RADIAL BASIS FUNCTION INTERPOLATION FOR 3-D ULTRASOUND

R.N. Rohling, A.H. Gee and L. Berman

CUED/F-INFENG/TR 327

July 1998

Cambridge University Engineering Department
Trumpington Street
Cambridge CB2 1PZ
England

E-mail: rnr20@eng.cam.ac.uk, ahg@eng.cam.ac.uk, lb@radiol.cam.ac.uk

Abstract

3-D freehand ultrasound imaging produces a set of irregularly spaced B-scans, which are typically reconstructed on a regular grid for visualisation and data analysis. Most standard reconstruction algorithms are designed to minimise computational requirements and do not exploit the underlying shape of the data. We investigate whether approximation with splines holds any promise as a better reconstruction method. A radial basis function (RBF) approximation method is implemented and compared with three standard methods. The RBF's are shown to accommodate both dense, overlapping data and sparse data, without introducing the kind of reconstruction artifacts common with the standard methods. The other potential advantages of RBF's, such as the direct computation of derivatives, make further investigation worthwhile.

1 Introduction

Ultrasound is a widely used imaging modality that has a large variety of clinical applications. Conventional 2-D diagnostic imaging is performed with a hand-held probe which transmits ultrasound pulses into the body and receives the echoes. The magnitude and timing of the echoes are used to create a 2-D grey-level image (B-scan) of a cross-section of the body in the scan plane. 3-D ultrasound imaging extends this concept, so that a volume of intensity data is created from pulse-echo information.

There are several methods of performing 3-D ultrasound imaging [4]. One method, freehand imaging, makes use of conventional ultrasound technology to build up a 3-D data set from a number of 2-D B-scans acquired in rapid succession. This is achieved by attaching a 3-D position sensor to the probe, so that each B-scan can be labelled with the position and orientation of the scan plane. The physician then moves the probe slowly and steadily over a particular anatomical region, so that the set of acquired B-scans encompasses the volume of interest with few gaps. Figure 1 shows the elements of a 3-D freehand system.

While the set of B-scans can be reviewed individually, reconstructing them into a regular array makes it possible to use conventional 3-D visualisation and data analysis tools. These tools include any-plane re-slicing, volume rendering, surface rendering, segmentation, and registration procedures. The reconstruction step is important because any loss of image quality or the introduction of artifacts must be avoided. For example, a small loss in image quality during reconstruction can make a barely detectable pathology in a B-scan become undetectable in the reconstructed data. This may result in a misdiagnosis. In this paper, we propose a new reconstruction technique to improve on existing techniques.

The paper is organised as follows. In Section 2, the existing reconstruction methods are reviewed and the new reconstruction method is described. In Section 3, the acquisition system, two *in-vivo* examinations, and a series of comparative tests are described. The results, followed by a discussion on the potential of the new technique, are given in Section 4. Finally, conclusions and suggestions for future work are given in Section 5.

2 Reconstruction Methods

2.1 Overview

Since the motion of the probe in freehand imaging is controlled by the physician, the B-scans can be at any relative position and orientation. This means the pixels (B-scan elements) lie at irregular locations in the array of voxels (volume elements). Therefore,

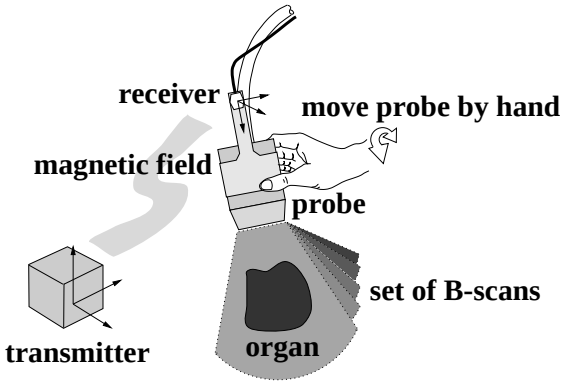

Figure 1: **3-D freehand ultrasound imaging**. Freehand imaging allows the physician to move the probe freely so the B-scans can have arbitrary relative locations and may overlap each other. As the probe moves, the B-scans are captured and stored in computer memory. A magnetic position sensor measures the position of a receiver (mounted on the probe) relative to a transmitter (fixed with respect to the patient's body). The relative location of each scan plane is then calculated and the B-scans combined into a single volume of data.

the reconstruction problem can be classified as unstructured, or scattered, data interpolation [1] [14].

A survey of the literature reveals a number of different methods for reconstruction of 3-D freehand ultrasound data sets. Most of these methods are very simple because they are designed to minimise the time and memory required for reconstruction. This is because physicians want to visualise the 3-D data sets immediately after acquisition, so the reconstruction should take ideally only a few seconds. Although almost all freehand systems perform reconstruction at some stage, details of the reconstruction method are often unpublished since the method is considered *ad hoc*. Nevertheless, the methods that have been published can be classified into the following categories: voxel nearest neighbour interpolation, pixel nearest neighbour interpolation, and distance-weighted interpolation.

2.2 Voxel Nearest Neighbour Interpolation

The concept of voxel nearest neighbour (VNN) interpolation is easy to understand: each voxel is assigned the value of the nearest pixel. There are no parameters to set. A naive implementation would traverse the array one voxel at a time and calculate the value of the nearest pixel, but this would be computationally inefficient . Using the fact that the nearest pixel lies along a line normal to the nearest B-scan greatly speeds up the reconstruction, making it one of the fastest of all methods. Moreover, a new 3-D ultrasound imaging system, developed at the University of Cambridge [16], can rapidly produce slices of the set of B-scans without reconstructing an entire voxel array. By cleverly using the voxel nearest neighbour interpolation method with dedicated graphics hardware, slices can be generated interactively.

While this reconstruction method has the advantage of avoiding gaps in the voxel array, reconstruction artifacts can be observed in slices through the voxel array. When a slice

[1] Since the pixels all lie on planes, the data is actually best described as semi-structured.

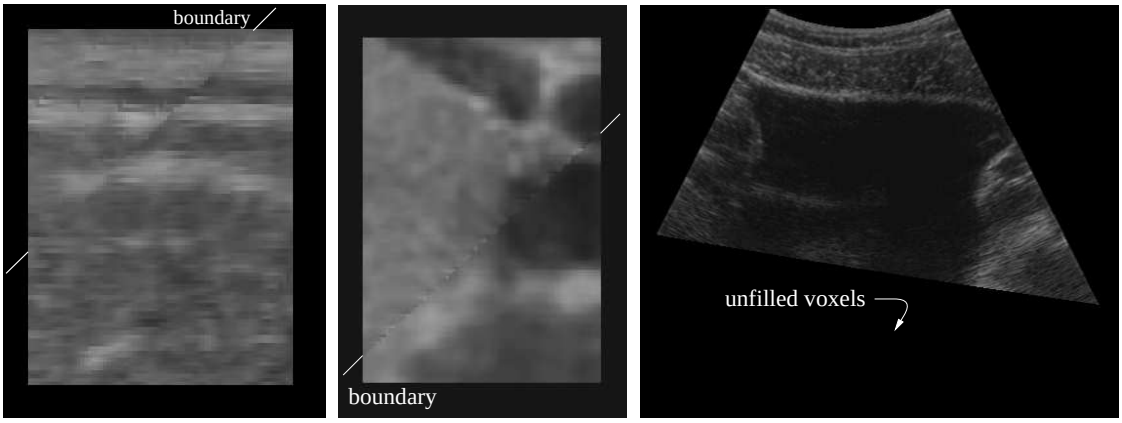

(a) VNN (b) PNN (c) DW

Figure 2: **Reconstruction artifacts.** In (a) a voxel nearest neighbour interpolation is used to reconstruct an examination of the neck. A small slice of the voxel array is shown. The slight misalignment between the projections of the two nearest B-scans arises from registration errors. In (b) a pixel nearest neighbour interpolation is used to reconstruct an examination of the thyroid. The boundary between the voxels filled by the first "bin-filling" stage and voxels filled by the second "hole-filling" stage is evident. In (c) a distance-weighted interpolation is used to reconstruct an examination of the bladder. The slice is truncated because the neighbourhood for selecting points is set too small, leaving some voxels unfilled.

plane intersects several of the original B-scans, we can consider the interpolated image as a collage of projections from the intersected B-scans. Registration errors, including tissue motion and sensor errors, contribute to slight misalignment of the B-scans. This results in a piece of the collage slightly mismatching its neighbours. The lines of intersection between the pieces then become visible — see Figure 2(a).

2.3 Pixel Nearest Neighbour Interpolation

Pixel nearest neighbour interpolation (PNN) is one of the most popular reconstruction methods [5, 8, 9, 13, 18]. The basic algorithm consists of two stages. In the first stage, the algorithm simply runs through each pixel in every B-scan and fills the nearest voxel with the value of that pixel. Multiple contributions to the same voxel are usually averaged, although the maximum value has also been suggested [13]. The parameters to set at this stage are therefore the weights on the multiple contributions.

If the voxel size is small compared to the distance between the acquired B-scans, gaps can occur in the voxel array. In practice, this situation is inevitable with voxel arrays of similar resolution to the B-scans. The second stage fills these remaining gaps in the voxel array. A variety of methods have been used, including averaging of filled voxels in a local neighbourhood [9, 13] and interpolating between the two closest non-empty voxels in the transverse direction to the B-scans [8]. Other publications [5, 18] do not describe the method for filling gaps, so we assume that they adopt similar approaches. or choose

the voxel size sufficiently large to avoid gaps. The parameters to set at this stage are the weights of the nearby voxels used to fill the gaps.

In summary, this method can be considered as a two stage process: the first stage of "bin-filling" the voxel array with pixels is very fast, but the "hole-filling" second stage may take longer, depending on the particular method chosen. Unfortunately, artifacts can be generated by this two stage process. For example, a slice plane passing though regions of both first stage and second stage filled voxels may show the boundary between the highly detailed "bin-filled" voxels and the smoothed "hole-filled" voxels — see Figure 2(b).

2.4 Distance-Weighted Interpolation

Like the voxel nearest neighbour interpolation method, distance-weighted (DW) interpolation proceeds voxel by voxel but assigns a value to each voxel based on a weighed average of some set of pixels from nearby B-scans. The parameters to choose are the weight function and the size and shape of the neighbourhood.

The simplest approach is to consider a fixed spherical neighbourhood of radius R, centred about each voxel [2]. All pixels is this neighbourhood are weighted by the inverse distance to the voxel and then averaged. This is similar to Shepard's method [19]. Unfortunately, it has the disadvantage of not reproducing any of the local shape properties implied by the data because the resulting interpolant typically has local extrema at the data sites [14]. It also requires R to be set prior to reconstruction. If R is set too small, gaps may result — see Figure 2(c). Yet if R is set too large, the voxel array will appear highly smoothed, since the effect of inverse distance weighting can be quickly overwhelmed by the much larger number of data points falling into the larger local neighbourhood. Nevertheless, with dense B-scans and a small value of R, excellent results are claimed [2].

