *a priori*.

The third step deals with the reconstruction itself (*i.e.*, resampling the information into a regular grid). We have to cope with a well-known problem of scattered data interpolation. A good survey about scattered data interpolation was provided by Franke (1982). In our domain, the reconstruction cannot be merely seen as an interpolation problem; the process can be divided into two stages, as suggested by Rohling et al. (1999) a bin-filling stage and a hole-filling stage. The former is a compounding operation that combines the information of planes that intersect each voxel. The latter is an interpolation operation because a value has to be inferred for the voxels that have not been filled during the former step. Although the interpolation can be carried out by several methods (Franke 1982), the use of global methods[1] is not feasible, due to the high dimensionality that the interpolation procedure would suffer. Hence, it is necessary to pose simple solutions to minimize the time and memory requirements. Some well-known approaches are VNN, PNN and DW (Rohling et al. 1999).

Now, we elaborate on each step shown in Fig. 2.

[1] By global method we mean that the interpolation is dependent on all data points. They are theoretically the optimum.

METHODS

Hardware and software

For probe tracking, an electromagnetic sensor (miniBird, Ascension Technologies, Burlington, VT) was used. This position sensor provides position (x, y, z) and orientation (α, β, γ) information between the receiver and the transmitter.

US imaging was performed with a Hitachi EUB-515. Images were digitized in RGB PAL format at 25 frames/s. Both the position sensor and US system video output were attached to a Silicon Graphics O2 workstation. The US slices and position sensor measures were matched using Stradx freehand 3-D US system (Cambridge University 3-D Ultrasound Research Group, Cambridge, UK; Prager et al. 1999). The calibration of the system was performed using a single-wall phantom as described by Prager et al. (1998). After the images and position were recorded, the volume reconstruction was performed with custom software on a 800-MHz processor Pentium III. The custom software was implemented as a VTK (visualization toolkit) (Schroeder et al. 1998) class. VTK libraries were also used for data visualization.

Preprocessing

In this study, we have performed a non-causal FIR filtering of the position and orientation data obtained from the positioning system (after been synchronized by Stradx). This simple procedure has drawn results at an acceptable quality. The patient, however, was asked to minimize motion and to refrain from breathing for the short period of the scanning so as to minimize motion artefacts. This paper does not focus on this issue; therefore, a simple procedure has been adopted. More sophisticated procedures to assure the accuracy of the position sensor measures (mainly based on registration ideas) (Rohling et al. 1997; Rohling et al. 1998) have been described in the literature.

Volume coordinate system

The selection of the coordinate system of the reconstructed volume is one of the most critical steps in the overall process of volume reconstruction. Under this denomination, we include the following design issues: the volume axes, the volume axes origin, the volume extent and the size of each voxel. The elements involved in this stage are depicted in Fig. 3.

The optimum coordinate system often depends on

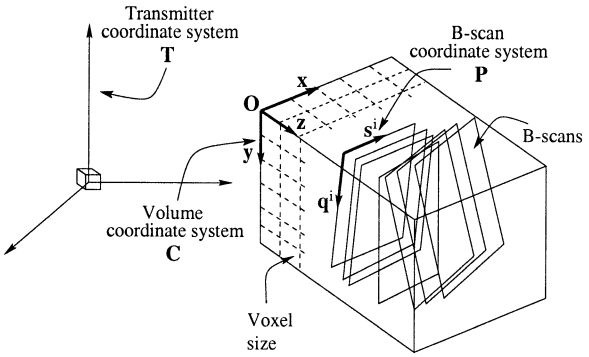

Fig. 3. Volume construction.

the application domain of the reconstructed volume; however, two general requirements can be highlighted: 1. the volume grid should enclose as large an information density as possible 2. the volume grid should have as homogeneous an information density as possible throughout the whole volume.

The first requirement accounts for the need for having a small volume that depicts all the tissue information presented in the original B-scans. The second requirement attempts to reduce the number of interpolated voxels; it would be risky to try to reconstruct areas of a volume where there is a great lack of information from the original B-scans, because such a reconstruction may lead a clinician to confusion if a diagnosis was made on a large area of interpolated data. Therefore, having a formal method that deals with these concepts would help to choose the main features of a volume reconstruction method that best fits the ROI while retaining only the relevant data from the original B-scans.

Volume axes. Principal component analysis (PCA) is a statistical tool that is used to find the linear combination with the largest variance of a data population. Typically, this tool has been used to reduce the number of variables to be treated by discarding the linear combinations with smaller variances and retaining only those with larger variances. In our case, the data population is the 3-D position of every B-scan pixel. As is well-known, if the purpose is to find the best way to get a compressed volume, the Karhunen–Loeve transform, or, in its discrete equivalent formulation, the Hotelling transform or PCA, is the optimum solution (Jain 1989).

Formally speaking, let us suppose that the 3-D position of each pixel in the transmitter coordinate system, ${}^T\boldsymbol{x}$, is a random vector ${}^T\boldsymbol{X}$ of three components with covariance matrix $\boldsymbol{\Sigma}$. The actual distribution of ${}^T\boldsymbol{X}$ is irrelevant except for the covariance matrix; however, if ${}^T\boldsymbol{X}$ is supposed to be normally distributed, more meaning can be given to the principal components. The key idea of PCA is to determine the orthogonal linear transformation, $\boldsymbol{\Phi}$, that transforms the original 3-D position into a new space:

$$ {}^C\boldsymbol{X}_{\mathrm{PCA}} = \boldsymbol{\Phi}^T\boldsymbol{X} \tag{2} $$

so that the covariance matrix of ${}^C\boldsymbol{X}_{\mathrm{PCA}}$ is a diagonal matrix, the components of which are the eigenvalues of the data covariance matrix $\boldsymbol{\Sigma}$. The columns of the transformation $\boldsymbol{\Phi}$ turn out to be the eigenvectors of $\boldsymbol{\Sigma}$. The straight meaning of the aforementioned transformation is that the components of the transformed position, ${}^C\boldsymbol{X}_{\mathrm{PCA}}$, are the ones that have maximum variance being mutually uncorrelated. The fact of being uncorrelated will allow us to trim the volume in one dimension with minimum impact on the others. A complete formulation of the mathematical background can be found in Anderson (1984).

The covariance matrix $\hat{\boldsymbol{\Sigma}}$ of our position data has been estimated from the sample positions, ${}^T\boldsymbol{x}_p$, for each pixel in the transmitter coordinate system:

$$ \hat{\boldsymbol{\Sigma}} = \frac{1}{N-1}\sum_{p=1}^{N} ({}^T\boldsymbol{x}_p - {}^T\bar{\boldsymbol{x}})({}^T\boldsymbol{x}_p - {}^T\bar{\boldsymbol{x}})' \tag{3} $$

where ${}^T\bar{\boldsymbol{x}}$ is the sample mean vector.

$$ {}^T\bar{\boldsymbol{x}} = \frac{1}{N}\sum_{p=1}^{N} {}^T\boldsymbol{x}_p \tag{4} $$

p indexes the set of pixels that contribute to the PCA analysis.

