3D ULTRASOUND PROBE CALIBRATION WITHOUT A POSITION SENSOR

A. H. Gee, N. E. Houghton,
G. M. Treece and R. W. Prager

CUED/F-INFENG/TR 488

September 2004

University of Cambridge
Department of Engineering
Trumpington Street
Cambridge CB2 1PZ
United Kingdom

Email: ahg/neh27/gmt11/rwp @eng.cam.ac.uk

3D Ultrasound Probe Calibration Without a Position Sensor

Andrew Gee, Neil Houghton, Graham Treece and Richard Prager

University of Cambridge
Department of Engineering
Trumpington Street
Cambridge CB2 1PZ

Abstract

We present a technique for 3D ultrasound probe calibration which does not require measurements from the position sensor. The principle of operation is that the beam is aligned with a set of coplanar wires strung across a rigid frame. The probe and frame are mounted on a precision-manufactured mechanical instrument which allows adjustment and measurement of their relative pose. Semi-automatic image processing facilitates alignment of the beam and wires to within a tolerance of around 200 μm, despite the considerable beam thickness. The calibration process requires just a single view and relatively little user expertise. In a series of experiments with different ultrasound probes, we demonstrate the technique's high accuracy and precision. The latter is partly due to the elimination of the position sensor, a significant source of measurement noise, from the end-user calibration process.

1 Introduction

Probe calibration is an essential prerequisite for all 3D ultrasound applications where a position sensing device is attached to the probe [4, 5]. This includes freehand 3D ultrasound, where the sensor records the relative positions and orientations of the scan planes, allowing reconstruction of a 3D data volume from the individual B-scans and their positions. It also includes those applications of mechanically swept 3D probes where the data needs to be registered with some fixed, external coordinate system, for example for surgical navigation. In these applications, a position sensor is required to locate the probe with respect to the external coordinate system.

The most common position sensors in use today fall into two categories. First, there are the electromagnetic devices comprising a fixed transmitter and a small receiver. Pick-up coils in the receiver allow determination of its position and orientation with respect to the transmitter. Secondly, there are the optical systems comprising a number of fixed cameras and a target, typically an arrangement of infrared LEDs. The multiple views of the target allow its position and orientation to be deduced with respect to the fixed cameras. What all these systems have in common is that they report the position and orientation of the target or receiver attached to the probe: this is not the same as the position and orientation of the scan plane emanating from the probe face. Hence the need for probe calibration, which involves determining the geometrical relationship between the coordinate system of the target or receiver (henceforth referred to as the *sensor*) and that of the B-scan — see Figure 1.

Probe calibration is a critical step in the 3D ultrasound process. For surgical navigation applications, it has a direct impact on the ability to locate features accurately in 3D space. For 3D ultrasound reconstruction, errors in probe calibration translate into gross deformations of the reconstructed volume, affecting the measurement of distance, angle, area and volume. It is crucial to calibrate the probe as accurately as possible: even an error of a degree or so can cause points to be located several millimeters from their true positions.

There is now a significant literature on the topic of 3D ultrasound probe calibration. What follows is a representative, but far from exhaustive, survey of this literature. The usual approach is to scan some object (or *phantom*) with known geometrical properties. The simplest such object is a point [1, 3], typically the intersection of a pair of wires strung across a water bath. The point is scanned from many different directions and its location marked in each B-scan. The resulting

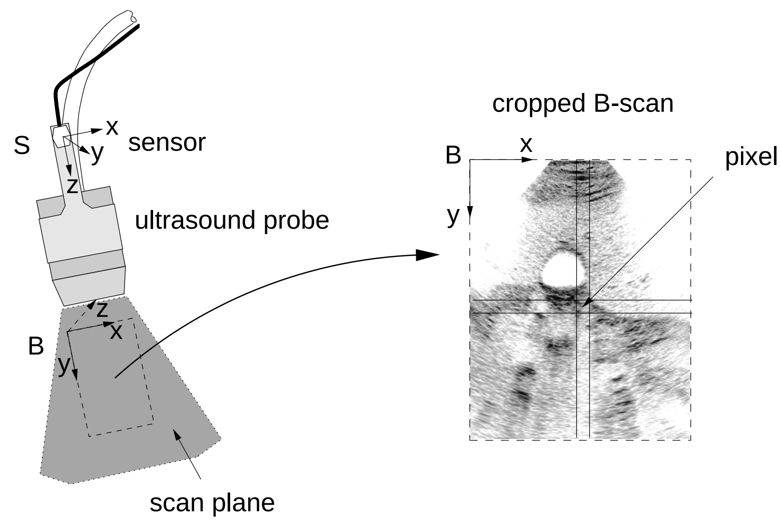

Figure 1: **Probe calibration**. Probe calibration involves finding the transformation $\mathbf{T}_{S \leftarrow B}$ between the B-scan coordinate system B and the sensor coordinate system S. 3D ultrasound systems typically crop the B-scan to a rectangular region of interest. The origin of B is at the top left corner of the cropped B-scan. The origin of S is at some arbitrary and unknown location inside the sensor casing. The calibration process may include estimation of the x and y dimensions of the pixels which make up the B-scan. Alternatively, and less accurately, if we assume that the pixels are square, the pixel dimension can be deduced from the nominal scan depth (in cm) and the height of the B-scan (in pixels).

set of points can be mapped into the sensor coordinate system using an assumed probe calibration, then into the world coordinate system using the position sensor readings. If the probe calibration is correct, the points will all map to the same location in world space. This observation can be expressed as a set of mathematical constraints on the elements of the calibration transformation, with each "view" of the phantom providing one or more equations. The set of equations is typically over-determined and therefore solved by least squares or iterative optimisation.

There are several difficulties with this simple approach. First, it is essential to scan the phantom from sufficient directions in a non-degenerate configuration, otherwise the set of equations is under-determined and no unique solution exists. Secondly, reliable segmentation of isolated points in ultrasound images requires manual intervention, making the calibration process slow and labour intensive. Thirdly, there is the problem of the ultrasound beam's considerable thickness in the elevational direction (ie. perpendicular to the scan plane): this can be as much as several centimeters. Consequently, a point may be visible in the B-scan even when the probe is not pointing directly at it. Unless it can be guaranteed that the point is at the elevational centre of the ultrasound beam, significant calibration errors are inevitable. Finally, there is the fact that the calibration technique utilises measurements from the position sensor, which are susceptible to noise of the order of 1.8 mm RMS for some electromagnetic devices. Noisy position sensor measurements result in inaccurate calibration.

Over the years, researchers have devised ingenious variations on the basic point-phantom theme, with the aim of overcoming the aforementioned difficulties. Lindseth et al. provide a useful taxonomy of the emerging techniques [7]. A common trend is to use not one but two position sensors, one attached to the probe, the other to the phantom. Knowledge of the phantom's position simplifies the calibration problem, reducing the number of equations, and hence views, required to solve it. An extreme example is the "2D alignment" paradigm [7, 11], which requires just a single view and is therefore immune to the problem of under-determined equations. However, it relies on the manual alignment of the scan plane with a set of coplanar phantom features, a procedure which is difficult to perform to any degree of precision: the beam thickness means that it is possible to move the probe a considerable distance before observing any change in the B-scan. 2D alignment techniques are highly sensitive to position sensor noise, since only one image is used, with no averaging across multiple views to reduce the effects of the noise.

The manual alignment problem is avoided by a class of techniques based on freehand scanning of wires arranged in N or Z patterns [2, 7, 8]. Given sufficient wires, the absolute location of the scan plane can be deduced from the positions of the wires in the B-scan (this locates the scan plane relative to the phantom) and the reading of the auxiliary position sensor (this locates the phantom in world space): thus, only one view is required to solve the calibration problem. With fewer wires, several views may be required. Beam thickness is still an issue, since the wires appear not as sharp dots but as smeared-out blobs: calibrating to the centre of the beam requires locating the precise centre of each blob. The effects of position sensor noise can be reduced by using more than the minimum number of views, but then segmentation may become troublesome, since the automatic detection of blobs in ultrasound images is never perfectly reliable.

In contrast, our own "Cambridge" phantom [10] appears as a clear, straight line in each and every B-scan. Images of lines contain much redundant information and are therefore particularly amenable to automatic segmentation. Only a single position sensor is required, and several hundred views can be used to suppress the effects of sensor noise. The phantom's design circumvents the beam thickness problem, but the scanning protocol [12] demands considerable user expertise to ensure that the calibration equations are not under-determined. In our experience supporting users of the Stradx freehand 3D ultrasound system [9, 12], incorrect calibration caused by inadequate scanning of the phantom is by far the most common source of error.