Another distance-weighted interpolation method uses a non-uniformly shaped neighbourhood to account for the asymmetric shape of the point spread function of the ultrasound beam [15]. Essentially, the resolution is much higher within the B-scan than between the B-scans, so the authors create a "thick" B-scan by convolving the 2-D B-scan image with a truncated 3-D Gaussian kernel. The 3-D Gaussian kernel is chosen to be wide perpendicular to the B-scan and narrow within the B-scan plane. A similar "thick" slab of weight values of the truncated 3-D Gaussian kernel is computed. The voxel array accumulates both weight and intensity values separately, at those voxels intersected by the "thick" B-scans [2]. The final voxel intensity values are calculated by dividing the accumulated pixel intensity values by the accumulated weight values stored at that voxel. In summary, each voxel value represents a weighted average of nearby B-scans, with an asymmetric Gaussian weight function attached to each B-scan. While this method has the advantage of being incremental, gaps may still remain if the truncated Gaussian weight function is smaller than the gaps between B-scans.

Another distance-weighed method which uses a non-uniform neighbourhood is based on pair-wise B-scan grouping [21]. In this method, each voxel is filled only by the nearest two B-scans which fall on either side of it. For each of the two B-scans, lines are drawn perpendicular to the B-scan plane, passing though the voxel. Bi-linear interpolation within each B-scan determines the contributing pixel value from that B-scan. The voxel is then set to the inverse distance-weighted average of the two contributing pixel values. This

[2] The authors also include a third value related to the "age" of the intensity value, so that new overlapping B-scans will overwrite old B-scan values.

method has the advantage of retaining the resolution of each B-scan in the voxel array and avoiding gaps. Although it is clear which pairs of B-scans to use with simple motions, such as linear sweeps, it is not clear which pairs to use with more arbitrarily located B-scans. It has been suggested that the pairs be chosen time sequentially [3], but this may not be optimal for all types of scanning motions.

2.5 Radial Basis Function Interpolation

It is perhaps intuitively obvious to attempt a traditional functional interpolation method for the reconstruction problem. In other words, choose a particular function (such as a polynomial), determine the polynomial coefficients to make it pass through the pixel values, then evaluate at regular intervals to produce a voxel array. There have been no previously published attempts at functional interpolation of 3-D freehand ultrasound data, since there are severe computational demands to overcome. For example, simple volume splines [14] require solving a set of N linear equations, where N is the number of pixels summed over all B-scans. This requires $O(N^3)$ calculations. For modern computer workstations, N can not be much larger than 500 before this becomes prohibitively slow and numerical inaccuracies arise [14]. This makes a typical ultrasound reconstruction problem of several million pixels infeasible.

There is a great deal of literature on interpolation, spanning a variety of disciplines, and progress has recently been made in the area of tackling very large data sets. Two good review papers are [6, 14]. In order to select an appropriate method for solving the 3-D ultrasound reconstruction problem, a list of requirements is listed below:

- the method must interpolate scattered trivariate data
- the method must be fast, i.e. $O(N)$ complexity and non-iterative
- the interpolating function must be smooth
- both interpolation and approximation must be possible
- large overshoots must be avoided

The smoothness requirement arises from an assumption that the input data (the set of B-scans) is smooth itself. This assumption arises from knowledge of the acquisition process. While the underlying anatomy is not smooth, it is measured by a finite width ultrasound beam with a smooth intensity profile. This is tantamount to convolution of the anatomical function with the beam profile, creating a smoothed image. Further signal processing, such as filtering of the pulse-echo signals, gives even more confidence that the image data to be interpolated is smooth.

The approximation requirement arises from the existence of measurement errors. These errors include tissue motion (such as breathing, pulsative and whole body motions), position sensor errors and calibration errors. This means we want to be able to change the interpolating function into an approximating function that passes close to, but not exactly through, the data points.

The final requirement of eliminating overshoots relates to the desire to have the range of interpolated voxel values in the same [0,255] grey-level range as the B-scans. This requirement is related to the approximation requirement, since an approximating function can often reduce overshoots compared to interpolating functions.

After surveying the recent advancements in trivariate interpolation of large data sets, we discovered a method ideally suited to our requirements. This method was recently developed by researchers at the University of Illinois for interpolation of multivariate geographical data sets [3] [12]. They dubbed the method "completely regularized splines with tension". The Illinois method is summarised as follows:

Consider a set of pixel values $p_j, j = 1, \ldots, N$ that are located at the positions $\mathbf{x}_j$, where $\mathbf{x}_j = (x_j, y_j, z_j)$ is expressed with respect to the voxel array. The basic idea is to find a spline $S(\mathbf{x})$ that passes as closely as possible to the data points and is as smooth as possible. These two requirements can be combined together such that we find the $S(\mathbf{x})$ that fulfils

$$\sum_{j=1}^{N} |p_j - S(\mathbf{x}_j)|^2 + wI(s) = \text{minimum}. \tag{1}$$

The first component is the deviation of the spline from the data points, and the second is a smoothness function $I(s)$. The weight w determines the relative cost of the two components.

The solution can be expressed as

$$S(\mathbf{x}) = T(\mathbf{x}) + \sum_{j=1}^{N} a_j \mathbf{R}(\mathbf{x}, \mathbf{x}_j) \tag{2}$$

where $T(\mathbf{x})$ is the trend function and $\mathbf{R}(\mathbf{x}, \mathbf{x}_j)$ is a radial basis function (RBF) whose form depends on the choice of $I(s)$ [20].

For the 2-D case, if $I(s)$ is chosen to minimise the cost of the second (partial) derivatives only, then the familiar thin plate spline results. If the same $I(s)$ is used for the 3-D case, the first derivatives of the RBF are divergent at the data points [10]. By carefully choosing a more general $I(s)$, it is possible to obtain a simple analytic expression for the RBF with regular derivatives of all orders everywhere [11]. This choice results in $T(\mathbf{x}) = a_0$, a constant, and

$$\mathbf{R}(\mathbf{x}, \mathbf{x}_j) = \frac{\phi^3}{4\pi} \left[\frac{1}{\phi r} \text{erf}\left(\frac{\phi r}{2}\right) - \frac{1}{\sqrt{\pi}} \right] \tag{3}$$

where $r = |\mathbf{x} - \mathbf{x}_j|$ is the distance from $\mathbf{x}$ to $\mathbf{x}_j$, and erf is the error function [1]. The parameter ϕ is a generalised tension parameter, and it controls the distance over which the point influences the resulting hypersurface. The multiplicative constant $\phi^3/4\pi$ can be omitted, since it can be combined with the coefficients a_j. A plot of the RBF is shown in Figure 3.

The spline coefficients can then be found by solving the set of linear equations

[3] see also `http://www.cecer.army.mil/grass/viz/VIZ.html`

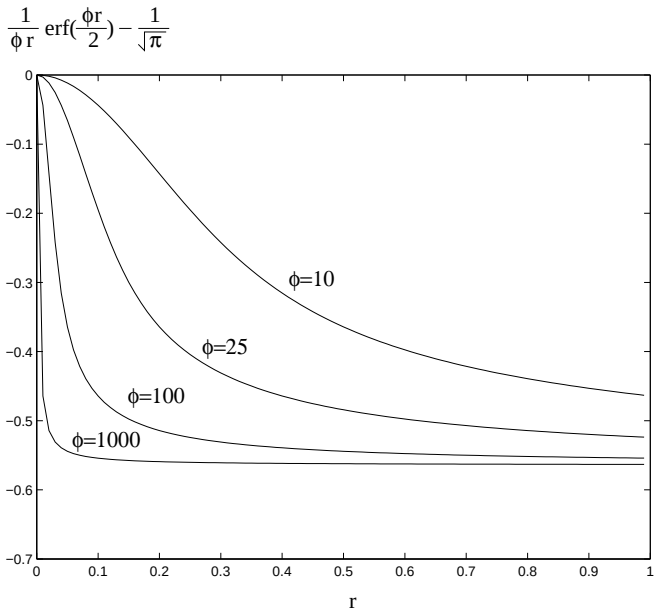

Figure 3: **Radial basis function.** Increasing the tension reduces the range of influence of the radial basis function.

$$a_0 + \sum_{j=1}^{N} a_j \left[\mathbf{R}(\mathbf{x}, \mathbf{x}_j) + \delta_{ij} w\right] = p_i \tag{4}$$

$$\sum_{j=1}^{N} a_j = 0 \tag{5}$$

where δ_{ij} is the Kronecker delta function.

There are therefore two parameters which can be adjusted to tune the nature of the interpolant: ϕ controls the tension, and w controls the level of approximation. The goal of tuning the interpolant is to find the optimal balance between the requirements of obtaining small deviations from the data points and avoiding overshoots. As Figure 4(a) shows, a high level of tension limits the distance at which each point influences the overall interpolant. Yet the RBF appears spiky and does not replicate the overall shape of the data. Low tension results in overshoot. Figure 4(b) shows that the overshoot from low tension interpolation may be effectively controlled by allowing a small amount of approximation. In fact, it will be shown that the combination of low tension and a small level of approximation works well with ultrasound data.

As previously mentioned, the RBF interpolant cannot be calculated using all the data points of an ultrasound examination at once. In order to localise the volume splines, the scattered input data must be divided into manageable **segments**. The basic idea is that the interpolating function in a local region is not influenced by data at some sufficiently distant point. The voxel array is therefore divided into many small, non-overlapping rectangular segments. Individual interpolating functions are then calculated for each segment until all of the voxel array is covered.

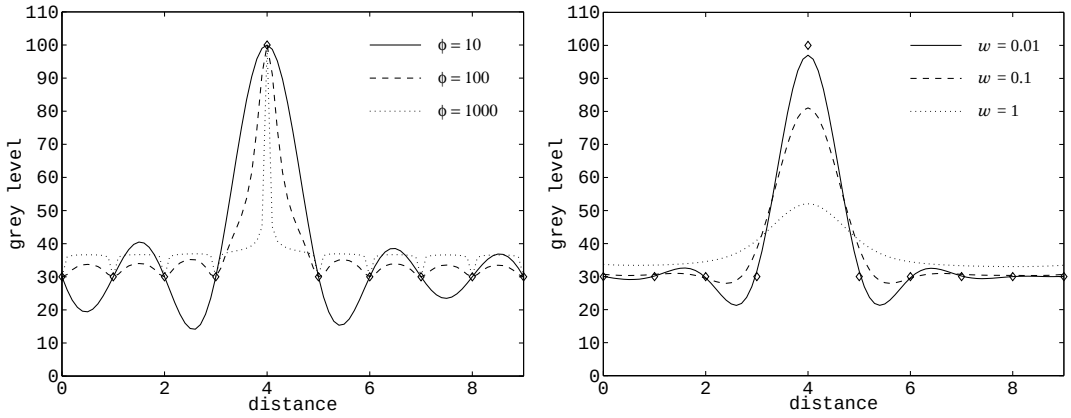

(a) Interpolation

(b) Approximation

Figure 4: **1-D radial basis function interpolation and approximation.** 10 sample data points obtained from an impulse function are shown as diamonds. In (a) w is set to zero so the RBF exactly passes through the data points. The average of the data points sets the trend part of the interpolant to a value of 37. With high tension, the interpolation quickly returns to the trend between data points. With low tension, overshoots appear. Increasing the tension can be thought of as changing the nature of the interpolant from a stiff plate to a membrane. In (b) the tension level ϕ was set to a value of 20. Increasing the level of approximation reduces the amount of overshoot but can result is large deviations from the data points. A small level of approximation produces a reasonable trade-off.