Volume axes origin. The transformation $\boldsymbol{\Phi}$ only defines a rotation to the transmitter coordinate system. As a matter of fact, the columns of matrix $\boldsymbol{\Phi}$ are the normalized axes of the reconstructed volume. In order to completely define the transformation ${}^C\boldsymbol{T}_T$, a center for the volume axes is needed. Although several approaches can be undertaken, in this study, the origin has been chosen to ensure that the coordinates of all the pixels are nonnegative. The position of the origin ${}^T\boldsymbol{x}_o$, is determined component-wise as follows:

$$ {}^T\boldsymbol{x}_{i_0} \leq {}^T\boldsymbol{x}_{i_p}, \tag{5} $$

where $i = 1, 2, 3$ indexes the vector components and $p = 1 \ldots N$ indexes every pixel.

Finally, the affine transformation ${}^C\boldsymbol{T}_T$ sought at this stage is:

$$ {}^C\boldsymbol{T}_T = \begin{bmatrix} \boldsymbol{\Phi} & {}^T\boldsymbol{x}_0 \\ 0\ 0\ 0 & 1 \end{bmatrix}. \tag{6} $$

Volume extent. The extent of the reconstructed volume turns out to be an important issue that has not been previously addressed in the literature. Some authors propose a volume size that encloses all the B-scans (Rohling et al. 1997; Meairs et al. 2000); however, this approach is clearly inefficient because outermost zones of the volume typically have a low information density, and any interpolation attempt increases the computation time and generates poor results. Henceforth, we will refer to this approach as the bounding box approach.

PCA provides a powerful tool to trim the volume size optimally without losing important information. Two facts are essential: the position components are uncorrelated and the pixel position variance for each dimension is known (the eigenvalues of the covariance matrix $\mathbf{\Sigma}$). A simple approach (hereafter referred to as the eigenvalue-driven approach) is to look for the volume size that fits into the bounding box and that retains the aspect ratio given by the eigenvalues in each direction. As we said before, the eigenvalues give an idea of how scattered the pixel positions are with respect to the mean value of the population; therefore, volume dimensions should preserve this ratio.

Assuming that the pixel position is normally distributed, ${}^{T}\boldsymbol{X} = N(0, \mathbf{\Sigma})$, an expression for the volume size can be worked out. For a normal distribution, surfaces of constant probability density are ellipsoids:

$$\boldsymbol{x}'\mathbf{\Sigma}^{-1}\mathbf{x} = C. \tag{7}$$

The principal axes of this ellipsoid correspond to the PCA of ${}^{T}\boldsymbol{X}$ (Anderson 1984). It can be found that the length of the ith ($i = 1, 2, 3$) principal axis of the ellipsoid with density C is:

$$l_i = 2\sqrt{\lambda_i C}, \tag{8}$$

where λ_i is the ith eigenvalue of the covariance matrix $\mathbf{\Sigma}$. The value of C can be calculated by integrating the density function within a prism that contains the ellipsoid defined by eqn (7) for a given probability r:

$$r = \frac{2^3}{(2\pi)^{3/2}(\lambda_1\lambda_2\lambda_3)^{1/2}} \times \int_0^{\sqrt{\lambda_1 C}}\int_0^{\sqrt{\lambda_2 C}}\int_0^{\sqrt{\lambda_3 C}} e^{-\frac{1}{2}\left(\frac{x_1^2}{\lambda_1}+\frac{x_2^2}{\lambda_2}+\frac{x_3^2}{\lambda_3}\right)} dx_1 dx_2 dx_3. \tag{9}$$

Therefore, the length of the trimmed volume along coordinate i that confines a given information density, say r, is:

$$l_i = 2\sqrt{2\lambda_i}\cdot erfinv\,(r^{1/3}), \tag{10}$$

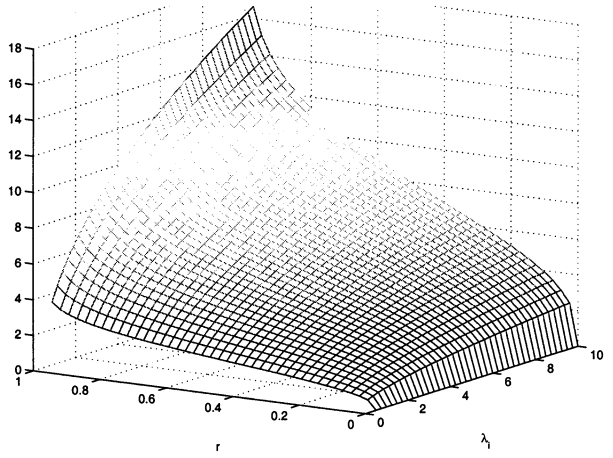

Fig. 4. Variation of volume length (l) with information density (r) and data dispersion (λ).

where *erfinv* is the inverse error function. Figure 4 shows the variation of the volume length related to r and λ. It is clear that, for each coordinate, the greater the density of information (r) and the data dispersion (λ), the larger the reconstructed volume length along this coordinate.

Voxel size. Voxel size computation aims at determining the grid spacing. On one side, it is obvious that the greater the grid spacing, the lower the interpolation requirements. On the other, the greater the grid spacing, the greater the loss of resolution in the reconstructed volume. Therefore, a satisfactory trade-off between the two tendencies must be achieved.

The original B-scans have a resolution given by the probe parameters; consequently, a reasonable criterion would be to choose the greatest voxel size so that after the B-scans are resampled into the volume, the planes of the B-scans do not suffer any loss of resolution (aliasing).

For the sake of clarity in the exposition, let us reduce the dimension of our problem. Instead of a volume, think of a plane grid and, instead of the B-scans, think of lines sampled with some sampling interval. Each sample would be a pixel intensity value. Figure 5 depicts the situation that is being posed; a 2-D object is sampled by means of lines. Suppose that we have a single line parallel to the x-axis; in this case, the maximum x-axis spacing would be the same as the line spacing (Δ). With only this line, no information would exist to find the spacing in the y-axis. As the line rotates around the x-axis counterclockwise, the optimum spacing would be reduced by a cos θ factor, where θ is the angle between the line and the x-axis. If we continued decreasing the spacing in the x-axis we would eventually reach a null spacing (*i.e.*, an infinite sampling frequency). This is obviously nonsense, so we take, as a threshold for finding the spacing in the x-axis, the value $\theta = \pi/4$. Similar considerations can be made about the y-axis. In this case,