In this paper, we present an improved calibration technique which combines the strengths of the previous approaches while avoiding their weaknesses. It is a 2D alignment technique, requiring just a single view and no operator expertise to select appropriate scanning directions. The phantom's design eliminates the beam thickness problem to within a tolerance of around 200 μm. Uniquely, the position sensor plays no part in the end-user calibration process: hence the apparently paradoxical title of this paper. The consequential elimination of sensor noise has a significant effect on the calibration precision.

The paper is organised as follows. In Section 2, we present the phantom and describe its construction; explain how careful metrology enabled accurate estimation of its critical dimensions; and outline the semi-automatic algorithms used to segment the phantom's features in the B-scans. In Section 3, we describe a set of experiments which demonstrate the new technique's superior accuracy, precision and ease of use. Finally, in Section 4 we draw some conclusions and suggest further work to enhance the phantom's utility in everyday clinical scenarios.

2 The calibration phantom

2.1 Principle of operation

Figure 2 shows the calibration instrument. At the highest level of abstraction, there is a planar phantom (the set of wires in Figure 2(b)) with which the ultrasound beam is to be aligned. The rest of the instrument comprises a water tank inside which the phantom is immersed, a gantry on which the ultrasound probe is mounted, and a set of adjustable stages to allow precise alignment of the phantom and the ultrasound beam.

The alignment of two planes requires adjustments across three degrees of freedom, two angular and one linear. Thus the water tank sits on top of two angular stages, the upper one allowing rotation about the x_w axis[1] (pitch), the lower one rotation about the y axis (yaw). Immediately below these is a linear stage, allowing translation in the z (in-out) direction. Below this is one further linear stage, allowing translation in the x (left-right) direction. This is not required to align the planes, but is provided to ensure that certain reflectors attached to the wires (which will

[1]The x_w axis is aligned with the wires in Figure 2(b). It is not the same as the instrument's main x axis. The latter is fixed, while the former's orientation depends on the yaw adjustment.

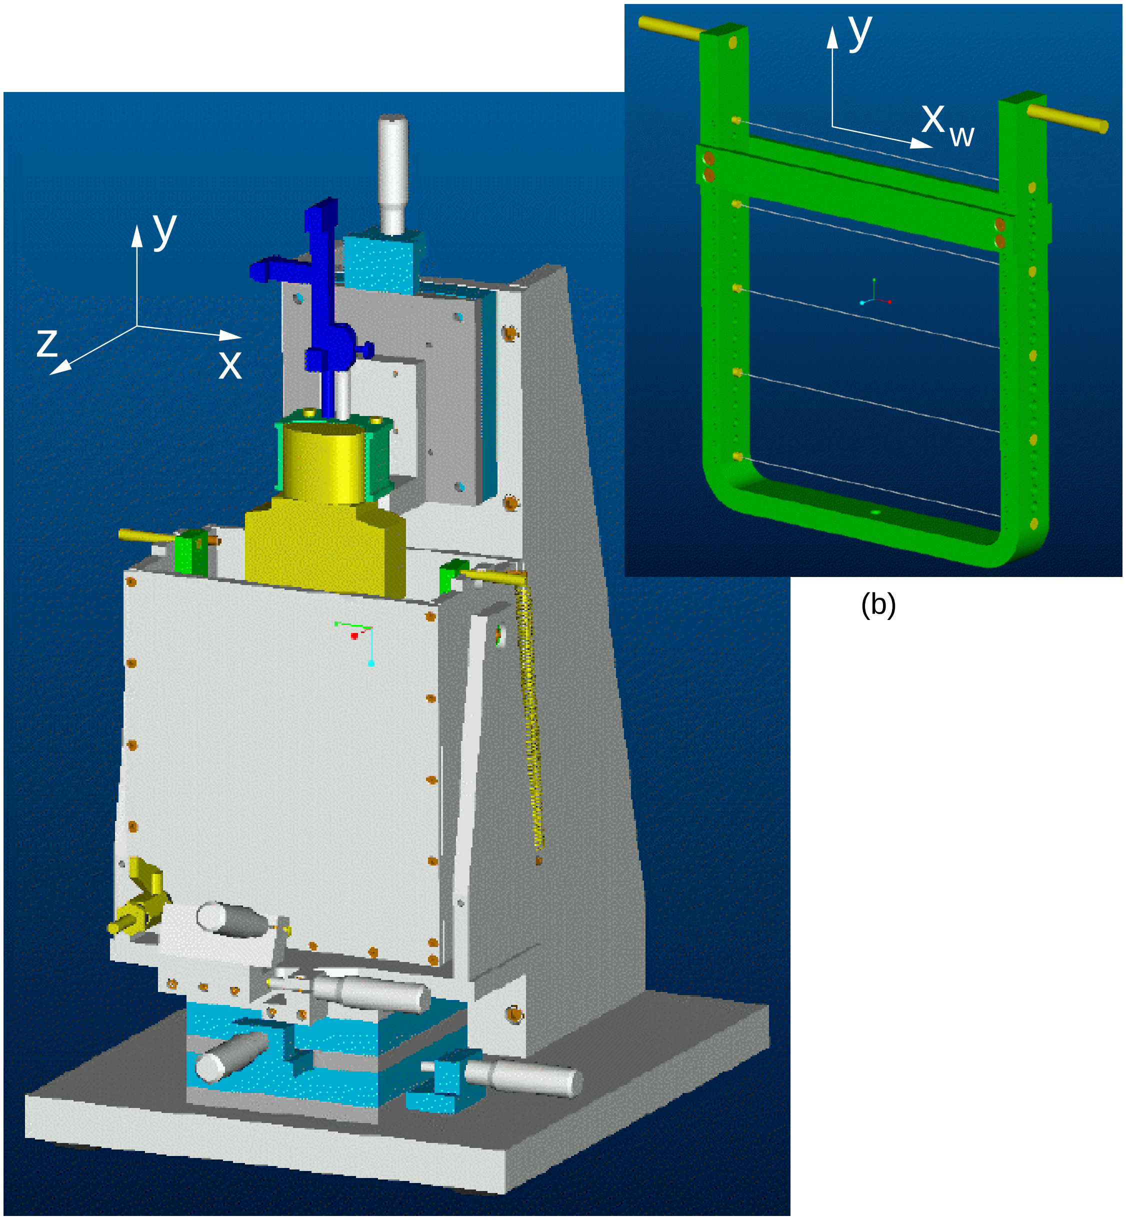

(a)

(b)

Figure 2: **The calibration instrument**. (a) shows a general view of the instrument. The probe and sensor are mounted on a gantry, whose position can be adjusted in the vertical (y) direction. The probe face is immersed in a 20 cm deep water tank, which sits on top of four adjustable stages. Two of the stages are linear, allowing adjustment of the tank's position in the x (left-right) and z (in-out) directions. The other two are angular, allowing rotation of the tank about the y axis (yaw) and the x_w axis (pitch). A removable frame (b) is mounted rigidly inside the tank. A set of coplanar wires are stretched across the frame, forming the plane with which the ultrasound beam is aligned. The various adjustments allow precise alignment of the wires and the beam.

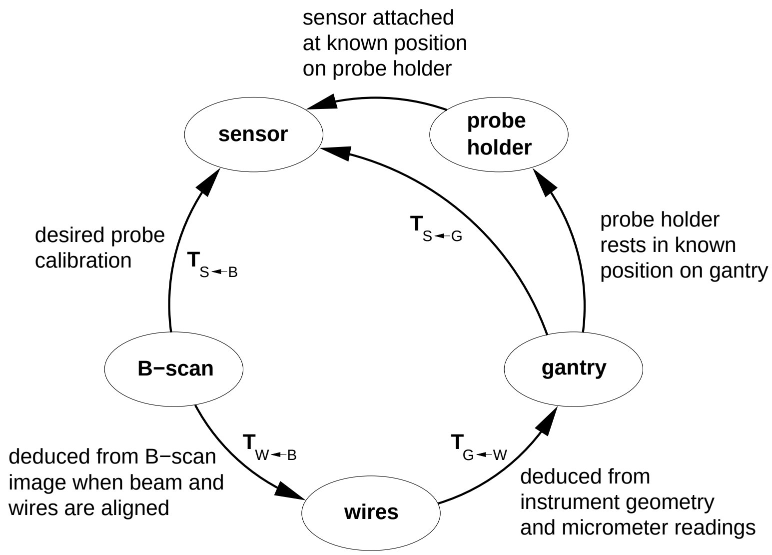

Figure 3: **The calibration process**. The figure illustrates how the probe calibration $\mathbf{T}_{S \leftarrow B}$ is the concatenation of three known transformations, as given in equation (1).

be described shortly) can be moved to directly below the probe face. For the same reason, the position of the gantry can be adjusted in the y (up-down) direction.