In order to get smooth connections among the RBF's of neighbouring segments, overlapping **windows** are used. This means that a window is established around each segment under consideration such that it encompasses not only all the data points in the segment but also a sufficient number of neighbouring data points. All data points in the window are then used to calculate the RBF for that segment. Since the windows overlap each other, the RBF for each segment will closely (but not necessarily exactly) match its neighbours' RBF's. With a reasonable amount of overlap, the differences will be negligible compared to the quantisation of the grey levels in the 8-bit voxel array.

Segments of fixed size, each containing fewer than a maximum set number of data points (K_{max}), were proposed in [11]. Windows were expanded from $3 \times 3 \times 3$ surrounding segments to $5 \times 5 \times 5$ and so on, until a sufficient number (K_{min}) of neighbouring points were included. In later work, the segmentation method was developed using an oct-tree representation of the voxel array to improve the ability to interpolate data with a heterogeneous spatial distribution [12]. This means the segments were continuously divided into eight subsegments until each contained no more than K_{max} points. In this way, segments of variable size were created to account for the clustering of data points. The window around each segment expanded equally in all directions until it encompassed at least K_{min} points.

If the data is only mildly heterogeneous, this windowing method works well. However, the problem of 3-D ultrasound reconstruction involves highly clustered data; all of the input data lie within the planes of the B-scans. For this reason, we have developed a more

flexible windowing method.

The basic problem with expanding a window until a fixed number of data points is encompassed, is that all the data points may lie on only one side of the segment. This can create problems with continuity of the interpolating function between segments. We propose to expand the window around the segment until data points are found in all directions around the segment. Examples of the conventional and improved windowing methods are shown in Figure 5.

In the 3-D case, the window is expanded in all directions at first, but each of the six faces of the window stops expanding only when a sufficient number of data points fall within the region defined by that direction — see Figure 6. In practice, expansion of the window in a given direction also stops when it reaches the extents of the voxel array.

The Illinois software was designed to work with 1000 to 10000 data points, so a number of practical changes were made. In order to accommodate several million data points from an ultrasound examination, points in the oct-tree are stored using recursive dynamic memory allocation. We also designed the oct-tree segmentation to allow segments to be divided into four (quad-tree) or two (binary-tree) subsegments when one dimension of the segment is only one voxel wide. This improves the ability to divide the volume into manageable segments. However, some very large segments can remain. For example, a single empty segment between two nearly parallel B-scans can extend occasionally the full length and width of the B-scans. The window around this segment, expanding one voxel at a time, will suddenly encompass an entire neighbouring B-scan in a single step. This produces an intractable number of data points, so we choose to chop these large, empty segments into smaller segments too.

In summary, we have a localised trivariate spline with $O(N)$ complexity. The method is efficient enough to handle very large data sets and flexible enough to handle highly heterogeneous data. The smooth interpolating function has regular derivatives of all orders, and has good accuracy compared to other interpolation methods [11]. It also has a variable tension and can be tuned between interpolation and approximation.

3 Comparisons

3.1 Apparatus

The freehand acquisition system comprises an ultrasound scanner, a position sensor and a computer for data acquisition. A Toshiba Powervision 7000 scanner (Toshiba America Medical Systems, Tustin, California) was used with a Polhemus FASTRAK magnetic position sensor (Polhemus Incorporated, Colchester, Vermont) mounted on the probe. Calibration to determine the location of the position sensor with respect to the probe was performed using the Cambridge phantom [17]. The Stradx acquisition software [16] was used in conjunction with an 8-bit frame grabber and a Silicon Graphics Indy workstation (Silicon Graphics Incorporated, Mountain View, California). The images were matched to the position sensor readings and recorded at a rate of 25 frames/s.

3.2 *In-vivo* Examinations

Two *in-vivo* examinations were performed on a healthy human subject. First, an examination of a thyroid gland was performed with a 7 MHz linear array probe. A depth setting of 40 mm was used giving a resolution of 0.087 mm/pixel. Each B-scan was cropped to

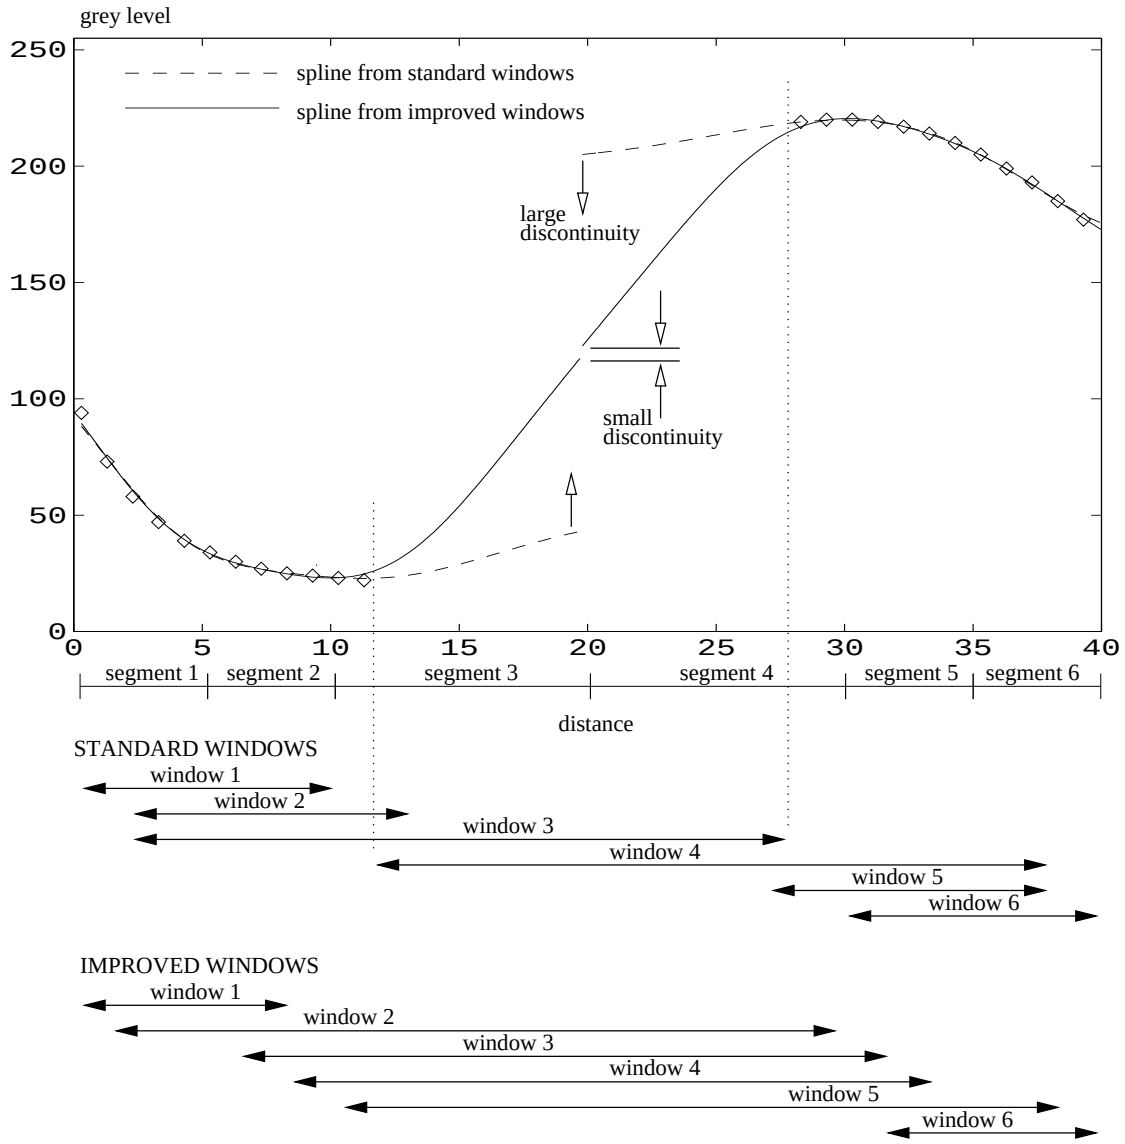

Figure 5: **Example of improved windowing method.** This contrived 1-D example illustrates the potential problem of the standard windowing method. Segments are created by subdividing the length of interest in a binary-tree so that each segment contains no more than 5 (K_{max}) data points. The RBF's are calculated (with tension $\phi = 3$, and smoothing $w = 0.01$) for each segment using the data that falls in the surrounding window. The standard windowing method expands in both directions from the segment in single increments until 10 (K_{min}) data points in total are obtained. The standard windows 3 and 4 only contain data on one side of the segment. This creates a large discontinuity in the interpolating function between segments 3 and 4. The new windowing method expands in both directions until at least 4 points are found on each side of the segment, or the limits of the data set are encountered. This way, the interpolants for segments 3 and 4 contain data from both sides of the segments and therefore meet much closer together.

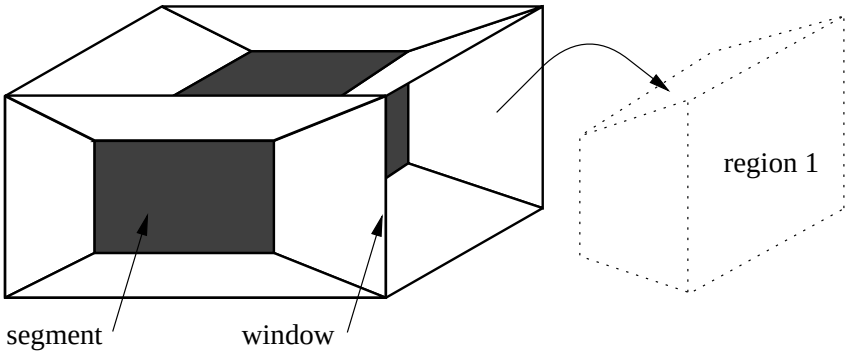

Figure 6: **Segment and surrounding window.** The segment is the shaded volume surrounded by the window. The window must encompass data points in each of the six regions around the segment, one of which is highlighted. This is done by expanding each of the six faces of the window until each region contains a minimum number of data points. Determination of which region a point falls into is calculated efficiently using the cross products of the vectors to the corners of the planes separating the regions.