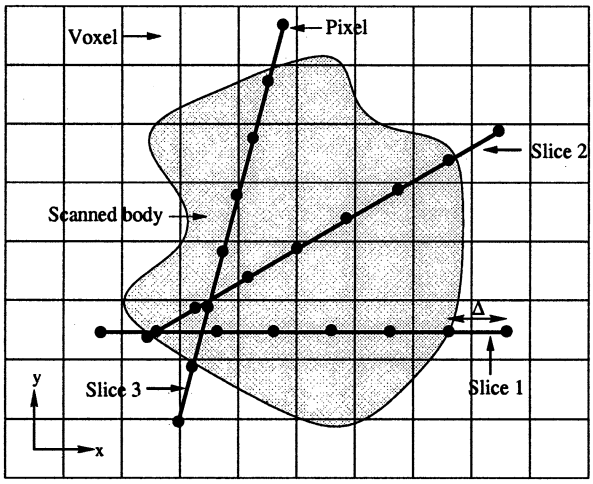

Fig. 5. Voxel size calculation. A volume is represented by a plane grid. A body is scanned by means of slices represented by lines and intensity pixels by dots. The spacing between dots is Δ.

the spacing increases with sin θ, with $\pi/4 < \theta < \pi/2$. An actual situation encompasses several lines (indexed by $i = 1, \ldots, m$) with orientations θ_i. The above procedure can be continued to each line, so the spacing, Δx and Δy, is finally given by:

$$\begin{cases} \Delta_x = \min_i \left(\Delta|\cos\theta_i|\right) & \text{if } \theta_i \leq \pi/4 \quad \forall i = 1, \ldots, m \\ \Delta_y = \min_i \left(\Delta|\sin\theta_i|\right) & \text{if } \theta_i > \pi/4 \quad \forall i = 1, \ldots, m. \end{cases} \tag{11}$$

The $\theta = \pi/4$ threshold has not been chosen at random. Assuming that the scanned body space variability is fully captured by the line resolution and that we have a high enough line density, the worst case, as far as spacing is concerned, arises when the lines lie parallel to the grid diagonal.

The 3-D case has a similar expression from the former discussion. It is more convenient to use a vectorial notation though. The *i*th B-scan (out of, say, M scans) is represented by of $\boldsymbol{s}^i$ and $\boldsymbol{q}^i$ (*i.e.*, unitary vectors that define the B-scan coordinate system). The volume axes are given by $\mathbf{x}$, $\mathbf{y}$ and $\mathbf{z}$ (see Fig. 3). $\Delta_{s^i} = \Delta_{\boldsymbol{s}}$ and $\Delta_{q^i} = \Delta_{\boldsymbol{q}}$ are the B-scan resolutions for directions $\boldsymbol{s}^i$ and $\boldsymbol{q}^i$, respectively, $\forall i = \{1, \ldots, M\}$. The grid spacing in each direction can be written as:

$$\Delta_l = \min_i \left(\Delta_{k^i}| < k^i, 1 > |\right)$$

$$\text{with } k^i = \begin{cases} \boldsymbol{s}^i & \text{if } | < \boldsymbol{s}^i, 1 > | \quad \geq \cos\pi/4 \\ \boldsymbol{q}^i & \text{if } | < \boldsymbol{q}^i, 1 > | \quad > \cos\pi/4 \end{cases}$$

$$\forall_i = 1, \ldots, M \tag{12}$$

where $\mathbf{1}$ denotes $\boldsymbol{x}$, $\boldsymbol{y}$ or $\boldsymbol{z}$ direction and $<, >$ denotes the inner product.

3-D reconstruction

Most 3-D freehand systems use similar algorithms to construct a regular volume after a volume grid has been established. Basically, all these methods transform each pixel position to a voxel position using eqn (1). Two situations can arise; 1) several pixels fall onto the same voxel; 2) several voxels are not intersected by any B-scan and are empty.

To face these two situations, the reconstruction process is divided into two subprocedures, namely, data fusion (bin filling) and hole filling.

Data fusion. A single voxel may contain several pixels due to different reasons. One of them is the fact that the voxel size may be larger than the B-scan pixel size. Another possibility is that each voxel may be intersected by several B-scans. Both situations generate redundant information that we have to manage somehow. As is stated by Rohling et al. (1997), every freehand system has to deal with compounding in some manner, because it is almost unavoidable that the scan planes intersect. A great activity has been carried out regarding spatial US compounding (Burckhardt 1978; Wagner et al. 1988) and, more specifically, in 3-D freehand ultrasound (Detmer et al. 1994; Rohling et al. 1998; Leotta and Martin 2000).

Averaging has been traditionally the compounding operation, because the speckle signal may be partially filtered out; theoretically, the improvement in the signal to noise ratio (SNR) may reach proportionality with the number of samples that are averaged. However, when the number of overlapping samples is small, the maximum sample-compounding technique seems preferable (Burckhardt 1978). Specifically, compounding pursues three objectives: speckle reduction, tissue boundary enhancement and shadowing reduction. It has been proven that averaging is an optimal way for speckle reduction; nonetheless, its optimality is not so clear as far as boundary enhancement is concerned. A tissue boundary exhibits different brightnesses depending on the insonation angle, so the averaging operation may dramatically blur the results. Burckhardt (1978) has shown that the maximum sample provides a similar improvement in SNR for a small number of pixels falling into the same voxels, and allows a clear boundary enhancement and shadowing reduction.

Hole filling. The hole-filling step aims at finding a value for those voxels that have not been filled in the data-fusion step. This present stage is an interpolation step and many of the attempts have focused on low-order interpolation. Given the great amount of data available,

NN approaches seem to efficiently tackle the problem. The NN approach involves sequencing through the regular volume (*i.e.,* through the original B-scans) and assigning to the empty voxel the value of the closest pixel.

In this paper, however, we have used an approach based on a basis function interpolation (Franke 1982) similar to the one proposed by Meairs et al. (2000). A Gaussian convolution kernel is placed on each voxel after the data-fusion step and the voxel value is computed as a weighted sum of all the voxels filled in the former step. Meairs et al. (2000) propose an ellipsoidal kernel, the main axis of which is oriented along the frame normal of the current image plane to avoid a smoothing of the image data in the horizontal direction; however, nothing is said about the selection of the variance of the Gaussian kernel. Moreover, it is not clear what the orientation of the image plane is for those voxels which turn out to be unfilled in the data fusion step.

Our approach is based on a spherical Gaussian kernel, the variance of which is a function of the variance of the intensity of the nearby pixels. This is an adaptation to the US case of what has been called sigma filter elsewhere (Lohmann, 1998). The variance of the intensity grey–scale images carries information about the extent of speckle formation (Dutt and Greenleaf, 1996). A speckle formation-driven interpolation allows us to detect likely structures out of the completely resolved speckle zones, reducing the smoothing effect of the Gaussian interpolation kernel by means of sharpening the kernel shape. Dutt and Greenleaf (1996) came up with a normalized variance-dependent parameter, f, that tracks the statistic of the speckle image:

$$f = \frac{\pi^2 D^2}{24\sigma_I^2}, \tag{13}$$

where D is a compression constant and σ_I^2 is the image variance. The statistical analysis carried out by the authors states that f parameterize scatterer densities; low densities result in fully resolved scatterers ($f = 0$) and large densities result in fully formed speckle ($f = 1$). The image variance, σ_I^2, can be estimated using the voxels filled during the bin-filling stage. The compression constant D stems from the logarithmic compression that the ultrasonic signal envelope must undergo to fit into the display's dynamic range. The authors propose a deterministic approach to resolve this parameter, although new blind inverse methods for gamma correction can be also applied (Farid 2001).