The instrument is manufactured to high precision, such that the rigid body transformation between any two of its components can be deduced from the readings of the micrometers that drive the five stages. Furthermore, the probe holder is designed such that whenever the probe is replaced on the gantry, the sensor ends up in precisely the same position relative to the gantry: the rigid body transformation $\mathbf{T}_{S \leftarrow G}$ between the gantry and sensor is thus known.

To calibrate the probe, the ultrasound beam is first aligned with the plane of the wires by adjusting the micrometers until the B-scan image indicates precise alignment. The micrometer readings are then used to calculate the rigid body transformation $\mathbf{T}_{G \leftarrow W}$ between the wires and the gantry. If we also know the transformation $\mathbf{T}_{W \leftarrow B}$ between the cropped B-scan and the wires (this can be deduced from the B-scan image itself, as described in Section 2.4), we can calculate the probe calibration $\mathbf{T}_{S \leftarrow B}$ by direct concatenation of the various components:

$$\mathbf{T}_{S \leftarrow B} = \mathbf{T}_{S \leftarrow G} \mathbf{T}_{G \leftarrow W} \mathbf{T}_{W \leftarrow B} \tag{1}$$

This equation, which is the crux of the calibration process, is illustrated in Figure 3.

We have thus far glossed over the details of how the ultrasound beam can be precisely aligned with the plane of the wires, despite the significant beam thickness. We achieve this by mounting pairs of wedges on the wires, as illustrated in Figure 4. We have previously described the use of wedges in the context of our earlier "Cambridge" calibration phantom [10]. The idea is that the wedges produce symmetrical reflections only when centred in the beam. It is thus possible to gauge the local alignment of the beam and wires in the vicinity of the wedges. The planes' *global* alignment can be determined from three pairs of wedges mounted at suitable points in the B-scan, as shown in Figure 4(b).

Several of the phantom's design features simplify the alignment procedure significantly. A simple jig is used to mount the top pair of wedges at the precise centre of the top wire[2]. The yaw and pitch stages are arranged such that the two axes of rotation intersect at this point. Thus,

[2]The jig comprises two identical lengths of plastic tube which slot over the wires: one sits on the left of the wedges, the other on the right. The length of the tubes is such that there is a marginal interference fit between the tubes, wedges and the sides of the frame. Thus, by fitting the jig while the top wedge is mounted on the wires, it is guaranteed that the wedge is at the centre of the wires. The jig is then removed.

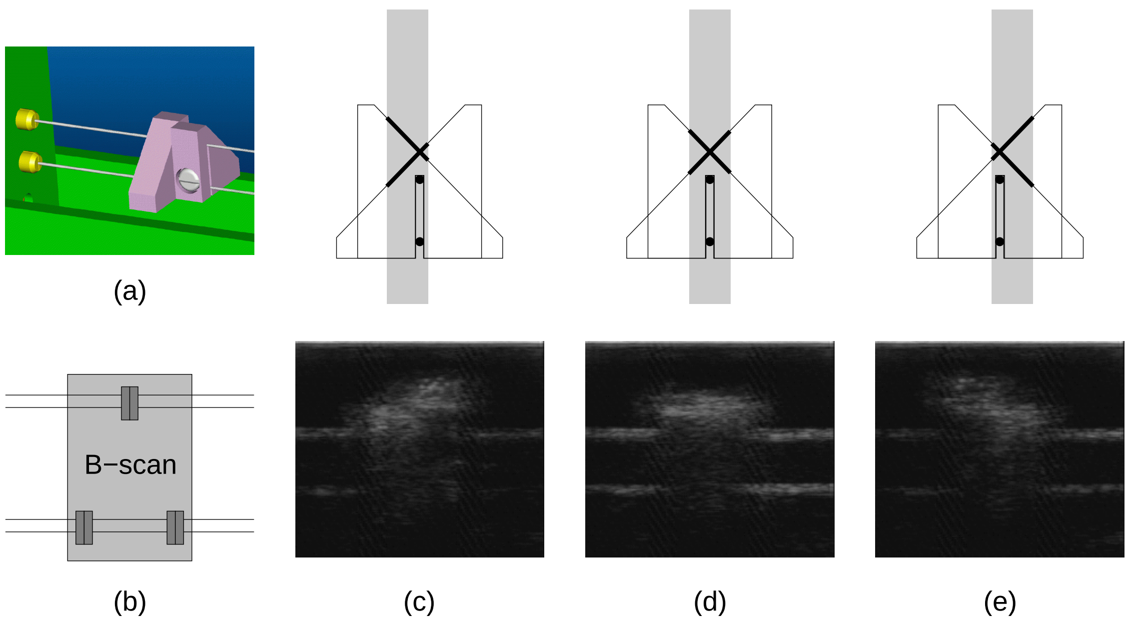

Figure 4: **Wedges**. The ultrasound beam can be aligned with the plane of the wires, despite the significant beam thickness, by mounting pairs of wedges on the wires, as shown in (a). In (c), the grey box shows the ultrasound beam, which in this case is not properly centred on the wires. The parts of the wedges which lie in the beam are highlighted in bold. The reflection from the right wedge is thus higher than that from the left wedge, producing an asymmetric image in the B-scan. In (d), the beam is centred on the wires and the image is accordingly symmetric. In (e), the beam has moved to the other side of the wires, so the left wedge appears higher than the right wedge. Note that the image of the wires is not a good indicator of alignment: it barely changes from (c) to (e). A single pair of wedges indicates only *local* alignment of the two planes. By arranging three pairs of wedges on the wires, as shown in (b), the global alignment can be verified.

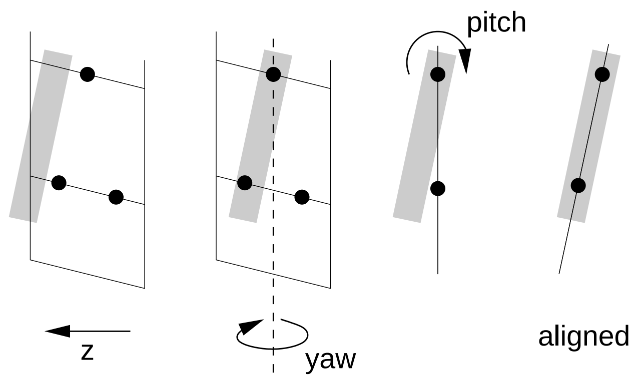

Figure 5: **The alignment procedure**. The ultrasound beam is shown end on, as a grey box, in all four views. The wedge-pairs are the three solid circles. The frame is first translated in the z direction, until the top wedge-pair is aligned with the beam. Next, the frame is rotated about the yaw (y) axis until the bottom two pairs show the same degree of misalignment. Finally, the frame is rotated about the pitch (x_w) axis until the bottom wedges are aligned with the beam.

adjusting the yaw and pitch has no effect on the position of the top wedge. The alignment is then achieved by adjusting the z stage until the top wedge is centred in the beam, then the yaw stage until the bottom two wedges show the same degree of misalignment, then the pitch stage until the bottom two wedges are centred in the beam — see Figure 5. Little or no iteration is required.

2.2 Mechanical design details

The calibration instrument comprises a base plate, on which are mounted two assemblies: an anodised aluminium tower, which supports the y stage and thus the probe gantry; and a stack of four stages (x, z, yaw and pitch, in that order), which supports the water tank. The stages are off-the-shelf components, supplied with micrometers that can be read to a precision of 10 μm. All other major components were manufactured from anodised aluminium in the workshops of Cambridge University's Engineering Department.

The angular stages operate as follows. Consider the pitch axis in Figure 2(a): it runs between the two points where the tank is suspended in its cradle. Note the spring running down the right side of the tank. There is a second, matching spring on the left side. The springs dispose the tank to rotate about the pitch axis such that its bottom moves out of the page, towards the reader. This motion is opposed by the pitch micrometer (the highest one on the stack), which is in contact with a small sphere attached to the front face of the tank. The angle of rotation is thus deduced by dividing the micrometer reading by the measured distance between the pitch axis and the contact point. The yaw stage is also spring loaded and works in a similar manner.