328×409 pixels. A single sweep of the organ with a slow and steady motion resulted in a dense set of 219 nearly parallel B-scans. Figure 7 shows the outlines of the set of B-scans and a typical B-scan.

The second examination was performed with a fan shaped sweep over the extents of the bladder. A 3.75 MHz convex curvilinear array probe was used with a depth setting of 140 mm, giving a resolution of 0.34 mm/pixel. Because the probe produces sector shaped B-scans, the acquired 480×413 pixel images were masked so that only the ultrasound intensity data was used in each of the interpolation methods. Figure 8 shows the outlines of the set of B-scans and a typical B-scan.

These two examinations were chosen because they allow tests of the reconstruction methods with different organs, probes, depth settings, ultrasound machine settings and types of probe motion.

3.3 Tests

The tests are designed to evaluate the ability of the reconstruction methods to interpolate the ultrasound data and fill in gaps. Since the true underlying anatomical function is unknown, we have decided to test the different methods by artificially removing data from the two examinations. The four different reconstruction methods are then evaluated on their ability to predict the intensity values at the locations where the data was removed. In other words, a good reconstruction method will interpolate the removed data points with values very near to the data that was originally there.

First, a B-scan near the middle of the sweep is selected. The voxel array (with voxels equal in size to the pixels) is aligned exactly with this B-scan such that pixels fall exactly onto voxels. A percentage of the pixels are then removed randomly from the B-scan, creating gaps of various sizes. The rest of the pixels and all other B-scans in the reconstruction are used in the interpolation to fill in all voxels in the voxel array. The values of the removed (original) pixels can now be compared to the values of the voxels aligned with them. The voxel array stores the interpolation results as floating point numbers to avoid

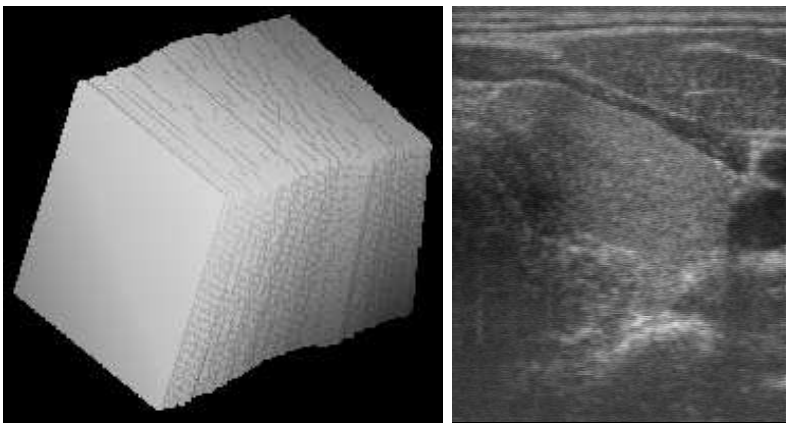

(a) Outlines (b) B-scan

Figure 7: **Thyroid examination.** The outlines of the B-scans are shown in (a) with a typical B-scan shown in (b).

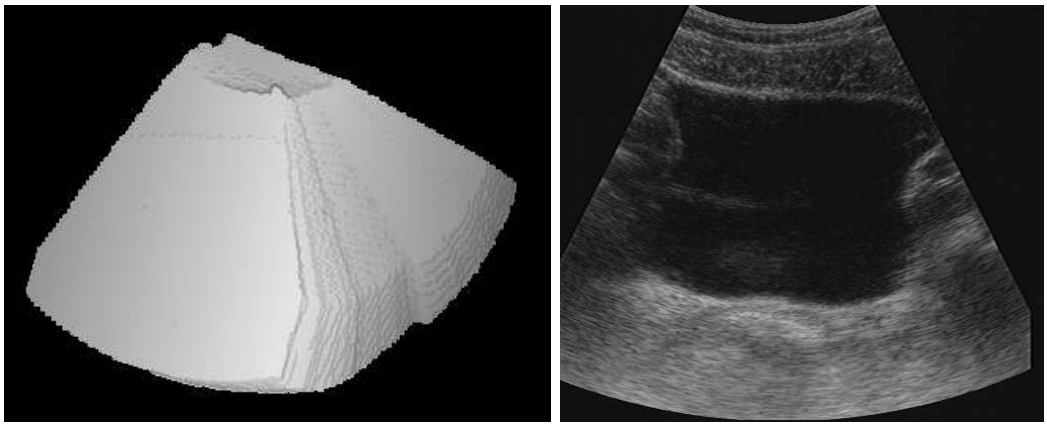

(a) Outlines (b) B-scan

Figure 8: **Bladder examination.** The outlines of the B-scans are shown in (a) with a typical B-scan shown in (b).

the influence of quantisation. The average of the difference between the interpolated and the original data over all missing data points is calculated by

$$V = \frac{1}{M} \sum_{i=1}^{M} |p_i - v_i| \tag{6}$$

where p_i is the original pixel that was removed from the reconstruction, v_i is the interpolated value of the voxel aligned with p_i and M is the number of removed pixels. A low value of V indicates a good ability to interpolate over the gaps.

The tests are performed with eight different percentages of removed data: 0%, 25%, 50%, 75%, 100%, 300%, 500% and 700%. For the 25%, 50%, 75% and 100% tests, pixels are removed only from the selected B-scan n. The 300% test refers to removing all of the pixels of B-scan n and all of B-scans $n-1$ and $n+1$. The 500% and 700% tests further remove B-scans $n \pm 2$ and $n \pm 3$ respectively. The 0% test is also included, since a reconstruction method may not exactly replicate the original data points. For example, a functional approximation method will miss the data points. For the 0% test alone, V is calculated over all pixels of the selected B-scan.

Each of the eight tests are repeated for ten different B-scans to give mean and variance estimates of V. In this way, 80 voxel arrays are created for each of the four reconstruction methods, producing 320 voxel arrays in total. A further 320 voxel arrays are created in the same way for the bladder examination.

3.4 Reconstruction Methods

The voxel nearest neighbour interpolation method is implemented by traversing the voxel array and filling each voxel with the value of the nearest pixel. The pixel nearest neighbour interpolation method is implemented in two steps. The first step assigns each pixel to the nearest voxel in the array. Multiple contributions to a single voxel are averaged together. The second step fills the remaining gaps. Empty voxels are filled by taking the average of the filled voxels in a $3 \times 3 \times 3$ neighbourhood. The remaining unfilled voxels are then filled by averaging originally filled voxels in a $5 \times 5 \times 5$ neighbourhood and so on, until all voxels are filled. This is similar to the method described in [13].

The reported high quality reconstructions described in [2] led us to choose the inverse distance-weighted method with a spherically shaped local neighbourhood. The one parameter to choose in this method is the radius, R, of the neighbourhood. If R is set to cover the largest gaps in the 700% tests, it would be much too large for the tests with smaller gaps and excessive smoothing would occur. R must be set to reasonable but not arbitrary values for each of the tests.

In [2], a 7 MHz linear array probe was used to examine the carotid artery, a similar examination to the thyroid gland. Two R values were suggested: 0.25 mm and 0.5 mm. In our examination of the thyroid gland, the distances between the centres of the B-scans range from 0.09 mm to 0.52 mm, with a mean of 0.32 mm. Using a neighbourhood of radius R equal to 0.25 mm is therefore not large enough to fill in all gaps, so 0.5 mm is used in the 0%, 25%, 50%, 75% and 100% tests [4]. For the 300%, 500% and 700% tests,

[4] It appears that setting R to half of the maximum centre distance between B-scans would fill the entire volume without gaps. But because the scan planes are not exactly parallel, larger gaps than 0.52 mm exist. In fact, the value of 0.5 mm is just sufficient to cover all gaps.

R must be increased to cover all gaps. We choose to increase R by the mean (0.32 mm) of the B-scan spacing. This means R was set to 0.82 mm (0.5 mm + 0.32 mm), 1.14 mm (0.5 mm + 2×0.32 mm), and 1.46 mm (0.5mm + 3×0.32 mm) for the 300%, 500% and 700% tests respectively.

To be consistent, the values of R for the bladder reconstruction are set in a similar fashion. R is set to the maximum distance between B-scans for the 0%, 25%, 50%, 75% and 100% tests, then increased by the mean distance between B-scans for the 300%, 500% and 700% tests. Since the examination used a fan-shaped sweep, the B-scan spacing is measured at the bottom centre of the B-scan, where the gaps are larger. The B-scan spacing ranges from 0.41 mm to 0.80 mm, with a mean of 0.60 mm. The value of R is therefore set to 0.80 mm for the 0%, 25%, 50%, 75% and 100% tests, and 1.40 mm (0.80 mm + 0.60 mm), 2.00 mm (0.80 mm + 2×0.60 mm) and 2.60 mm (0.80 mm + 3×0.60 mm) for the 300%, 500% and 700% tests respectively.

The RBF method is implemented using the oct-tree segmentation and improved windowing technique described in Section 2.5. K_{max} is set to 30 data points, and each region of the window is required to contain at least 5 data points.

For the 300%, 500% and 700% tests of the thyroid examination, large gaps and extremely long computation times result. The window is therefore expanded faster in the two directions perpendicular to the selected B-scan (towards the neighbouring B-scans) than in the four directions in the plane of the B-scan. This avoids making the window too large and enveloping an intractable number of data points. We do this simply to speed up the calculations, since such a large overlap is not required to get smooth transitions among the segments. As a reassurance, at least 20 points per region are required for these 300%, 500% and 700% tests, compared to the minimum of 5 points per region for all other tests.

The tension and approximation parameters are tuned manually by testing V for a randomly selected B-scan. A fairly low tension ϕ combined with a small amount of data smoothing w results in a good trade-off between minimising overshoot and passing very near the data points. This, in turn, produces a small value of V. For the thyroid examination, we set $\phi = 17$ and $w = 0.01$. The bladder examination uses $\phi = 25$ and $w = 0.1$. These values fall within the range of values typically used for geographic data interpolation [12].

4 Results

The overall trends of the test results are illustrated in Figure 9 and tabulated in Tables 1 and 2. To give an indication of the distributions of V over the 10 trials, box-and-whisker plots are shown in Figures 10 and 11. A small but representative set of images of the interpolated data are shown in Figures 12, 13, 14 and 15.