In this work, we have used f as a mapping parameter between the image variance and the kernel variance. Assuming that an interval of acceptable variances (σ_{min}^2, σ_{max}^2) can be defined, the adaptive variance of the Gaussian kernel, σ_K^2, is worked out as:

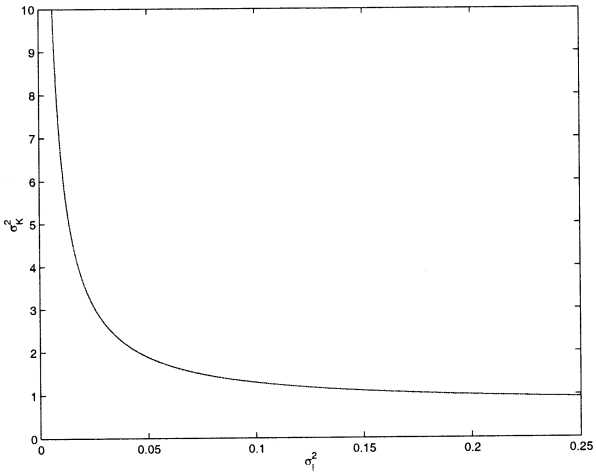

Fig. 6. Variation of Gaussian kernel variance vs. image variance. To preserve the structures in the reconstructed volume, the kernel variance decreases as the image variance increases. The intensity grey-scale value has been normalized between [0, 1] and a compression constant $D = 0.22$ was assumed. The desired kernel variances were set to lie within the interval $\sigma_{\min}^2 = 0.796$ and $\sigma_{\max}^2 = 10$.

$$\sigma_K^2 = \sigma_{\min}^2 + (\sigma_{\max}^2 - \sigma_{\min}^2) f, \tag{14}$$

The image variance, σ_I^2, is estimated using the voxels filled during the bin-filling step. This could be interpreted as an unsharp masking filtering over the kernel variance driven by the underlying speckle pattern. Figure 6 shows the variation of the kernel variance as the image variance increases. For computational purposes, instead of building the exact kernel for each variance, a logarithmic quantification of the former plot is devised. Hence, a bank of precomputed kernels is used to speed up the interpolation process. As a rule of thumb, we choose the size of the Gaussian kernel, k, so that it encloses 98.76% of the area within the Gaussian curve (*i.e.,* Trucco and Verri 1998).

$$k \geq 5\sigma_K. \tag{15}$$

The kernel function is then sampled at unitary distances and the values are normalized to maintain the DC component after the filtering.

RESULTS

A set of four US sweeps has been used to test the reconstruction process. US volumes from a thyroid and a liver were available from the Cambridge University 3-D

Ultrasound Research group. A kidney and a fetus[2] have been scanned under our supervision. All the scanned volumes are single-sweeps, except for the fetus case, which is formed by two overlapping sweeps.

Evaluation of the volume grid construction procedure

Evaluating the PCA process capability to define a coordinate system is a difficult task, due to the lack of a ground truth orientation; it is our opinion that there could be as many optimum orientations as ultrasonographers were asked to provide one, so the orientation given by other methods, like the Key-frame approach, could be equally valid. However, more objective differences between methods may arise when a trimming of the output volume is carried out to process only the relevant areas of the scanning. In this case, as we will show, the effectiveness of the trimming process is highly dependent on the chosen coordinate system.

Evaluation methodology. We will analyze five different coordinate systems and two different trimming procedures. About coordinate systems, these will be the PCA proposed in this paper, three choices of the Key-frame approach proposed by Barry et al. (1997) and the transmitter coordinate system given by the positioning system (hereafter denoted by T). The three choices of Key frames will be the B-scan in the middle of the recorded sequence, the B-scan located right at the end of the first quarter of the sequence starting from the beginning, and the B-scan located at the end of the third quarter from the beginning. These methods will be referred to as *KF#2Q, KF#1Q* and *KF#3Q,* respectively. In the examples, the numbers after the # sign will be the actual slice number of the Key frame within each dataset. About trimming procedures, Key frame and T will be trimmed with a deterministic method described later in this subsection, and PCA will be trimmed both with the eigenvalue-driven procedure and with the deterministic method. PCA methods will be referred to, respectively, as $PCA - eig$ and $PCA - det$.

For each data set, a careful delineation of the ROI has been carried out, assisted by the supervision of an expert using the 3-D slicer (Gering et al. 1999). The resulting ROIs will be taken as our ground truth (*i.e.,* as the target to be preserved in the reconstructed volume). The best procedure (both in terms of coordinate system and of trimming procedure) will be that that most compactly comprises the whole ROI (*i.e.,* the method with which, as soon as the whole ROI is captured within the reconstructed grid, the physical volume[3] of such grid is the least). For illustration purposes, Fig. 10 shows a 3-D model of the ROI with the original B-scans superimposed for each of the four datasets that will be used in the paper.

To monitor the trimming procedure, we will obtain ROC-like curves (Metz 2000), in which the magnitude in the horizontal axis will be the physical volume of the reconstructed grid (V_{grid}), and the magnitude in the vertical axis will be the portion of the ROI captured by the reconstructed grid. Such a magnitude will be hereafter referred to as captured ROI and denoted by V^c_{ROI}. Both axes will be normalized by the physical volume of the ROI (V_{ROI}), for the curves of the four datasets to be comparable irrespective of the actual size of the ROI. Such curves will be referred to as ROC curves, for simplicity. The physical volume V_{grid} for which $V^c_{\text{ROI}} = V_{\text{ROI}}$ will be hereafter referred to as optimal grid physical volume and denoted by V_{OGPV}.

To apply a trimming algorithm to those coordinate systems that are not based on an eigenvalue computation, a deterministic approach has been applied. This deterministic trimming approach has been introduced to compare the PCA with the Key frame. The chosen deterministic method follows an *ad hoc* scheme, where the length of each axis is reduced based on the deviation of the center of the volume bounding box from the center of mass of the data. Departing from the bounding box limits, each dimension of the bounding box is trimmed according to a term, d_j, which controls the extent of the trimming process in this direction as a function of the deviation of the center of mass with respect to the bounding box limits for the axis. Everything is weighted by a term p that is used to control the whole process. In our particular implementation, the length of the trimmed volume along the jth coordinate ($j = x,y,z$) is given by:

$$l_j = [(BB_j - mc_j)(d_j + 0.5) - (bb_j - mc_j)(1 - d_j)] * p, \quad (16)$$

where BB_j and bb_j are the respective jth coordinates of the lower and upper bounding box extremes, and mc_j is the jth coordinate of the center of mass. p is a coefficient that is used to control the trimming strength, and d_j is a parameter that accounts for the deviation from the center of mass following an arctangent shape. Specifically:

$$d_j = \frac{\pi + 2 \arctan\left[5\left(\frac{mc_j - bb_j}{BB_j - bb_j} - \frac{1}{2}\right)\right]}{4\pi}. \quad (17)$$

d_j is bounded over the interval $(0, 0.5)$ as mc_j tends to bb_j or BB_j.