The tank is sealed on the inside with silicone and includes a release valve (on the bottom left of the tank, as seen in Figure 2(a)) for easy drainage. The frame (Figure 2(b)) is braced with two aluminium bars to resist the tension in the wires. Each side of the frame is drilled with 21 holes at a pitch of 7.5 mm: wires can be strung between any pair of holes. The wires are easily tensioned by pulling on them while simultaneously tightening a small brass clamp which fits inside the holes: two of these clamps can be seen in Figure 4(a). The frame can be lifted out of the tank and replaced with another, perhaps with a different arrangement of wires and wedges to suit a different ultrasound probe. A semi-kinematic mount guarantees highly repeatable positioning of the frame within the tank. There are just three contact points: two dowels at the top of the frame, which rest in V-shaped notches cut into the top of the tank, and a hole drilled in the bottom of the frame, which engages a pin protruding from the bottom of the tank.

The remaining important design details concern the gantry and the probe holder. The concavity on the inside of the probe holder is machined to match the curvature of the probe handle.

The face of the concavity is lined with a self-adhesive foam strip, the probe is pressed onto the strip and secured tightly with two cable ties, as shown in Figure 11(b). This design allows very rigid mounting of the probe in its holder, while being completely reversible: just snip the cable ties and work the probe away from the adhesive strip. This is an important consideration given the terms of manufacturers' guarantees. There are two holes drilled into the probe holder. These receive two pins mounted vertically on the gantry's horizontal base plate. Both pins are visible in Figure 2(a), protruding slightly out of the top of the probe holder. One of the pins is visible in Figure 11(b). The holes in the probe holder are necessarily a tight fit around the pins, but one of the holes is circular and the other is slotted, and both have a short taper at the bottom. These details allow the probe holder to be dropped into place with little effort. As with the frame, the positioning of the probe holder is highly repeatable: its bottom face rests on the gantry's base plate, while the pins constrain the remaining degrees of freedom.

Protruding from the top of the probe holder in Figure 2(a) is a pin, on which is mounted the position sensor, in this case an AdapTrax[3] infrared LED target for the Polaris[4] optical tracking system. The pin has a neck to constrain the vertical position of the sensor, and a flat to receive the sensor's mounting screw. The flat constrains the twist of the sensor about the pin. Thus the sensor can be attached to the probe holder, in the same position, to a high degree of precision. Note how the sensor and the probe are mounted independently on the probe holder. This means that, even if the probe is removed from the holder and replaced, the rigid body transformation between the sensor and the gantry is unchanged. Furthermore, we have manufactured three holders for different probes. While they necessarily differ in the design of the concavity which locates around the probe handle, they are identical when it comes to the positioning of the sensor pin with respect to the two holes. The gantry to sensor transformation $\mathbf{T}_{S\leftarrow G}$ is therefore constant.

2.3 Metrology

Although the phantom was manufactured to close tolerances, a set of measurements was made with a Mitutoyo[5] coordinate measuring machine (CMM) to remove any residual geometric errors. First, the x, y and z adjustments were set to nominal values, midway along each micrometer's range of motion. Then, with reference to Figure 6, the purpose of the CMM measurements was to establish:

- The yaw and pitch adjustments which align the z axes of the wire and gantry coordinate systems.
- For these adjustments, the residual rigid body transformation $\mathbf{T}^*_{G\leftarrow W}$ between the wire and gantry systems. $\mathbf{T}^*_{G\leftarrow W}$ comprises three translational components and a single rotational component (roll about the z axis).

With this information, and knowing the positions and orientations of the various adjustment axes, we can readily calculate the rigid body transformation $\mathbf{T}_{G\leftarrow W}$ for any set of micrometer readings, not just the nominal settings. This is required for the calibration process (1).

The results of the CMM measurements are summarised in Table 1. The bottom row of the table indicates the uncertainty in the measurements. The possible causes of error include the CMM itself, manufacturing errors in the wire frame[6] and squareness errors in the mounting of the various stages (these come into play only when the stages are adjusted away from their nominal settings). The contribution of each source of error was estimated using appropriate *a priori* knowledge, such as the manufacturer's specification of the CMM and the precision of the machines used to manufacture the frame. The errors were assumed to be independent and consequently combined in a root-sum-of-squares sense to obtain the values in Table 1.

A further set of measurements was made to estimate the rigid body transformation $\mathbf{T}_{S\leftarrow G}$ between the gantry and the position sensor, which is also central to the calibration process (1).

[3] Traxtal Technologies, `http://www.traxtal.com/`

[4] Northern Digital Inc., `http://www.ndigital.com/`

[5] Mitutoyo Corporation, `http://www.mitutoyo.co.jp/global/`

[6] The measurements were made on the tank, with the wire frame removed. This was to ensure that they were not biased by any play in the frame mount.

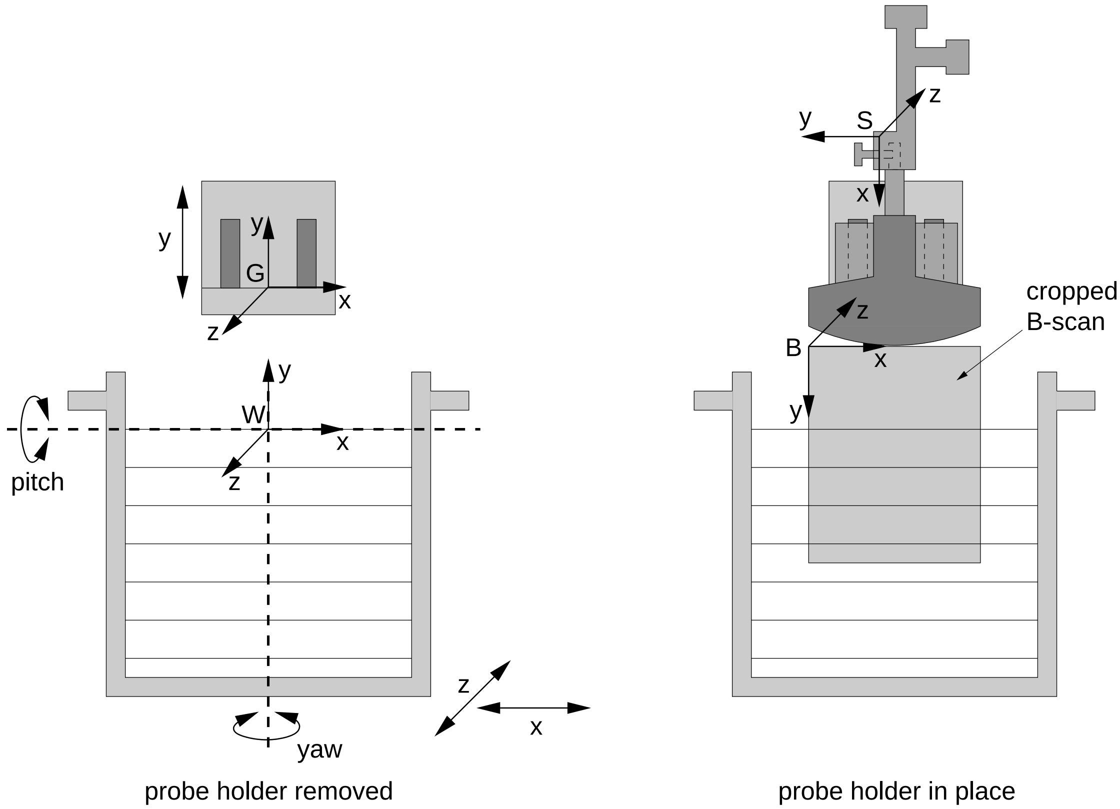

Figure 6: **Coordinate systems**. The calibration process depends on the relative pose of four coordinate systems. The origins and orientations of the wire, gantry and B-scan coordinate systems are arbitrary: we have defined them as illustrated in the figure. The wire coordinate system W has its origin at the centre of the top wire, where the yaw and pitch adjustment axes intersect. The gantry coordinate system G has its origin on the horizontal base plate, midway between the two pin centres. The B-scan coordinate system B has its origin at the top left corner of the B-scan. The orientation and origin of the sensor coordinate system S are defined by the manufacturer. The five adjustment axes are shown as double-headed arrows in the left hand figure.

	x	y	z	yaw	pitch	roll
micrometers (mm)	11.00	11.00	11.00	11.64	13.10	–
$\mathbf{T}^*_{G\leftarrow W}$ wrt G	-0.11 mm	-68.45 mm	29.90 mm	0.00°	0.00°	0.53°
uncertainty in $\mathbf{T}_{G\leftarrow W}$	± 123 μm	± 158 μm	± 158 μm	$\pm 0.052°$	$\pm 0.044°$	$\pm 0.049°$

Table 1: **Rigid body transformation between the gantry and wire coordinate systems**. The top row shows the nominal micrometer settings which align the two z axes. For these nominal settings, the middle row shows the residual rigid body transformation between the coordinate systems. The bottom row indicates the uncertainty in these values.

	x	y	z	azimuth	elevation	roll
mean $\mathbf{T}_{S\leftarrow G}$ wrt G	-17.30 mm	74.51 mm	1.25 mm	$-89.79°$	0.51°	177.91°
uncertainty in $\mathbf{T}_{S\leftarrow G}$	±0.90 mm	±0.71 mm	±0.65 mm	$\pm0.63°$	$\pm0.61°$	$\pm1.20°$

Table 2: **Rigid body transformation between the sensor and gantry coordinate systems**. The mean was calculated from ten measurements. The tabulated uncertainty is 1.96 times the standard deviation of the ten measurements, ie. 95% confidence limits assuming the measurements to be Normally distributed.