4.1 Voxel Nearest Neighbour Interpolation

For both examinations, V is zero at 0% data removal, since the voxels are set to their nearest neighbours: the original pixels of the B-scan. At 25%, 50% and 75%, the nearest neighbours of the voxels are mainly the remaining pixels of the selected B-scan. Therefore, the resulting interpolated images appear as a patchwork of irregularly shaped pieces and relatively large values of V result.

Test	VNN		PNN		DW		RBF	
	μ	σ	μ	σ	μ	σ	μ	σ
0%	$0.00^{\star}$	0.00	$0.00^{\star}$	0.00	$0.00^{\star}$	0.00	0.64	0.03
25%	$6.65^{\dagger}$	0.12	$7.05^{\dagger}$	0.12	$7.73^{\dagger}$	0.27	4.82	0.25
50%	$6.95^{\dagger}$	0.14	$7.31^{\dagger}$	0.13	$7.82^{\dagger}$	0.30	6.14	0.33
75%	$7.86^{\dagger}$	0.16	$7.82^{\dagger}$	0.13	$7.84^{\dagger}$	0.35	6.89	0.39
100%	8.84	2.60	$7.78^{\dagger}$	0.64	7.90	0.42	7.51	0.78
300%	$11.66^{\dagger}$	2.00	9.85	0.75	9.48	0.47	9.56	0.34
500%	$12.90^{\dagger}$	1.65	$11.27^{\dagger}$	0.73	10.43	0.53	10.67	0.29
700%	$13.99^{\dagger}$	1.32	$13.03^{\dagger}$	1.00	$10.95^{\star}$	0.52	11.32	0.23

Table 1: **Interpolation error V for the thyroid examination**. μ is the mean of V and σ is the standard deviation. $\dagger$ means that the assertion $\mu > \mu_{RBF}$ is statistically significant for a confidence level of 0.05. $\star$ means that the assertion $\mu < \mu_{RBF}$ is statistically significant for a confidence level of 0.05. The assertions are tested with the paired-sample t-test statistical method [7].

Test	VNN		PNN		DW		RBF	
	μ	σ	μ	σ	μ	σ	μ	σ
0%	$0.00^{\star}$	0.00	$0.00^{\star}$	0.00	$0.00^{\star}$	0.00	0.96	0.03
25%	$5.60^{\dagger}$	0.39	$5.01^{\dagger}$	0.22	$5.37^{\dagger}$	0.09	3.57	0.25
50%	$5.50^{\dagger}$	0.40	$5.08^{\dagger}$	0.25	$5.32^{\dagger}$	0.09	3.85	0.31
75%	$5.27^{\dagger}$	0.50	$5.19^{\dagger}$	0.35	$5.24^{\dagger}$	0.10	4.13	0.40
100%	4.13	0.38	$5.25^{\dagger}$	0.40	$5.11^{\dagger}$	0.14	4.29	0.37
300%	6.92	0.40	$7.03^{\dagger}$	0.15	$6.85^{\dagger}$	0.12	6.69	0.19
500%	$8.50^{\dagger}$	0.23	$7.80^{\dagger}$	0.14	$7.62^{\star}$	0.11	7.73	0.16
700%	$9.37^{\dagger}$	0.26	8.36	0.18	$8.07^{\star}$	0.09	8.37	0.16

Table 2: **Interpolation error V for the bladder examination**. See the caption of Table 1 for an explanation of the tabulated terms and symbols.

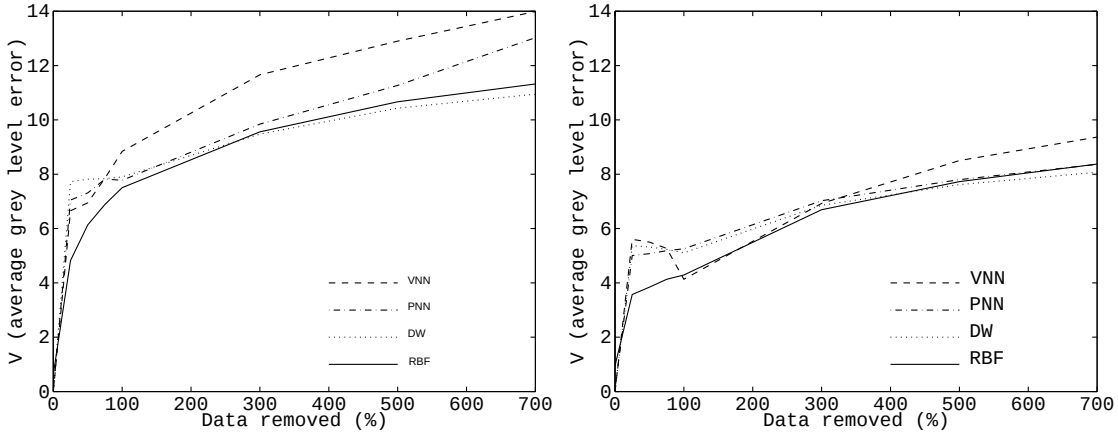

(a) Thyroid

(b) Bladder

Figure 9: **Average grey level error for the thyroid and bladder data sets**. For each of the four reconstruction methods, V is calculated for various percentages of removed data.

For the 100% to 700% tests, the interpolated image is made from the projection of the pixels from the nearest B-scan, so the patchwork appearance disappears. For the thyroid examination, the variance of V is large because some of the projections of neighbouring B-scans closely match the original B-scan, but others do not. This depends on both the level of registration error and the similarity of the pixel values. The values of V are generally greater than the other three reconstruction methods.

For the bladder examination, the 100% to 700% tests show a smaller variance of V than the thyroid examination, but the variance is still large compared to the other methods. The mean, however, is lowest at 100%, and increases again for the higher percentage tests. This may be explained by comparing the level of registration errors to the voxel size. If registration errors, such as position sensor error, are of similar absolute magnitudes for the two examinations, the errors will have a smaller effect on the bladder tests. This is because the voxel size is larger in the bladder tests than in the thyroid tests, so the effect of the errors is reduced. Yet, as more and more B-scans are removed, the relative registration errors increase, thereby making V increase again.

In general, however, the results look sharp and detailed for the 100% to 700% tests, since the projected data exhibits no blurring. Moreover, the boundaries between the portions of the projected data are not discernible for any tests of either examination. This suggests that the registration errors are mainly small and the images vary slowly from one B-scan to the next.

4.2 Pixel Nearest Neighbour Interpolation

For both examinations, V is zero at 0% data removal, since the nearest neighbours to the voxels are the original pixels. At 25%, 50% and 75%, the gaps are filled mainly with an average of the remaining pixels in the original B-scan. The interpolated image appears as

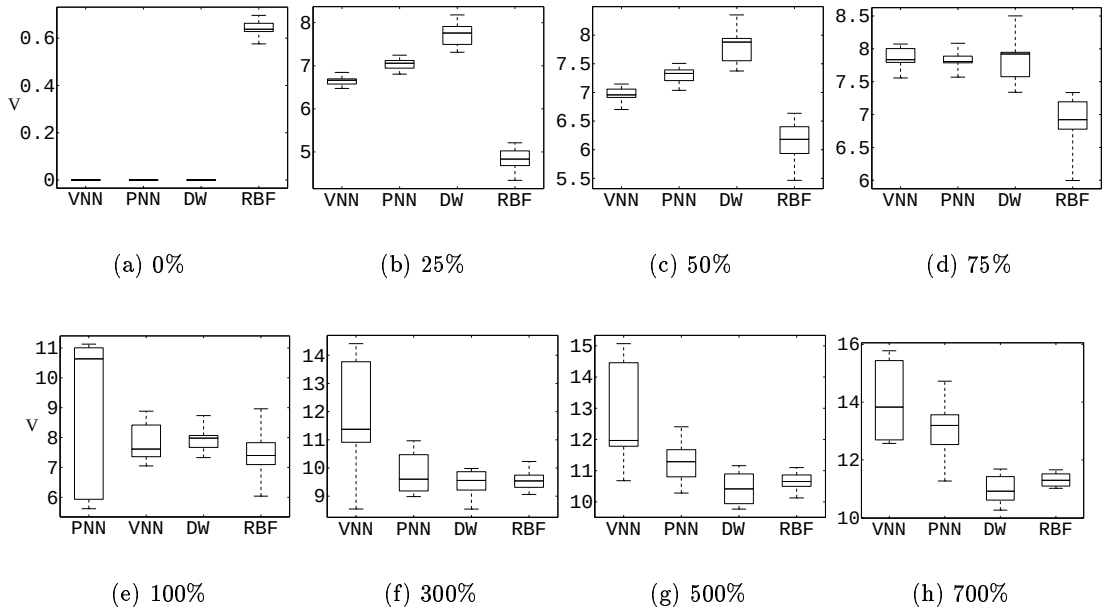

(a) 0% (b) 25% (c) 50% (d) 75%

(e) 100% (f) 300% (g) 500% (h) 700%

Figure 10: **Interpolation error V for the thyroid examination.** The box-and-whisker plots show the three quartiles and the extremes of V calculated over the 10 trials at each percentage of removed data. The median (50th percentile) is the line inside the box of the 25th and 75th percentiles of V. The whiskers show the minimum and maximum values.

a patchwork again, and relatively large values of V result.

For the thyroid examination, the mean of V increases progressively for the 500% and 700% tests compared to the relatively good performance in the bladder examination. This can be explained by noticing that the gaps in the bladder data set are much smaller at the top than the bottom, which limits the blurring effect of the "hole-filling" stage to the large gaps at the bottom of the bladder images. With lower amounts of blurring in the bladder tests, lower values of V result.

Unfortunately, the images exhibit significant reconstruction artifacts for both examinations, especially for the 500% and 700% tests — see especially Figures 13(d) and (h). A visible boundary exists between portions which are filled, for example, by averages of a $7 \times 7 \times 7$ neighbourhood, and portions filled by a neighbourhood of $9 \times 9 \times 9$, because they involve different amounts of smoothing. The boundaries between the voxels filled by the first "bin-fill" stage and second "hole-fill" stage are also visible.