[2] The corresponding author's first son!

[3] We will used interchangeably the terms reconstructed grid and reconstructed volume. When we need to refer to the size (in cubic units) of the reconstructed volume, we will use the term physical volume.

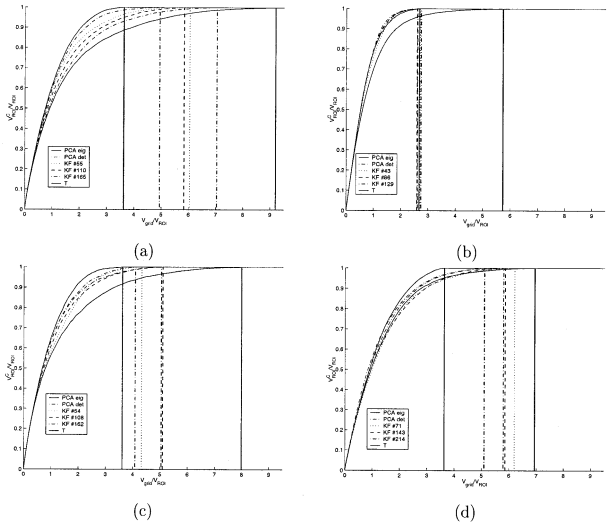

Fig. 7. Study of the evolution of the V^{c}_{ROI} with V_{grid}, both normalized by V_{ROI} (see main text for definitions). Four cases have been analyzed: (a) thyroid gland, (b) liver, (c) kidney, and (d) fetus. For each case, six methods have been studied, namely, PCA with eigenvalue trimming method, PCA with deterministic trimming method, Key frame with deterministic trimming method (the chosen Key-frame sequence number is used to label each case), and T coordinate system with deterministic trimming method. The vertical lines show the point at which each chosen coordinate system starts to enclose the ROI.

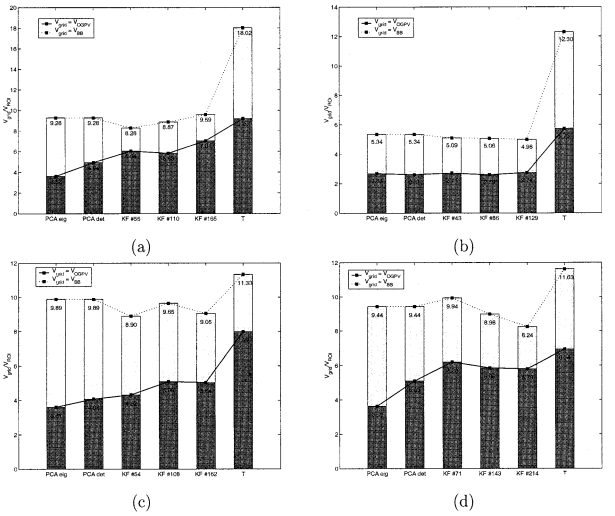

Fig. 8. Comparison between the V_{BB} and V_{OGPV} (see main text for definitions) for (a) thyroid, (b) liver, (c) kidney, and (d) fetus.

Coordinate system performance. Figure 7 shows the ROC curves that result from the different methods for the four datasets. These plots depict how fast the reconstructed volume comprises the ROI as a function of the growing physical volume of the reconstructed grid, as we increase either r or p, as convenient for the corresponding trimming method. Needless to say that the best procedure will grow faster in the vertical axis than the others, as they all move rightwards along the horizontal axis.

It can be seen that the PCA with eigenvalues method clearly outperforms the others in three of the four datasets (Figs. 7a, c, and d). It is only for the liver (Fig. 7b) that the proposed method it is not the best, even though differences are almost negligible. This can be explained very easily by noticing that, for this dataset, the ROI comprises almost the whole sequence (see Fig. 10b). If this is so, PCA does not add any particular advantage. On the other hand, if we focus on the fetus case (Fig. 7d), where two overlapping sequences (with different scanning directions) have been acquired to cover the whole fetus (see Fig. 10d), our method achieves the largest difference with respect to the other methods. The fetus case reveals the Key frame difficulty to choose the proper B-scan that clearly compiles the best trajectory. Moreover, the advantage of the eigenvalue with respect to the deterministic trimming methods is revealed, for this case, in terms of least physical volume enclosing the ROI, given the large amount of data that can be considered outlier (*i.e.*, outside the ROI). Similar conclusions can be drawn for the thyroid case (Fig. 7a), where the ROI is defined only within a subset of the dataset (see Fig. 10a), making the eigenvalue trimming method the best choice.

A deterministic trimming turns out to be a fairly poor solution, given that, though we are able to reduce the volume, there is no knowledge about how to perform this reduction depending on the axes and pixel distribution in the 3-D space (*i.e.*, how to trim the volume, preserving the ROI as much as possible for a given trimming ratio).

Another issue that we have addressed in our study is the relation between the least physical volume of the reconstructed grid V_{grid} that comprises the bounding box (say, $V_{\mathrm{grid}} = V_{BB}$) and V_{OGPV}. Fig. 8 shows this relation for each case and each coordinate system considered. The bargraphs reflect how worthy it is to use a trimming method in terms of the amount of data needed to enclose the ROI. Also, it is worth noting that PCA is not the method that achieves the least physical volume when the whole bounding box is required. That makes sense if we take into account that PCA tries to choose the best orientation to be as insensitive as possible to portions of B-scans not within the ROI; consequently, any attempt to enclose those B-scans will result in a greater volume than with other orientations that implicitly use the knowledge of the scanning policy (if it exists) as, for instance, the Key-frame approach. As we stated in the introduction, as

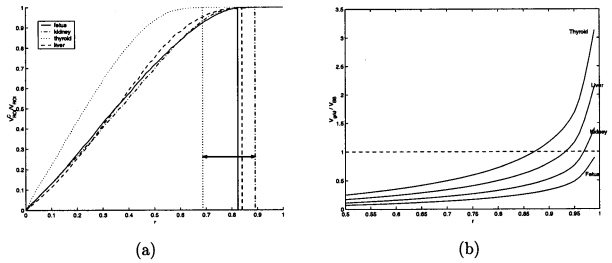

Fig. 9. (a) V^{c}_{ROI} normalized by V_{ROI} vs. r for each case. (b) V_{grid} normalized by V_{BB} vs. r. The horizontal dashed line represents the level where the trimming algorithm reaches V_{BB}. Note that, the closer the scanning trajectory to a straight line, the faster (with respect to r) V_{BB} is reached and, from that value r on, the theoretical trimming box is larger than the bounding box. It is obvious that the tails of the assumed Gaussian distribution does not hold any real voxel information.