This is the only point at which measurements are required from the position sensor, so we must bear in mind the significant uncertainty in its readings (around 350 μm RMS for the Northern Digital Polaris, compared with around 10 μm for the Mitutoyo CMM). The AdapTrax LED target was mounted on the probe holder, the holder was dropped onto the gantry and the whole assembly placed in the Polaris cameras' field of view. A second LED target was attached to a pointing device. In order to record the location of the pointer's tip, we needed to know the offset between the tip and the target. We performed this calibration in the usual way [6], by rotating the pointer with its tip held stationary against a rough surface. We then used the pointer to record the positions of various points on the gantry. The AdapTrax's position and orientation were used to refer all the pointer readings to the sensor coordinate system S.

First, we recorded three points around the base of each of the pins mounted on the gantry's base plate — see Figure 6. We located the pin centres at the corresponding circumcenters. We also recorded four points on the base plate, one near each corner, and used these to find a best fit plane and its normal. The pin centres were projected onto this plane: their midpoint defines the origin of the gantry coordinate system G. The line joining the pin centres defines the x axis, while the base plate normal defines the y axis: this fixes the orientation of the coordinate system. Recall that all pointer measurements are in the sensor coordinate system S, so we now know the complete transformation between the G and S systems. The entire set of measurements was repeated ten times. The results are summarised in Table 2.

The use of the position sensor has introduced considerably more uncertainty compared with the measurements in Table 1. If we were to use just a single measurement of $\mathbf{T}_{S\leftarrow G}$, we must expect errors according to the bottom row of Table 2. However, assuming there is no bias in the individual $\mathbf{T}_{S\leftarrow G}$ measurements[7], we can reduce the uncertainty as much as we wish by simple averaging. Thus, if we use the mean of the ten measurements, as in the top row of Table 2, the uncertainty is reduced by a factor of $\sqrt{10}$. If we were to repeat the measurements 100 times (this would take about a day), we could reduce the uncertainty by a factor of ten, bringing it more into line with the values in Table 1.

2.4 Image segmentation

Referring back to the key calibration equation (1), we have now established how suitable metrology can provide good estimates of $\mathbf{T}_{S\leftarrow G}$ and $\mathbf{T}_{G\leftarrow W}$. That leaves just the wire to B-scan transformation, $\mathbf{T}_{W\leftarrow B}$. With the wire and B-scan planes aligned, $\mathbf{T}_{W\leftarrow B}$ has just one rotational (roll about the z axis) and two translational (x and y) degrees of freedom. There is an additional half degree of freedom, a possible left-right flip: the wedges appearing at the bottom left corner of the B-scan might in fact be at the bottom right of the wire frame, depending on which way round the probe is put in its holder. All 3.5 degrees of freedom, along with the pixel dimensions, can be estimated from the B-scan image — see Figure 7.

Recall that the top wedge-pair is positioned at the precise centre of the top wire: this is the origin of the wire coordinate system W. By locating this point in the B-scan coordinate system B, as shown in Figure 7, we can estimate the translation between W and B. The apparent angle

[7] To this end, the pointer was recalibrated between each set of measurements, and the phantom moved to a different position within the Polaris cameras' field of view.

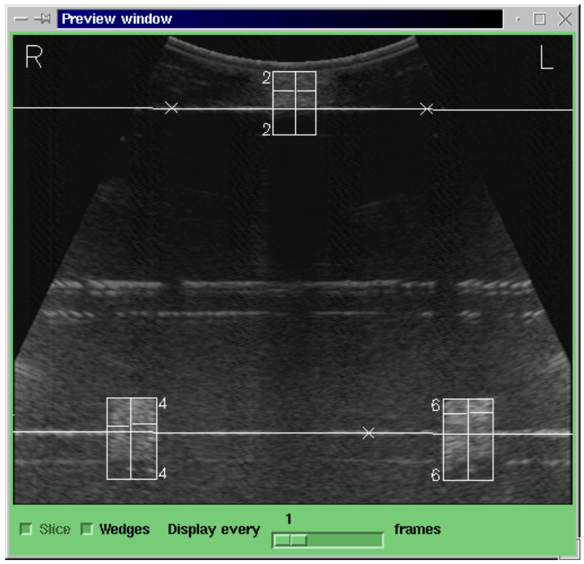

Figure 7: **Image segmentation**. Semi-automatic segmentation algorithms locate the three wedge-pairs, top wire and bottom wire in the B-scan. The intersection of the top wire with the central axis of the top wedge-pair reveals the origin of the wire coordinate system W relative to the B-scan coordinate system B. The angle of the wires to the horizontal determines the roll between W and B. The distance between the two wires in B, along with knowledge of their physical separation in W, provides an estimate of the y pixel dimension (mm/pixel). The x pixel dimension can be deduced from the spacing of the two lower wedge-pairs, or alternatively the pixels can be assumed to be square. A left-right flip (as indicated by the 'R' and 'L' at the top of the figure) is easily spotted by sliding a finger along the probe face while observing the B-scan.

of the wires to the horizontal determines the roll between the two systems. The vertical pixel dimension can be deduced from the distance between the two highlighted wires in B, along with knowledge of their actual separation in W. In this study, we assume square pixels: alternatively, the horizontal pixel dimension can be estimated from the separation of the two lower wedge-pairs.

A key component in Figure 7 is a semi-automatic algorithm to locate and characterise the three wedge-pairs. The user initiates the process by drawing a rectangular bounding box around each pair. The automatic part of the algorithm then determines the wedges' central axis and two indicators of asymmetry. Recall from Figure 4 that the ultrasound and wire planes are aligned when each wedge-pair produces a symmetrical reflection. Although misalignments of around 500 μm are detectable by eye, the automatic algorithm increases the sensitivity to around 200 μm and avoids any inter- and intra-operator variability. Details of the algorithm can be found in Figure 8 and its caption.

Figures 9 and 10 show examples of the wedge segmentation algorithm in action. Naturally, its performance depends on the quality of the incoming image data. We have therefore chosen to illustrate its capabilities with some of the poorest quality probes at our disposal, namely a 7 MHz linear array probe and a 3.5 MHz convex array probe, attached to a Toshiba[8] model SSA-270A/HG ultrasound scanner. This is a two generations old model: the probes, having seen many years of active clinical service followed by much abuse in our laboratory, have dead crystals with corresponding drop-outs in the B-scans (one of them passing directly through the centre of the top left wedge in Figure 9). Despite this, the wedge segmentation algorithm is quite capable of detecting misalignments of the order of 200 μm. Of the two indicators described in Figure 8, the left and right peak finder is preferred in practice, since it shows not only the degree of misalignment