4.3 Distance-Weighted Interpolation

For both examinations, V is zero for the 0% tests, since the weights on the original pixels that fall exactly on the voxels approach infinity. At 25%, 50% and 75%, the images can be considered as the superposition of two effects. The first effect arises from the voxels filled by the original data weighted by infinity. The second effect arises from the voxels in the gaps that are calculated from a weighted average of many neighbouring pixels. The resulting interpolated data is therefore a combination of the remaining pixels from the

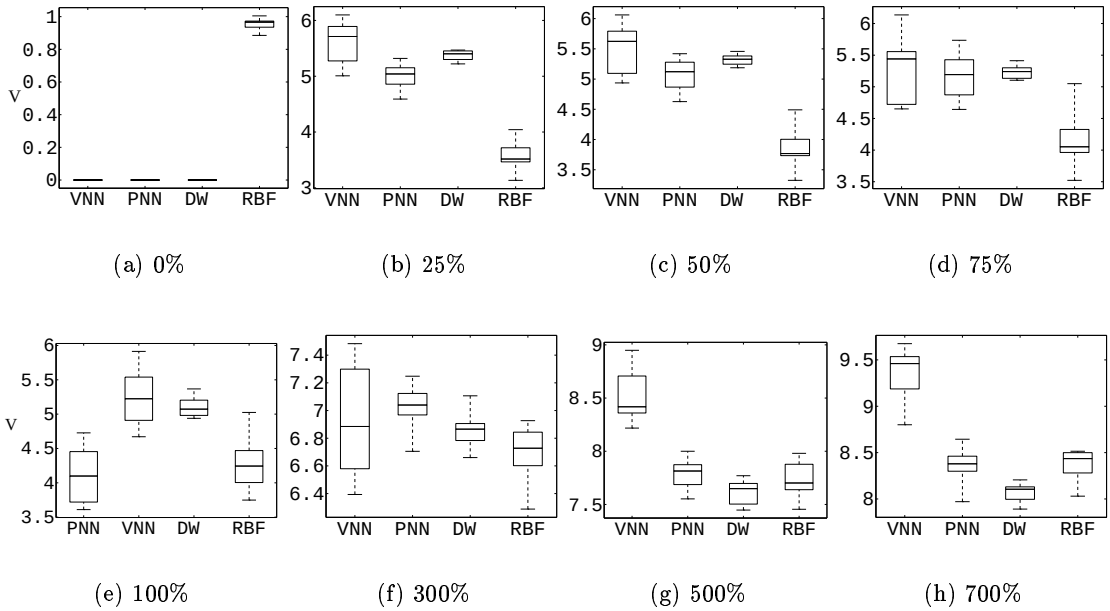

(a) 0% (b) 25% (c) 50% (d) 75%

(e) 100% (f) 300% (g) 500% (h) 700%

Figure 11: **Interpolation error V for the bladder examination.** See the caption of Figure 10 for an explanation of the box-and-whisker plots.

original B-scan and smoothed data in the gaps.

Since the radius R of the neighbourhood is tailored for each test of 100% and above (as explained in Section 3.4), the values of V for the thyroid examination are relatively small. However, for the 100% test on the bladder examination, the mean of V is considerably higher than the RBF method. This is because R is too large at the top of the image where the data is dense (but just large enough to fill the gaps at the bottom), resulting in excessive smoothing at the top. For 300% to 700%, the values of V improve, since the removal of data offsets the increase in R.

The images show progressive blurring as the percentage of removed data increases, but no reconstruction artifacts are generated. For the thyroid examination, the level of blurring is uniform. For the bladder examination, blurring is greater near the top of the images.

4.4 Radial Basis Function Interpolation

For both examinations, V is greater than zero for the 0% test, since the approximating function does not pass exactly through the original pixel values. For both examinations, however, the mean of V at 0% is less than one.

At 25%, 50% and 75%, the mean of V is considerably lower than the other methods and the resulting interpolated data appears the most detailed and least artificial. This is because the interpolation utilises both the remaining pixel data of the original B-scan and the neighbouring B-scans. The low values of V demonstrate the ability of a functional method to use the general shape of the underlying anatomical data to interpolate across the gaps.

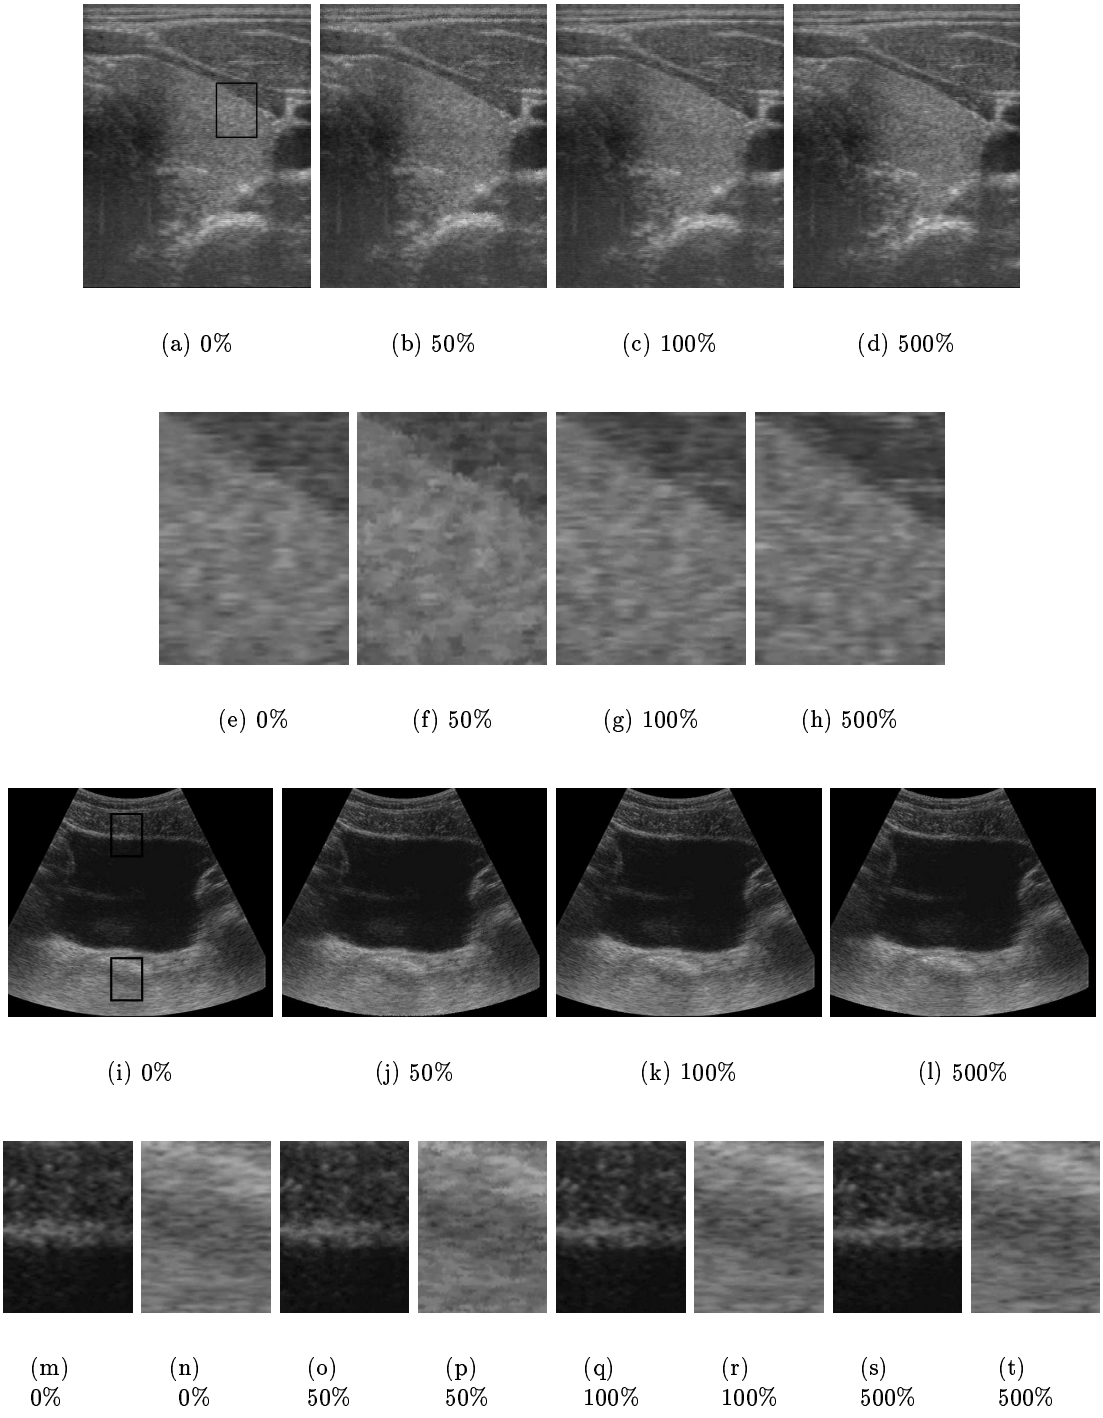

Figure 12: **Typical B-scans interpolated by the voxel nearest neighbour method**. All images are shown for a particular B-scan with various percentages of data removed. Images (a) to (d) are from the thyroid examination, with expanded views of the area indicated in image (a) shown in images (e) to (h). Images (i) to (l) are from the bladder examination, with expanded images of the top and bottom areas indicated in image (i) shown in images (m) to (t).

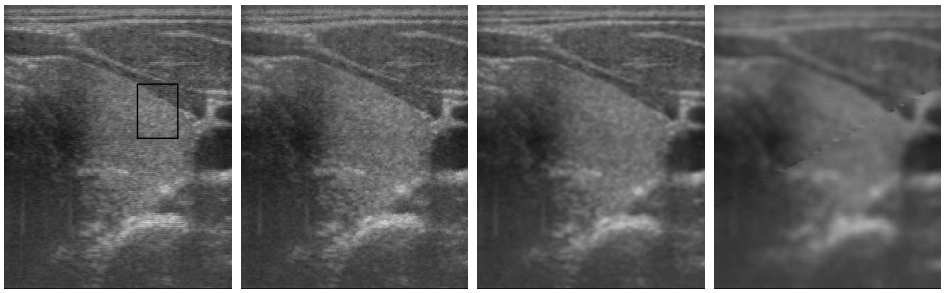

(a) 0% (b) 50% (c) 100% (d) 500%

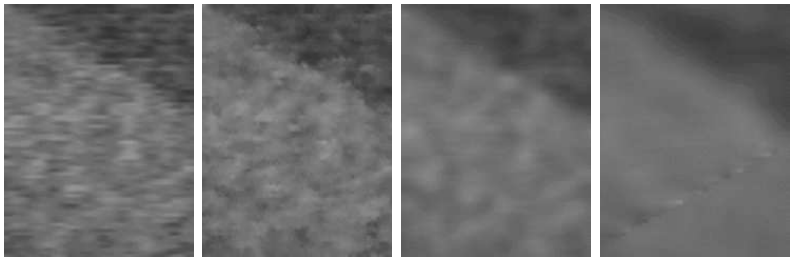

(e) 0% (f) 50% (g) 100% (h) 500%

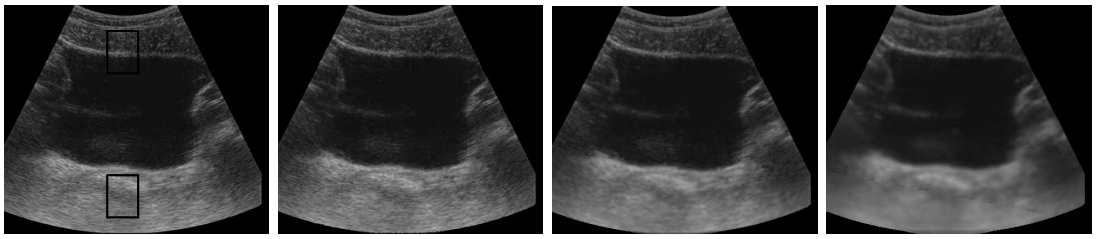

(i) 0% (j) 50% (k) 100% (l) 500%

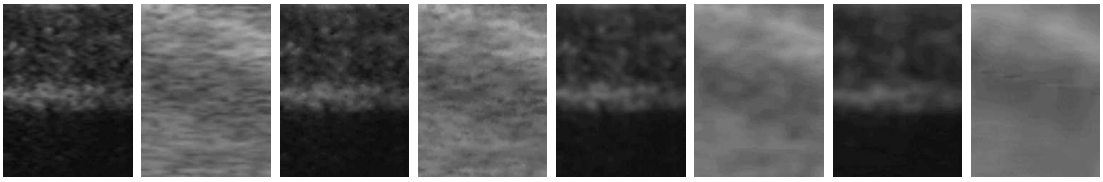

(m) 0% (n) 0% (o) 50% (p) 50% (q) 100% (r) 100% (s) 500% (t) 500%

Figure 13: **Typical B-scans interpolated by the pixel nearest neighbour method**. See the caption of Figure 12 for an explanation of the images.