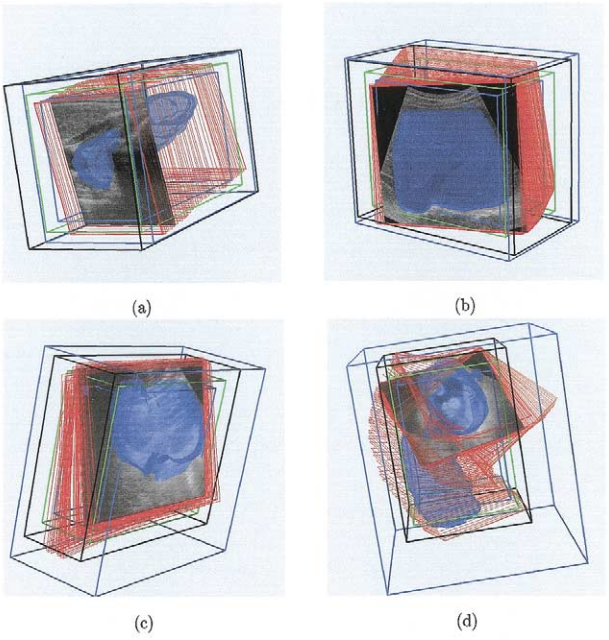

Fig. 10. Eigenvalue-driven trimming performance. The original B-scan planes, bounding box (outer blue) and boxes for r = 0.9, r = 0.8 and r = 0.7 (black, green and inner blue, respectively) are shown for the volumes (a) thyroid gland, (b) liver, (c) kidney, and (d) fetus. The difference between the bounding box volume and the trimmed volume are fairly clear.

far as a human observer is concerned, PCA is not the preferable option if the intention is to preserve all the data and the observer is willing to scroll through irrelevant B-slices. Our PCA approach becomes relevant when the trimming is carried out; in that case, as we have shown, our method outperforms the Key-frame approach in terms of physical size of the volume.

The evolution of the eigenvalue-driven physical volume with r is shown in Fig. 9. Figure 9a shows the evolution of the ratio $V^{c}_{\mathrm{ROI}}/V_{\mathrm{ROI}}$ for each case. From this plot, it can be seen that an interval ranging from r = 0.7 to r = 0.9 covers the ROI for the cases under study without adding information-lacking regions. For illustrative purposes, Fig. 10 shows the resulting reconstructed volume for the eigenvalue method. The original freehand B-scans are shown, as well as the boxes corresponding to the bounding box, r = 0.9, r = 0.8 and r = 0.7 (see figure caption for details). It is also worth analyzing the evolution of the V_{grid} normalized by the V_{BB}. This measure sheds light upon the sensitivity of the trimming approach with respect to the density of information r that we want to retain. Figure 9b plots this measure. The thyroid case turns out to be the most sensitive because its normalized volume quickly increases when r gets closer to *1*. On the other hand, the fetus is the least sensitive case. The conclusion that can be drawn is that, the greater the deviation of the scanning trajectory with respect to a straight line (either by several sweeps or misalignments of the original B-scans), the greater the bounding box and, consequently, the proposed trimming approach behaves more effectively.

Therefore, the PCA coordinate system and the eigenvalue driven approach constitute the basis of a formal framework for building and trimming at controlled distortion levels a regular volume out of irregularly sampled data. Results clearly show that the volume sizes so obtained are considerably smaller than those obtained by reported methods when the scanning trajectory is not correctly represented by a single slice within the dataset.

Voxel size. The selection of the voxel size is based on a criterion that retains the B-scans resolution into the volume grid. Table 1 shows the voxel sizes for different B-scan sequences and for different coordinate systems. The original B-scan resolution is also shown. For all the cases, it can be shown that the output resolution is bounded by the input resolution and a 1/2 factor of the input resolution. The upper bound is almost achieved in uniform scans like the thyroid case and the lower bound is more likely in several sweeps like the fetus case. Although the lower bound could be proposed as the general solution that guarantees our criterion, the proposed method tries to relax this lower bound by choosing the optimum for each particular case.

Evaluation of the resampling procedure

Different approaches can be used to evaluate the quality of the reconstruction process (bin- filling and hole-filling). For instance Meairs et al. (2000) state that one of the best ways to judge the quality of 3-D US volume acquisition and reconstruction is through volume visualization using rendering techniques. However, we

Table 1. Volume voxel sizes using the proposed method for the four datasets used in the paper

		Key-frame		PCA	
	B-scan pixel size (mm/pixel) $(s\ q)^T$	Vol. voxel size (mm/pixel)	Vol. dimensions $(x\ y\ z)^T$	Vol. voxel size (mm/pixel)	Vol. dimensions $(x\ y\ z)^T$
Thyroid	0.08713	0.08703	368	0.08713	672
	0.08913	0.08713	522	0.07923	519
		0.08713	595	0.08701	373
Liver	0.27986	0.27804	549	0.27658	578
	0.28526	0.27353	502	0.27499	527
		0.27986	391	0.27986	376
Kidney	0.26138	0.25967	635	0.25774	674
	0.28978	0.26138	556	0.26138	586
		0.26138	319	0.26138	297
Fetus	0.32335	0.22937	1127	0.25066	1216
	0.35377	0.25140	1043	0.22899	1433
		0.25140	1223	0.22946	1159

The output voxel size and the volume dimensions are shown both for the PCA and the Key-frame methods. The Key-frame corresponds to the middle B-scan of the recorded sequence. The volume dimensions correspond to V_{BB} for the given coordinate system.

believe that this approach remains subjective and that the visual cues associated to each particular rendering technique can potentially disguise pitfalls of the reconstruction method. One can even change the spatial position of the triangles in a triangle mesh according to different criteria to get a better visualization (San José et al. 2001), even though the original volume remains unaffected.

Another possibility could be to build a phantom in which the material had known acoustic properties so that values of the B-scan could be anticipated by analytical calculations and compared with the actual B-scans. This possibility, however, requires an effort that can be circumvented by other methods.

In this study, we have followed the testing approach introduced by Rohling et al. (1999). The main idea is to evaluate the ability of the reconstruction method to predict the intensity values at the locations where the original data have been removed. We will use the thyroid gland case as a figure of merit. This examination presents a regular enough sweep to allow an easy interpretation without any lack of generality. A B-scan near the middle of the sweep is selected. A Key-frame approach is used to align the volume coordinate system using the chosen B-scan. The voxel size is set so that x and y axes have a resolution equal to the B-scans. Therefore, the selected B-scan falls perfectly into voxels in the volume grid. To preserve this property, a voxel size in z should be properly chosen. The distance between the centers of the adjacent B-scan from the key B-scan is used as the voxel size in the z direction.