[8] Toshiba Corporation, `http://www.toshiba.com/`

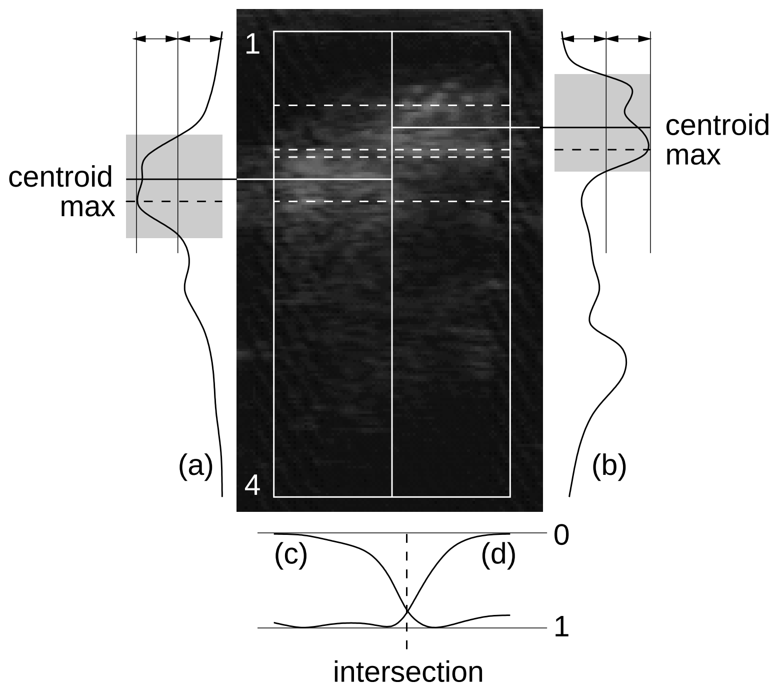

Figure 8: **Wedge segmentation**. The user clicks and drags to define an approximate bounding box around the wedge-pair. Consider the left half of this box. By summing the pixel intensities along horizontal rows, we obtain the vertical intensity profile (a). This has a broad peak around the reflection of the left wedge, though note that the profile's maximum is subject to noise and happens to lie towards the bottom of the peak. We therefore define a region of interest (the shaded rectangle) starting at the maximum and extending in both directions until the profile drops to the mean of its maximum and minimum values. We use the centroid of the profile within this region as our estimate of the left wedge's image location: this is the solid horizontal line in the left half of the bounding box. By repeating this analysis for the right wedge (b), we obtain a pair of horizontal lines whose relative vertical displacement is indicative of the local ultrasound beam misalignment. A second misalignment indicator is calculated by comparing the profiles (a) and (b) along their length, summing the absolute differences. The result is displayed as a number at the bottom of the bounding box: the number at the top is the smallest result discovered so far. Small numbers indicate a high degree of symmetry between the left and right profiles. Now consider a narrow region of interest around each centroid (the dotted horizontal lines). By summing pixel intensities along vertical columns within each region, we obtain the horizontal intensity profiles (c) and (d). The profiles are normalised by shifting and scaling so that they span the range 0 to 1. Their point of intersection is then a good indicator of the wedge-pair's central axis. To suppress the effects of noise, the (c) and (d) profiles are averaged across 200 images, 100 acquired with the right wedge appearing above the left one (as above), 100 with the left above the right. The user has to adjust the phantom's z stage to acquire these 200 images: this process takes about ten seconds.

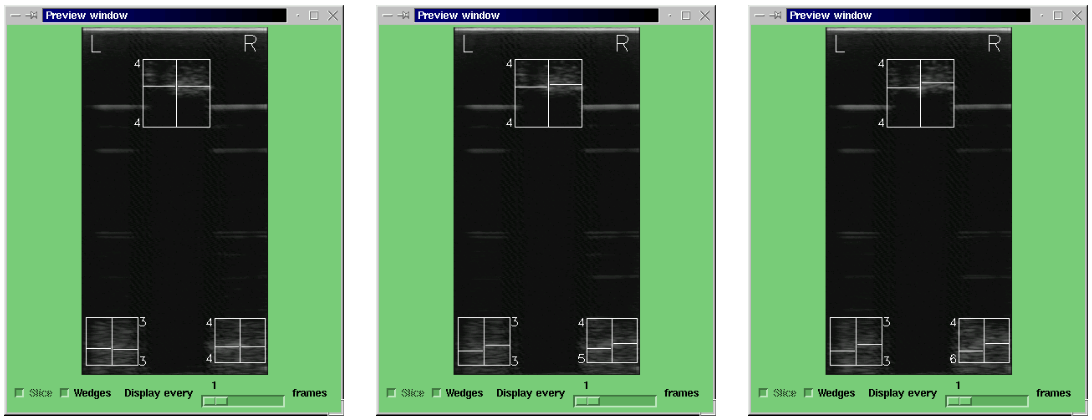

(a) Aligned (b) 200 μm (c) 400 μm

Figure 9: **Wedges imaged with a 7 MHz linear array probe at a depth setting of 6 cm**. The phantom's z stage was used to acquire images at precise degrees of misalignment. Note the tightness of the intensity peaks at the top of the B-scan, compared with the broadness of the peaks at the bottom, where the ultrasound beam is much wider. Nevertheless, the automatic left and right peak finding algorithm is evidently sensitive to misalignments as small as 200 μm.

but also its direction. However, where the ultrasound beam is so wide that the wedges lie entirely inside it, the assumptions behind the peak finders break down and we must instead turn to the profile difference indicators. Current generation scanners and probes produce images far superior to those in Figures 9 and 10, with improved elevational focusing. In particular, modern "2.5D" probes have especially narrow beams.

It only remains to describe the procedure for segmenting the top and bottom wires in Figure 7. The user clicks on three points, two on the top wire and one on the bottom wire. A straight line is fitted through the top two points and a parallel line is drawn through the bottom point. To eliminate any inter- and intra-operator variability, an automatic algorithm locates the wires in the vicinity of the user's clicks. The sole purpose of the user intervention is to define a 9×21 pixel region of interest centred on the click point. A vertical intensity profile is extracted within the region of interest and the wire located at the intensity centroid (the crosses in Figure 7), in exactly the same way as described for wedges in Figure 8.

3 Evaluation of the calibration phantom

3.1 Precision

The precision of the new technique was assessed by performing multiple calibrations with the Toshiba model SSA-270A/HG ultrasound scanner. To demonstrate the versatility of the phantom, one set of experiments featured the 7 MHz linear array probe at a depth setting of 6 cm, another the 3.5 MHz convex array probe at a depth setting of 12 cm. B-scans were transferred at 25 Hz from the scanner's composite analogue output to an 800 MHz PC running Stradx[9], where they were digitised using an inexpensive video framegrabbing card. The wedge segmentation algorithm ran comfortably in real time on this modest hardware, providing interactive visualisation of the ultrasound beam's current state of alignment. The experimental setup is shown in Figure 11.

The tank was filled with water at approximately 50°C, the temperature at which the speed of sound in water matches the ultrasound scanner's assumed speed of sound in soft tissue (1540 ms^{-1}). No attempt was made to maintain the water temperature during the hour or so required for each

[9] `http://mi.eng.cam.ac.uk/~rwp/stradx/`

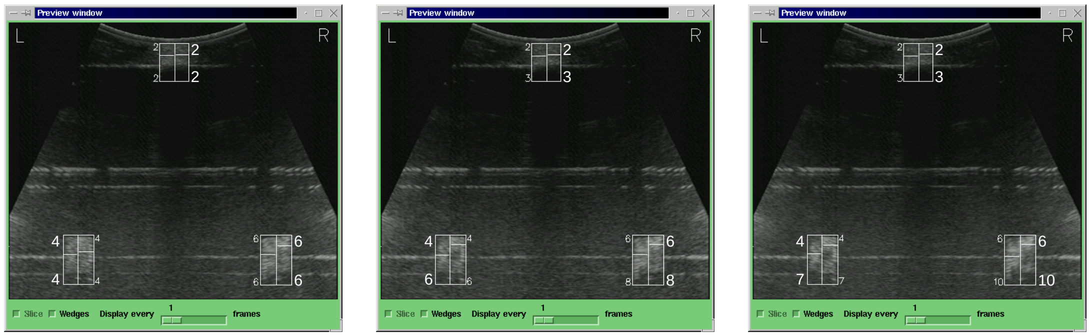

(a) Aligned (b) 200 μm (c) 400 μm

Figure 10: **Wedges imaged with a 3.5 MHz convex array probe at a depth setting of 12 cm**. In this case, the automatic left and right peak finding algorithm works correctly for the top wedge-pair, but produces fairly random results with the bottom two pairs. The beam is so wide at the bottom that the wedges lie entirely inside it. The assumption that the B-scan contains two tight reflections moving up and down relative to each other, as in Figure 4, is therefore no longer valid. Instead, we must look for less specific asymmetries between the left and right profiles, caused by intensity variations across the width of the beam. The profile difference indicators, shown alongside each wedge box (and mirrored in bold on the other side, since they are barely legible in the screenshot), give a minimum value at the point of alignment, and are sensitive to misalignments of the order of 200 μm.