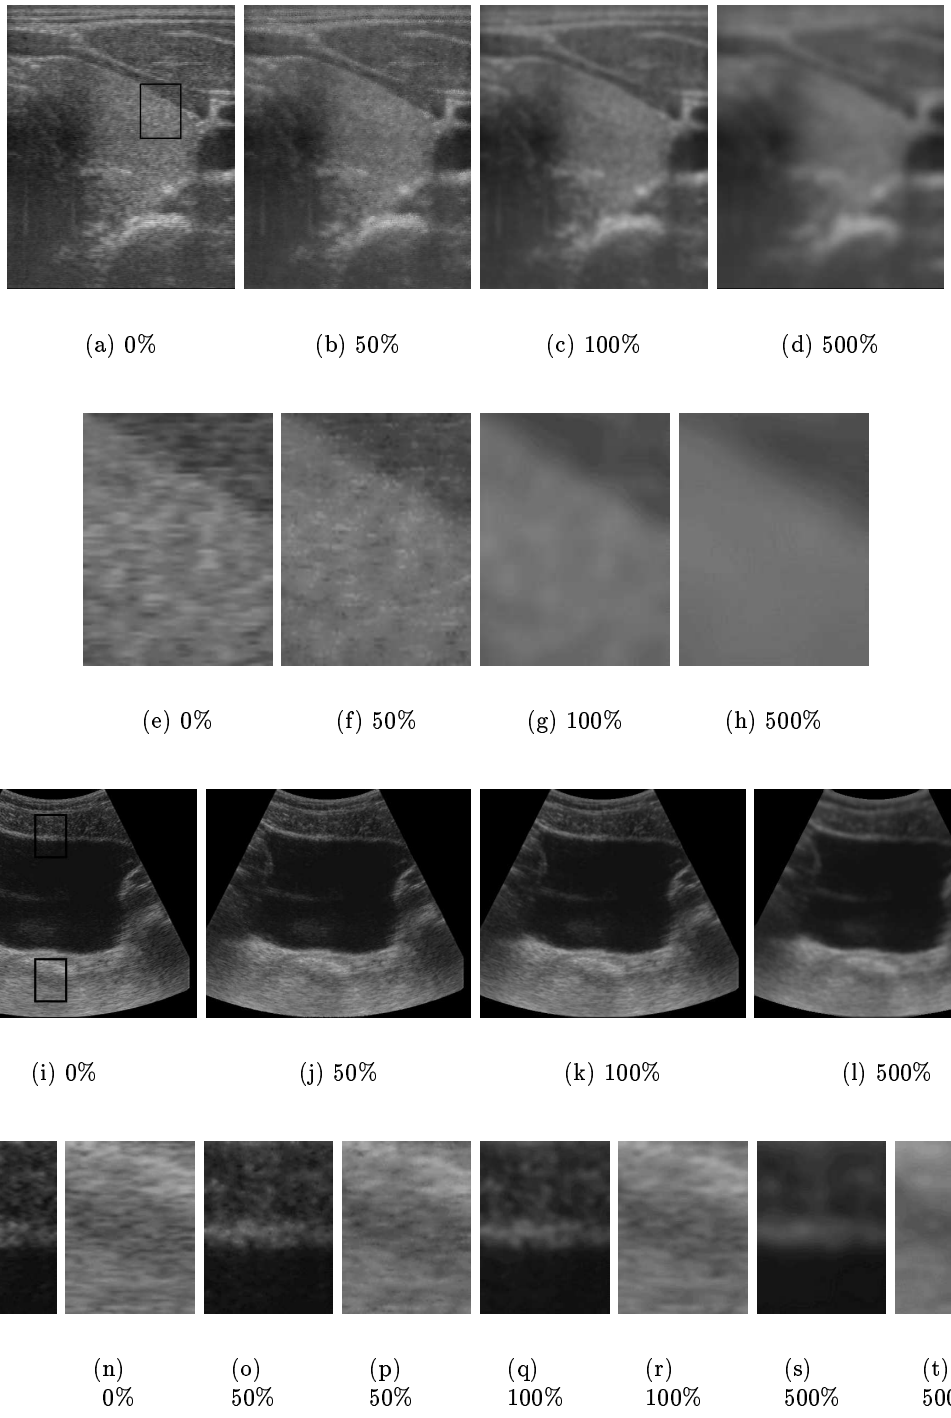

(a) 0% (b) 50% (c) 100% (d) 500%

(e) 0% (f) 50% (g) 100% (h) 500%

(i) 0% (j) 50% (k) 100% (l) 500%

(m) 0% (n) 0% (o) 50% (p) 50% (q) 100% (r) 100% (s) 500% (t) 500%

Figure 14: **Typical B-scans interpolated by the distance-weighted method**. See the caption of Figure 12 for an explanation of the images.

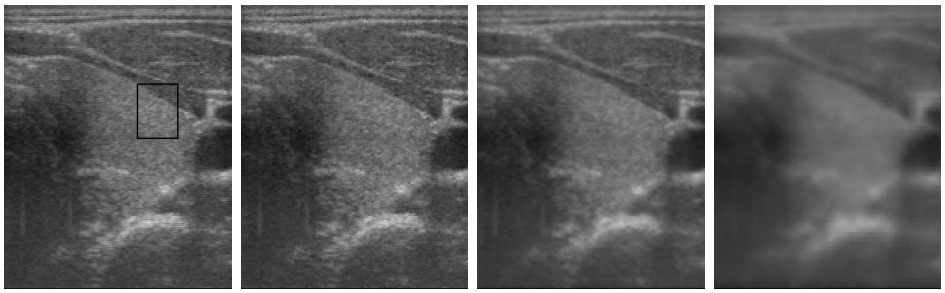

(a) 0% (b) 50% (c) 100% (d) 500%

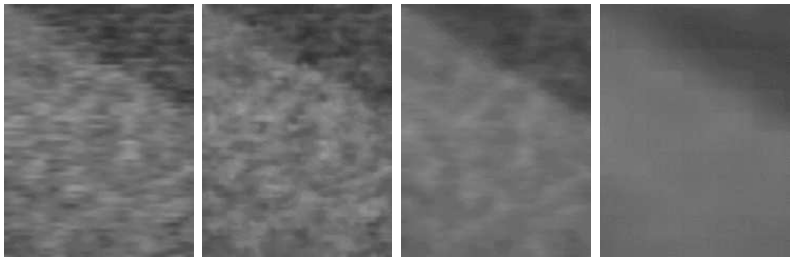

(e) 0% (f) 50% (g) 100% (h) 500%

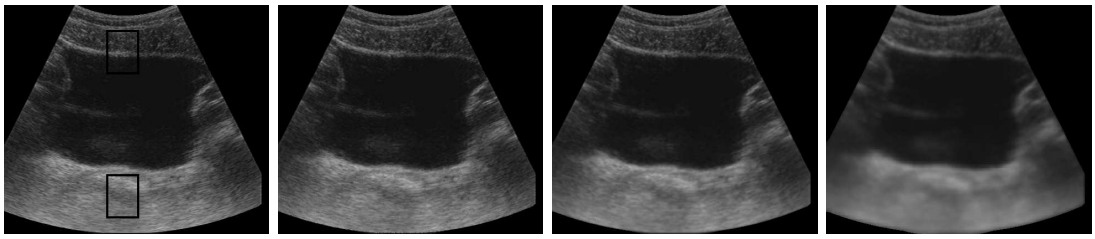

(i) 0% (j) 50% (k) 100% (l) 500%

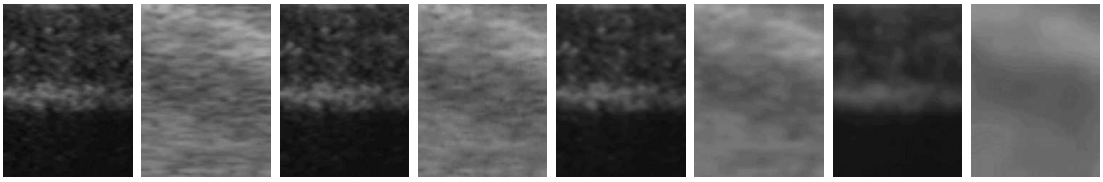

(m) 0% (n) 0% (o) 50% (p) 50% (q) 100% (r) 100% (s) 500% (t) 500%

Figure 15: **Typical B-scans interpolated by the radial basis function method**. See the caption of Figure 12 for an explanation of the images.

Yet at percentages of 100% and greater, the RBF is not always significantly better than the other methods. In particular, the VNN method matches it at 100% in the thyroid examination and at 100% and 300% in the bladder examination. The PNN method matches it at 300% in the thyroid examination and at 700% in the bladder examination. The DW method is even closer, matching the RBF method at 100%, 300% and 500% for the thyroid examination, and even slightly betters it at 700%. It also slightly betters the RBF method at 500% and 700% in the bladder examination.

One of the reasons the RBF loses some of its advantage is that the underlying shape of the anatomical data is lost when the gaps become too large. Another problem is that the RBF approaches the trend in the largest gaps. This is because the tension needs to be high enough so that no overshoots appear within the B-scan. Yet in the gaps, a higher tension results in a faster approach of the interpolating function towards the trend.

In general, however, the RBF method produces natural-looking images with no apparent artifacts. The boundaries between segments are not visible, since they are generally less significant than the grey level quantisation. No overshoots are generated either: the range of the interpolated data matches the range of grey levels in the original B-scans.