A percentage of pixels is randomly removed from the B-scan, creating holes of various sizes that should be interpolated. The rest of the pixels are used in the interpolation process to fill in all the voxels in the voxel array. The interpolated value is compared with the removed pixel value to assess the reconstruction accuracy. The average of the absolute difference between the interpolated and the original values over all missing values is used as a goodness measure

$$V = \frac{1}{L} \sum_{i=1}^{L} |p_i - c_i|, \tag{18}$$

where p_i is the original pixel intensity value and c_i is the voxel interpolated intensity value aligned with p_i. L is the number of discarded pixels that need an interpolation value. V also includes the quantization error, because the output values are stored with integer precision.

The tests have been carried out with six different percentages of removed data: 25%, 50%, 75%, 100%, 300% and 500%. The tests corresponding to 25%, 50%, 75% and 100% remove data from the chosen B-scan. Percentages over 100% mean that additional pixels from the two adjacent B-scans are removed. Specifically, for the 300% test, all the pixels from the chosen B-scan and the two closest adjacent scans to it have been removed. Evaluating the 300% and 500% accurately requires an examination with slices almost parallel and with one sweep so that the missing voxels are mostly aligned with the ground truth removed and a reliable comparison can be carried out. This is the reason for the choice of the thyroid case to perform the test. A kernel size of $3 \times 3 \times 3$ voxels has been used for the smallest percentages.

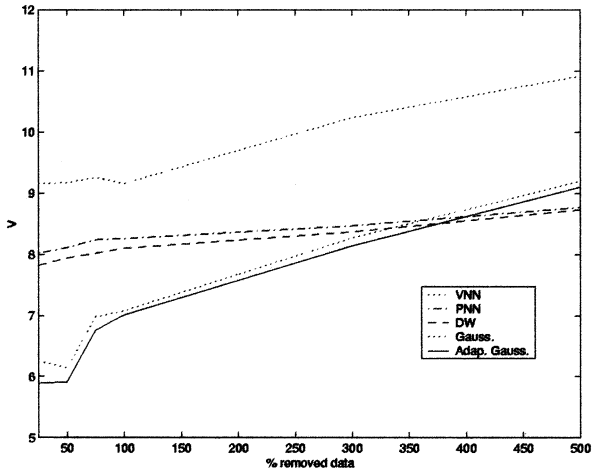

Fig. 11. Grey-level error for the thyroid case. Four well-known interpolation methods (VNN, PNN, DW, and Gaussian convolution kernel) are compared with the proposed method (adaptive Gaussian convolution kernel).

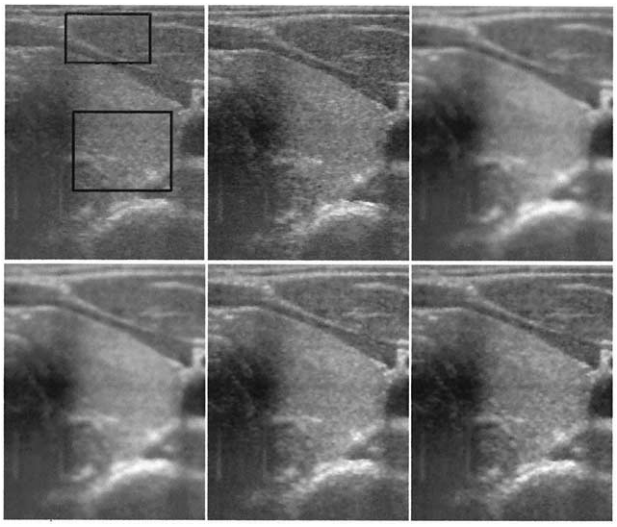

Fig. 12. Thyroid B-scans interpolated by several methods. All the images show the result of the interpolation for a 100% percentage of data removed. From left to right, row 1: the original slice (before data removal) and it is the *ground truth* VNN, PNN; row 2: DW, Gaussian and adaptive Gaussian.

For the 300% and 500% tests, kernel sizes were $5 \times 5 \times 5$ and $7 \times 7 \times 7$ voxels, respectively. A $7 \times 7 \times 7$ mask has been used to estimate the image variance that drives the adaptive Gaussian interpolation procedure.

Figure 11 plots the results for some well-known methods and the one proposed in this study. The results show that a Gaussian approach performs much better than well-known methods, as long as the percentage of data removed is below a certain threshold. Focusing on results below 300%, the method introduced in this paper clearly outperforms former proposals, closely followed by a traditional Gaussian interpolation. For percentages over 300%, the great lack of information makes those techniques that do not underestimate farther pixels, like PNN, behave better in terms of the error metric considered.

For illustrative purposes, Fig. 12 shows the set of interpolated images for the 100% case, (*i.e.,* when the whole slice has been interpolated using out-of-plane information for the methods discussed in the paper) together with the original image. The figure shows that both DW and PNN tend to oversmooth the data. About the other three methods, no important differences can be found; only minor comments can be made. For instance, VNN seems slightly to overemphasize the texture patterns. This is the case within the upper rectangle drawn in the original image. With respect to the Gaussian interpolators, as we said, their behavior is very similar; however, the nonadaptive version tends slightly to smooth the texture patterns. This is the case for the area within the rectangle located in the middle of the original image. These comments, however, do not lead to a clear preference for one particular method.

The second example highlights the main differences of the methods; Fig. 13 shows a slice of the reconstructed volume for the case of the kidney dataset. The PCA-based proposed method has been used to work out the reconstruction volume grid and the trimming algorithm has been set up with $r = 0.70$. The VNN method computes the nearest pixel in a spherically-shaped local neighborhood with radius $R = 5$ voxels. For the PNN and DW methods, the kernel size has been set up to $7 \times 7 \times 7$. For the adaptive Gaussian method, we have chosen $\sigma_{\min} = 0.8$ and $\sigma_{max} = 1.4$ as the desired marginal deviations. The kernel

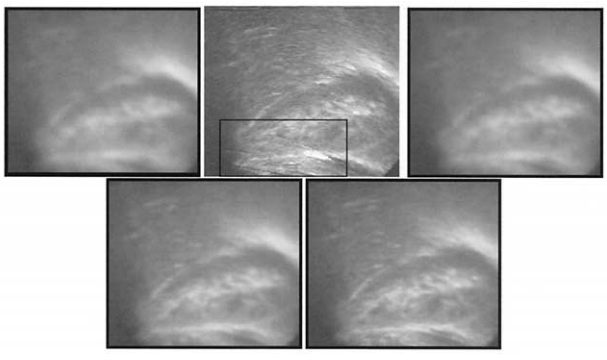

Fig. 13. Kidney B-scans interpolated by several methods. From left to right, row 1: PNN, VNN, DW; row 2: Gaussian, adaptive Gaussian.

variance will be bounded between these two values, (see eqn (14)). Then, the adaptive Gaussian kernel size is computed using eqn (15). As before, PNN and DW tend to oversmooth data. Both Gaussian kernels show a smoother behavior than VNN, especially the non-adaptive version, which creates a blurry image as compared with the other two. VNN shows a more textured pattern, as in the previous example. However, the price to pay is the presence of misalignments in the resampled image, due to the overlap of the original slices. This is particularly clear in the lower contour of the kidney, where, for an important part of it (enclosed within the rectangle shown), a textured diagonal pattern has been superimposed, distorting its presence and making it difficult to be traced. The adaptive Gaussian method turns out to be the one that shows the best trade-off between hole-filling capabilities and blurring. Misalignment artefacts are not visible. Consequently, although these misalignments are common to all the methods, kernel-based methods can more efficiently deal with them.