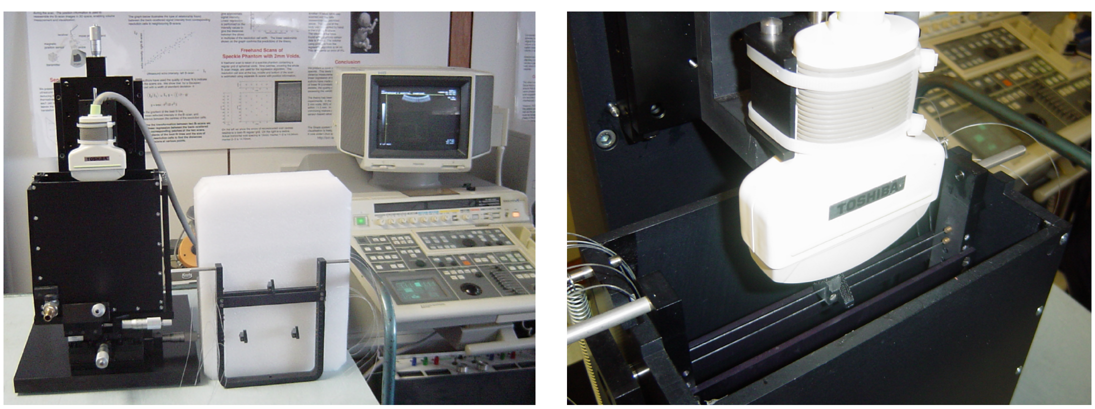

(a) (b)

Figure 11: **Experimental apparatus**. (a) shows a general view of the experimental apparatus. The phantom is on the left, with the 3.5 MHz convex array probe mounted on the gantry. The position sensor is not required for the experiments and is not attached to the probe holder. The wire frame with the three wedge-pairs is to the right of the phantom, ready to be put in the tank. In the background is the Toshiba model SSA-270A/HG ultrasound scanner. (b) shows a close-up view of the probe just before the tank is filled with water. The top pair of wedges is visible just beneath the probe face.

set of experiments, nor was any correction factor applied to the B-scans. The results are therefore indicative of what might be achieved without cumbersome monitoring of the water temperature.

Twenty calibrations were performed with the 7 MHz linear array probe at its 6 cm depth setting, each taking two to three minutes. Before each calibration, the probe holder was replaced on the gantry, the wire frame was replaced in the tank, and the z, pitch and yaw micrometers were scrambled. The x and y stages were used to place the top wedge-pair near the top centre of the B-scan, and then left undisturbed throughout the twenty trials. The bottom wedge-pairs were positioned near the bottom corners of the cropped B-scan (4.5 cm below the top pair, separated from each other by around 2 cm). The ultrasound machine's transmit power and receive gain were set to arbitrary (though not unreasonable) values before each calibration. However, the time gain control was kept constant, with high gain around the wedges and low gain elsewhere, and the focus was set at 7 mm for imaging the top wedge-pair (when adjusting the z stage) and 55 mm for the bottom wedge-pairs (when adjusting the pitch and yaw stages). Between the z and yaw adjustments, the z stage was moved in and out a little, so that the automatic segmentation algorithm could locate the top wedges' central axis (refer to Figure 8 and its caption). The left and right peak finding algorithm was used to determine correct alignment of all wedges. Typical B-scans are shown in Figure 9.

Twenty further calibrations were performed with the 3.5 MHz convex array probe at its 12 cm depth setting. The experimental procedure was exactly as for the 7 MHz probe, except that the bottom two wedge-pairs were moved lower down the frame and further apart (8.25 cm below the top pair, separated from each other by around 9 cm). The focus was set at 20 mm when imaging the top wedge-pair and 97 mm for the bottom wedge-pairs. The left and right peak finding algorithm was used for the top wedge-pair, the profile difference indicators for the lower two pairs. Typical B-scans are shown in Figure 10.

The results are given in Table 3. Instead of showing the variation in the parameters of the calibration transformation $\mathbf{T}_{S \leftarrow B}$, we have chosen instead to show the effect of this variation on the position of the B-scan. This is a common measure of precision: for example, it is closely related to the "calibration reproducibility" in [7]. It is a measure which is independent of the particular calibration phantom, and is therefore useful for comparing different techniques.

Each calibration places the B-scan coordinate system B at a slightly different position relative to the fixed sensor system S. We therefore define a nominal B-scan system B^*, aligned with the mean B-scan orientation across the twenty trials. We consider the positions $\mathbf{p}$ of the B-scan's centre and corners in B^*. For each point $\mathbf{p} = (x, y, z)$, we calculate a 3D error $\|\mathbf{p} - \bar{\mathbf{p}}\|$, where $\bar{\mathbf{p}}$ is the mean of $\mathbf{p}$ across the twenty trials. The mean 3D error and its 95% confidence limits are given in Table 3. We also show the 95% confidence limits of the individual components of $\mathbf{p}$.

The results for the new phantom are in the bottom two blocks of Table 3. For comparison, in the top block we reproduce from our earlier work [12] results for the "Cambridge" phantom [10]. These were obtained with a superior 5–10 MHz linear array probe at its 6 cm depth setting, with digital image transfer between the scanner and the PC. Since the B-scan dimensions are similar, the "Cambridge" results are best compared with those for the 7 MHz probe in the block below. The new technique's superior precision is immediately evident: the errors are reduced by around a factor of two. At the time of publication [12], the "Cambridge" phantom was demonstrably more precise than other techniques reported in the literature. The new technique also compares favourably with more recent work [7].

Some more subtle characteristics of the results warrant comment. For the 7 MHz linear array probe, the pattern of x and y errors suggests uncertainty in the size of the B-scan caused by incorrect estimation of the pixel dimension: the y error increases top to bottom, the x error increases left to right. This is most likely explained by changes in the speed of sound: recall that the pixel dimension is estimated from the separation of the two wires, which would appear to grow as the water cooled. For the 3.5 MHz curvilinear array probe, speed of sound errors cause the straight wires to be imaged as curves. This leads to uncertain estimates of their angle to the horizontal, and thus uncertain estimates of the roll of the B-scan about the top wedge-pair. Again, the x and y errors are compatible with this explanation: the centre of the B-scan is relatively close to the top wedge-pair, the bottom corners very distant.

Frequency (depth)		mean error (mm) 3D	95% confidence limits (mm) x	y	z	3D
5–10 MHz	centre	0.44	± 0.37	± 0.18	± 0.93	< 0.78
(6 cm)	top left	0.29	± 0.27	± 0.21	± 0.50	< 0.50
"Cambridge"	top right	0.48	± 0.25	± 0.30	± 0.95	< 0.84
	bottom left	0.54	± 0.52	± 0.18	± 1.08	< 0.94
	bottom right	0.67	± 0.50	± 0.34	± 1.42	< 1.18
7 MHz	centre	0.15	± 0.22	± 0.14	± 0.19	< 0.26
(6 cm)	top left	0.15	± 0.18	± 0.05	± 0.28	< 0.27
	top right	0.17	± 0.28	± 0.07	± 0.24	< 0.30
	bottom left	0.22	± 0.19	± 0.27	± 0.35	< 0.39
	bottom right	0.28	± 0.27	± 0.28	± 0.47	< 0.49
3.5 MHz	centre	0.24	± 0.43	± 0.33	± 0.14	< 0.43
(12 cm)	top left	0.43	± 0.55	± 0.55	± 0.49	< 0.75
	top right	0.35	± 0.50	± 0.50	± 0.28	< 0.62
	bottom left	0.58	± 0.82	± 0.90	± 0.38	< 1.01
	bottom right	0.58	± 0.95	± 0.68	± 0.48	< 1.02

Table 3: **Variation in B-scan location due to calibration imprecision**. The table shows how variation in $\mathbf{T}_{S \leftarrow B}$ affects the location of key points on the B-scan. The mean 3D error and 95% confidence limits are given in each case. The (x, y, z) axes are in the nominal B-scan coordinate system B^*. The results for the new technique are in the bottom two blocks. For comparison, the top block shows earlier results for the "Cambridge" phantom.

Next, note the similarity of the z errors for both types of probe: this is because the wedges are just as sensitive to misalignment of the low frequency beam than the high frequency one. In contrast, the x and y errors are about twice as high for the 3.5 MHz probe than for the 7 MHz one. This has two root causes. First, the increased size of the 3.5 MHz B-scan (14×12 cm, compared with 3×6 cm for the 7 MHz B-scan) amplifies the effects of any speed of sound errors. Secondly, the increased pixel dimension (0.28 mm/pixel, compared with 0.14 mm/pixel for the 7.5 MHz probe) affects the precision of the automatic segmentation algorithms used to estimate $\mathbf{T}_{W \leftarrow B}$. Recall that $\mathbf{T}_{W \leftarrow B}$ acts only in the plane of the B-scan, not perpendicular to it.

Finally, note how the z errors are minimised at the centre of the B-scan. This is because the wedges fix the z position of the top centre and bottom corners to within a close tolerance t. Imagine holding a sheet of paper and randomly jiggling it so that these three points move a maximum distance t perpendicular to the sheet. On average, the centre of the sheet will move much less than this.

3.2 Accuracy

The precision results indicate the expected uncertainty in a *single* calibration with the new phantom. We could reduce this error as much as we please, by averaging the results of many calibrations. For example, the mean of 100 calibrations will exhibit ten times less variation than in Table 3. But this does not tell us anything about accuracy: would the mean of a very large number of calibrations be biased in any way?