For the thyroid examination, the interpolated images become progressively blurred as the percentage of removed data increases. The level of blurring is uniform throughout the image. For the bladder examination, however, the denser data at the top of the images allows the RBF to retain a high level of detail. The blurring then progressively and smoothly increases toward the bottom of the images where the gaps are larger. The level of blurring therefore reflects the level of uncertainty in estimation of the underlying function. In comparison, the DW method shows, counter-intuitively, the greatest blurring near the top of the bladder images.

4.5 Discussion

The RBF method appears to work very well: low values of V are produced and no reconstruction artifacts are generated. It must be noted, however, that the computational demands are much greater than for the other three methods. The thyroid and bladder data sets can be reconstructed by the VNN, PNN and DW methods in a few minutes, but the RBF method requires a few hours. Yet because of segmentation, the RBF method is amenable to parallel processing. The segments can be interpolated in parallel because the memory storing the pixel values does not change (eliminating read errors). Each process also writes to a different segment in the voxel array (eliminating write errors). Since many modern ultrasound machines already have the capacity for parallel processing (the Toshiba Powervision 7000 we used for these examinations contains more than 60 Pentium processors), a practical implementation of the RBF method is not infeasible.

The greatest advantage of the RBF method occurs with small percentages of removed data. This is because the underlying shape of the data can be used to interpolate across small gaps and reduce the effects of outliers (caused by registration errors, for instance). This is particularly evident in the bladder examination, where the data is very dense and the shape of the data readily deducible from the remaining data points. The sub-100% data removal experiments are important in the context of one of the most popular visualisation techniques, any-plane slicing. Typically, the user wishes to view slices orthogonal or nearly orthogonal to the original B-scans. Thus, the slice plane is intersected by many of the B-scans and contains many true data points, as well as many gaps which need interpolating across. It appears that the RBF method performs especially well in this sort of situation.

The higher percentage tests are representative of what would happen with faster hand motion during the examination or a slower rate of acquisition by the computer. When the gaps become larger, the RBF method loses some of its advantage. The shape of the underlying data is eventually lost when the gaps become too large. Nevertheless, with few exceptions, the RBF method performs at least as well as the other methods. It is possible that the performance of the RBF method can be improved by choosing a different combination of tension ϕ and smoothing w. But any advantage at low percentages of data removal results in a trade-off at high percentages. Moreover, the improvements are small. The results are mainly unchanged as long as the tension ϕ is in the range of [10, 25] and the smoothing w is in the range [0.01, 0.1].

A much greater potential improvement of the RBF lies with the introduction of anisotropic tension. It is possible for the tension to be different in different directions. This is because the RBF functions are invariant to translations and rotations, but are not scale invariant [11]. This means a change in the scale of the dimensions is equivalent to a change in the tension parameter. By changing the scale of the three axes individually, the tension is changed for these three directions. The tension should ideally be high within the individual B-scans to avoid overshoots, and low in the direction orthogonal to the B-scans to fill the gaps between them. This will reduce the blurring in the gaps between B-scans that is evident in these tests.

Although the RBF method produces the lowest values of V, it must be admitted that the VNN method produces very detailed images at high percentages of removed data. Since the registration errors are small, the nearest neighbour projections appear to be seamless. Both examinations of the thyroid and bladder are of organs with a slowly varying shape, so the projections from different B-scans are similar. But if more complex shapes such as foetuses are scanned, or the registration errors increase, the boundaries will become more visible. Moreover, images interpolated across large gaps (which depict projections of distant B-scans) appear sharp, but are strictly incorrect since the anatomy is featured in the wrong place.

In summary, however, all four methods perform well on both examinations, and produce reasonable images when the gaps are kept at realistic levels. It is also true that all of the methods could be improved or tailored to a particular type of examination. Yet the point of this paper is not to declare a winner among the different reconstruction methods. Rather, it is to investigate typical reconstruction techniques on typical examinations and determine whether a functional interpolation method offers any advantages.

5 Conclusions and future work

The RBF reconstruction method is shown to perform at least as well as traditional reconstruction methods. The RBF method performs particularly well with small gaps, but loses some of its advantage when larger gaps are present. This is because the RBF method alone makes use of the underlying shape of the data to interpolate across the gaps. Yet with large gaps, this shape can no longer be followed and the RBF produces smoothed averages of the nearby B-scans.

The RBF method can reconstruct simultaneously areas with densely overlapping B-scans and areas with large gaps, while making a smooth transition between them. The resulting quality of the interpolated data is good, with no visible reconstruction artifacts produced in any of the tests. Unfortunately, the RBF method is currently the slowest reconstruction method and is not yet practical in its current form. Yet it is easily parallelised

and, given more efficient coding techniques, a practical implementation is feasible.

There is considerable scope for future development of the RBF method. First, the tension should be adjusted to accommodate the nature of the freehand ultrasound data. Since the pixel data all lie in planes, and are usually separated by gaps larger than the pixel size, the tension should be higher within the B-scan and lower perpendicular to the B-scan. It is also possible to introduce weights on the pixels [12] when, for example, estimates of registration error or signal fall-out are available. Finally, it is possible to find the optimal combination of tension and smoothing by statistical techniques such as cross-validation [12].

Since the RBF method alone uses a functional approach, it also offers a number of unique opportunities. For example, derivatives of any order can be computed directly from the interpolating function. This may be useful for visualisation, registration and segmentation. The functional representation may also be useful for data compression. For lower resolution voxel arrays, the functions can also be passed directly through low-pass filters for proper anti-aliasing. Since the function usually passes further away from outliers (when a non-zero smoothing parameter is used), the regions with high predictive error can be determined. This may be useful, for example, for investigations into registration errors.

In summary, this paper demonstrates that the RBF interpolation method holds considerable promise for reconstructing 3-D freehand ultrasound data. The RBF method is shown to be at least as good as the traditional reconstruction methods for two representative examinations. Many opportunities also exist to exploit the unique properties of the RBF method.

Acknowledgements

The original program code for the regularized splines with tension was kindly provided by Lubos Mitáš and Helena Mitášová. Lubos Mitáš is at the National Center for Supercomputing Applications at the University of Illinois. Helena Mitášová is at the Geographic Modeling Systems Laboratory at the University of Illinois. Development of the original spline code was carried out in cooperation with the United States Army Construction Engineering Research Laboratories. The freehand acquisition system was developed by Richard Prager and Patrick Gosling at the University of Cambridge. The 3-D renderings were produced using the 3DViewnix visualisation package. The many helpful discussions with Jonathon Carr are also appreciated. Robert Rohling is supported by Churchill College and an ORS award.

References

[1] M. Abramowitz and I.A. Stegun. *Handbook of Mathematical Functions.* Wiley, Dover, New York, 1965.

[2] C.D. Barry, C.P. Allott, N.W. John, P.M. Mellor, P.A. Arundel, D.S. Thomson, and J.C. Waterton. Three-dimensional freehand ultrasound: image reconstruction and volume analysis. *Ultrasound in Medicine and Biology*, 23(8):1209–1224, 1997.

[3] J. Carr. *Surface Reconstruction in 3D Medical Imaging.* Ph.D. Thesis, University of Canterbury, Christchurch, New Zealand, 1996.

[4] A. Fenster and D. B. Downey. 3-D ultrasound imaging — a review. *IEEE Engineering in Medicine and Biology Magazine*, 15(6):41–51, 1996.

[5] D. Fine, S. Perring, J. Herbetko, C. N. Hacking, J. S. Fleming, and K. C. Dewbury. Three-dimensional (3D) ultrasound imaging of the gallbladder and dilated biliary tree: reconstruction from real-time B-scans. *British Journal of Radiology*, 64:1056–1057, 1991.

[6] R. Franke. Scattered data interpolation: tests of some methods. *Mathematics of Computation*, 38(157):181–200, January 1982.

[7] R.V. Hogg and J. Ledolter. *Engineering Statistics*. Macmillan Publishing Company, New York, New York, USA, 1987.

[8] F. Hottier and A. Collet Billon. 3D echography: status and perspective. In K.H. Hohne, H. Fuchs, and S.M. Pizer, editors, *3D Imaging in Medicine: Algorithms, Systems, Applications*, pages 21–41. Springer-Verlag, Berlin, Germany, 1990.

[9] H.A. McCann, J.C. Sharp, T.M. Kinter, C.N. McEwan, C. Barillot, and J.F. Greenleaf. Multidimensional ultrasonic imaging for cardiology. *Proceedings of the IEEE*, 76(9):1063–1072, 1988.

[10] L. Mitáš and H. Mitášová. General variational approach to the interpolation problem. *Computers and Mathematics with Applications*, 16(12):983–992, 1988.

[11] H. Mitášová and L. Mitáš. Interpolation by regularized spline with tension: I. theory and implementation. *Mathematical Geology*, 25(6):641–655, 1993.

[12] H. Mitášová, L. Mitáš, W.M. Brown, D.P. Gerdes, I. Kosinovsky, and T. Baker. Modelling spatially and temporally distributed phenomena: new methods and tools for GRASS GIS. *International Journal of Geographical Information Systems*, 9(4):433–446, 1995.

[13] T.R. Nelson and D.H. Pretorius. Interactive acquisition, analysis and visualization of sonographic volume data. *International Journal of Imaging Systems and Technology*, 8:26–37, 1997.

[14] G.M. Nielson. Scattered data modeling. *IEEE Computer Graphics and Applications*, 13(1):60–70, January 1993.

[15] R. Ohbuchi, D. Chen, and H. Fuchs. Incremental volume reconstruction and rendering for 3D ultrasound imaging. *Proceedings of SPIE — The International Society for Optical Engineering*, 1808:312–323, 1992.

[16] R. W. Prager, A. H. Gee, and L. Berman. 3D ultrasound without voxels. In *Proc. Medical Image Understanding and Analysis*, Leeds, UK, 1998.

[17] R.W. Prager, R. Rohling, A. Gee, and L. Berman. Rapid calibration for 3-D free-hand ultrasound. *Ultrasound in Medicine and Biology*, 1998. In press.

[18] R. Rohling, A. Gee, and L. Berman. 3-D spatial compounding of ultrasound images. *Medical Image Analysis*, 1(3):177–193, April 1997.

[19] D Shepard. A two-dimensional interpolation function for irregularly spaced data. In *Proceedings of the 23rd National Conference of the ACM*, pages 517–523, 1968.

[20] A. Talmi and G. Gilat. Method for smooth approximation of data. *Journal of Computational Physics*, 23:93–123, 1977.

[21] J. W. Trobaugh, D. J. Trobaugh, and W. D. Richard. Three-dimensional imaging with stereotactic ultrasonography. *Computerized Medical Imaging and Graphics*, 18(5):315–323, 1994.