CONCLUSIONS

In this paper, a theoretical framework for the reconstruction of a regular US volume out of irregularly-sampled volume data has been proposed. We have focused on two main problems, namely, how to create a regular volume grid and how to find the value of each of the voxels in the grid. With respect to the former, the main contribution of the paper is the proposal of a well-known statistical tool to find the optimum coordinate system; in our experiments, we have shown that the method proposed in the paper encloses the ROI with a smaller overall physical volume than other methods previously proposed in those cases that the scanning trajectory differs from a pure straight line; as a matter of fact, we have shown that the greater the reduction of the volume size of our method, the greater the deviation of the scanning trajectory with respect to a straight line. In addition, we have proposed a method to find the size of each dimension of the voxel so as to reduce aliasing effect while maintaining low storage requirements. With respect to the second problem, we have proposed an adaptive Gaussian kernel to carry out the interpolation procedures, and numerical results show that this procedure achieves the minimum distortion for the datasets used in the paper. In addition, graphical results show that this method achieves a moderate blur while avoiding misalignments artefacts that stem from the superposition of slices. Although more extensive data analyses should be done to infer possible correlations between the scanning trajectories and the appropriate value of the parameter r to perform the trimming, it is our belief that this paper shows a formal methodology to reconstruct volumes from freehand US scanning while maintaining moderately low processing and storage needs. The graphic tools here developed (an example of which is shown in Fig. 10) allow practitioners to fine-tune the final volume size at will very easily.

Acknowledgment—The authors acknowledge the Comisión Interministerial de Ciencia y Tecnología for research grants (IFD97-0881-C02 and TIC01-3881-C02) and Junta de Castilla y León for research grants (VA78/99 and VA91/01). The authors appreciate the support given by Julio Diaz, M.D., Hospital Río Carrión, Palencia, Spain in performing the scanning, by James Ellsmere, M.D., Brighman and Womens Hospital, Boston, USA, in delineating the ROI, and by Miguel Martín-Fernández, University of Valladolid, Spain, in Vtk and Tcl/Tk issues. The first author acknowledges the useful discussions with Dr. Carl-Fredrik Westin, Brighman and Womens Hospital, Boston, USA, about contents and presentation of the paper. The authors also thank the anonymous reviewers for their help in improving the quality of the paper.

REFERENCES

Anderson TW. An introduction to multivariate statistical analysis, 2nd ed. New York: John Wiley & Sons, 1984.

Barry CD, Allott CP, John NW, et al.. Three-dimensional freehand ultrasound: Image reconstruction and volume analysis. Ultrasound Med Biol 1997;23(8):1209–1224.

Burckhardt CB. Speckle in ultrasound B-mode scans. IEEE Trans Sonics Ultrason 1978;25(1):1–6.

Detmer PR, Bashein G, Hodges T, et al.. 3-D ultrasonic image feature location based on magnetic scan head tracking: In vitro calibration and validation. Ultrasound Med Biol 1994;20:923–936.

Dutt V, Greenleaf JF. Adaptive speckle reduction filter for log-compressed B-scan images. IEEE Trans Med Imaging 1996;15(6):802–813.

Farid H. Blind inverse gamma correction. IEEE Trans Image Proc 2001;10(10):1428–1433.

Franke R. Scattered data interpolation: Tests of some methods. Math Comput 1987;38(157):181–200.

Gering D, Nabavi A, Kikinis R, et al. An integrated visualization system for surgical planning and guidance using image fusion and interventional imaging. In: Medical image computing and computer-assisted intervention, lecture notes in computer science. Berlin: Springer-Verlag 1999:809–819.

Jain AK. Fundamentals of digital image processing. Princeton, NJ: Prentice Hall, 1989.

Leotta DF, Detmer PR, Martin RW. Performance of a miniature magnetic position sensor for three-dimensional ultrasound imaging. Ultrasound Med Biol 1997;23(4):597–609.

Leotta DF, Martin RW. Three-dimensional ultrasound imaging of the rotator cuff: Spatial compounding and tendon thickness measurement. Ultrasound Med Biol 2000;26(4):509–525.

Lohmann G. Volumetric image analysis. New York: John Wiley & Sons, 1998.

Meairs S, Beyer J, Hennerici M. Reconstruction and visualization of irregularly sampled three- and four-dimensional ultrasound data for cerebrovascular applications. Ultrasound Med Biol 2000;26(2): 263–272.

Metz CE. Handkook of medical imaging. Vol. 1. SPIE Press 2000.

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

Prager RW, Gee A, Berman L. Stradx: Real-time acquisition and visualization of freehand three-dimensional ultrasound. Med Image Anal 1999;3(2):129–140.

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

Rohling RN, Gee AH, Berman L. Three-dimensional spatial compounding of ultrasound images. Med Image Anal 1997;1(3):177–193.

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 comparison of freehand three-dimensional ultrasound reconstruction techniques. Med Image Anal 1999;3(4):339–359.

San José R, Alberola-López C, Ruiz-Alzola J. Reshaping polygonal meshes with smoothed normals extracted from ultrasound volume data: An optimization approach. Proc. SPIE 2001;4325:462–472.

Sanches JM, Marques JS. A Rayleigh reconstruction/interpolation algorithm for 3D ultrasound. Pattern Recog Lett 2000;21(10):917–926.

Schroeder W, Martin K, Lorensen B. The visualization toolkit: an object-oriented approach to 3D graphics. Princeton, NJ: Prentice-Hall, 1998.

Treece GM, Prager RW, Gee AH. Regularised marching tetra-hedra: Improved iso-surface extraction. Computer Graphics 1999;23(4): 583–598.

Treece GM, Prager RW, Gee AH, Berman L. Surface interpolation from sparse cross sections using region correspondence. IEEE Trans Med Imaging 2000;19(11):1106–1114.

Treece GM, Prager RW, Gee AH, Berman L. 3D ultrasound measurement of large organ volume. Med Image Anal 2001;5(2):41–54.

Trucco E, Verri A. Introductory techniques for 3-D computer vision. Princeton, NJ: Prentice Hall, 1998.

Wagner RF, Insana MF, Smith S. Fundamental correlation lengths of coherent speckle in medical ultrasonic images. IEEE Trans Ultrason Ferroelec Freq Control 1988;35:34–44.