Probe calibration accuracy has traditionally been very difficult to assess, largely because of the very indirect relationship between the calibration action (eg. the repeated scanning of a fixed point) and the end result (the spatial relationship between the sensor and scan plane). Authors have therefore resorted to more general measures of overall system accuracy. A common approach is to scan a known object in three dimensions, segment pertinent features in the 3D reconstruction and then calculate measures such as navigation accuracy (how closely a point in the reconstruction coincides with its true 3D location) and distance reconstruction accuracy (how well distances measured in the reconstruction correspond to the true lengths). While these accuracy measures

Frequency (depth)		mean error (mm) 3D	95% confidence limits (mm) x	y	z	3D
7 MHz (6cm)	centre	0.14	± 0.14	± 0.18	± 0.18	< 0.25
	top left	0.14	± 0.14	± 0.18	± 0.18	< 0.25
	top right	0.14	± 0.14	± 0.18	± 0.18	< 0.25
	bottom left	0.14	± 0.15	± 0.18	± 0.19	< 0.25
	bottom right	0.15	± 0.15	± 0.18	± 0.19	< 0.26
3.5 MHz (12cm)	centre	0.14	± 0.15	± 0.18	± 0.19	< 0.25
	top left	0.15	± 0.14	± 0.19	± 0.19	< 0.26
	top right	0.15	± 0.14	± 0.19	± 0.20	< 0.26
	bottom left	0.16	± 0.17	± 0.19	± 0.21	< 0.28
	bottom right	0.16	± 0.17	± 0.19	± 0.22	< 0.29

Table 4: **Asymptotic accuracy.** The table shows the effect of possible bias in $\mathbf{T}_{G \leftarrow W}$, within the bounds established in Table 1. Translational uncertainty dominates and the errors are consequently fairly uniform across the B-scan.

are certainly useful to know, calibration errors are only one of many contributing factors. Other sources of error include the position sensor, the segmentation of pertinent features, the assumed dimensions of the known object, the speed of sound, any interpolation in the reconstruction algorithm, analogue video corruption and temporal calibration [12].

Unusually, for this new technique we are able to assess the accuracy *of the calibration process alone*. There is a direct relationship between the calibration action and the end result, as expressed in equation (1). Any bias in $\mathbf{T}_{S \leftarrow B}$ can be estimated by considering the individual contributions of $\mathbf{T}_{S \leftarrow G}$, $\mathbf{T}_{G \leftarrow W}$ and $\mathbf{T}_{W \leftarrow B}$.

Neglecting speed of sound errors, which can be avoided by careful control of the water temperature, it is reasonable to assume no significant bias in the calculation of $\mathbf{T}_{W \leftarrow B}$. The algorithm for finding the central axis of the wedge-pair relies on nothing more than local symmetry, while the wires are located at their intensity centroids, which should coincide with their true centres for such thin wires as used in this work. Imaging distortion in the scanner itself will bias $\mathbf{T}_{W \leftarrow B}$, but the same distortion will affect any 3D reconstruction. Since one would want the calibration to compensate as far as possible, it would be incorrect to consider this as a source of calibration error. Any bias in $\mathbf{T}_{G \leftarrow W}$ will lie within the bounds established in Table 1. As explained in Section 2, any bias in $\mathbf{T}_{S \leftarrow G}$ can be suppressed by averaging a large number of measurements, recalibrating the pointer each time and moving the phantom around the Polaris cameras' field of view.

The asymptotic accuracy is therefore a function of the values in Table 1. Their effect on the B-scan position was determined by evaluating equation (1) 1000 times, randomly perturbing $\mathbf{T}_{G \leftarrow W}$ within the tabulated bounds each time. These Monte Carlo experiments were performed for representative alignments of the 3.5 MHz and 7 MHz probes. The results can be found in Table 4. Bear in mind that this is the *asymptotic* accuracy. To achieve a calibration approaching this accuracy, one would have to average a very large number of measurements to obtain $\mathbf{T}_{S \leftarrow G}$, be careful not to subsequently disturb the position of the sensor on the probe holder, and then carefully control the water temperature while performing a large number of calibrations, replacing the probe holder and frame each time. Finally, one would average the calculated $\mathbf{T}_{S \leftarrow B}$ values to obtain a high precision result.

A more realistic estimate of calibration accuracy is given in Table 5. Here we combine the uncertainty in $\mathbf{T}_{G \leftarrow W}$ with additional uncertainty in $\mathbf{T}_{S \leftarrow G}$. We assume that $\mathbf{T}_{S \leftarrow G}$ is determined by averaging 25 individual measurements: this could be carried out in a couple of hours. Again, the values in Table 5 are based on 1000 Monte Carlo experiments, randomly perturbing both $\mathbf{T}_{G \leftarrow W}$ and $\mathbf{T}_{S \leftarrow G}$ within their respective bounds. We must stress that Table 5 gives a measure of the phantom's practical *accuracy*, ie. the expected bias in the calibration result. If only one calibration is performed, the result is also subject to precision uncertainty as given in Table 3.

Frequency		**mean error (mm)**	**95% confidence limits (mm)**			
(depth)		3D	x	y	z	3D
7 MHz	centre	0.31	± 0.45	± 0.23	± 0.43	< 0.54
(6cm)	top left	0.28	± 0.39	± 0.24	± 0.40	< 0.49
	top right	0.27	± 0.39	± 0.23	± 0.37	< 0.48
	bottom left	0.35	± 0.50	± 0.24	± 0.50	< 0.61
	bottom right	0.34	± 0.50	± 0.23	± 0.48	< 0.59
3.5 MHz	centre	0.33	± 0.49	± 0.23	± 0.47	< 0.58
(12cm)	top left	0.33	± 0.38	± 0.30	± 0.52	< 0.58
	top right	0.30	± 0.38	± 0.27	± 0.43	< 0.52
	bottom left	0.45	± 0.61	± 0.30	± 0.70	< 0.78
	bottom right	0.42	± 0.61	± 0.27	± 0.64	< 0.74

Table 5: **Accuracy with realistic estimation of $\mathbf{T}_{S\leftarrow G}$.** The table shows the effect of possible bias in $\mathbf{T}_{S\leftarrow G}$, within the bounds established in Table 2 and assuming an estimate based on the average of 25 individual measurements (ie. five times less uncertainty than in Table 2). Also included is the uncertainty in $\mathbf{T}_{G\leftarrow W}$, within the bounds established in Table 1.

	x	y	z	azimuth	elevation	roll
mean $\mathbf{T}_{S\leftarrow G}$ wrt G	-17.30 mm	74.51 mm	1.25 mm	$-89.79°$	$0.51°$	$177.91°$
uncertainty in $\mathbf{T}_{S\leftarrow G}$	± 39 μm	± 44 μm	± 126 μm	$\pm 0.064°$	$\pm 0.043°$	$\pm 0.395°$

Table 6: **Repeatability of mounting the sensor on the probe holder**. This was estimated by re-mounting the sensor twenty times and measuring its position relative to the Polaris cameras. All the readings were then transformed to a new coordinate system, such that their mean coincided with the mean $\mathbf{T}_{S\leftarrow G}$ in Table 2. We can thus observe the effect of mounting errors on the elements of $\mathbf{T}_{S\leftarrow G}$. The tabulated uncertainty is 1.96 times the standard deviation of the twenty measurements, ie. 95% confidence limits assuming the measurements to be Normally distributed.

There is one final source of error to consider. Suppose $\mathbf{T}_{S\leftarrow G}$ is measured carefully, averaging many individual observations, but the sensor is then removed and replaced on the probe holder. Unless it returns to exactly the same position as before, there will be an error in $\mathbf{T}_{S\leftarrow G}$ that will bias future calibrations. To estimate this error, we replaced the sensor on the convex array probe holder twenty times, recording its position each time. To reduce the effects of sensor noise, each position was based on the average of twenty Polaris readings. The results are shown in Table 6, indicating a high degree of repeatability in all directions except z and roll. These results are as one would expect given the design of the mount, which is faithfully illustrated in Figure 6. The roll (the twist of the sensor on its mounting pin) is constrained only by the mounting screw impinging on the flat, while the z position is sensitive to any play of the pin within its hole, unlike the x position, which is well constrained by the mounting screw. The effect of this extra source of bias is shown in Table 7. This includes all the sources of bias identified thus far, and is therefore indicative of the phantom's accuracy in routine, everyday use. As before, the values are derived from the results of 1000 Monte Carlo trials.

4 Conclusions and further work

The principal distinguishing feature of the new calibration technique is its ability to operate without the position sensor. This significantly improves the calibration precision, so that reliable results can be obtained from a single measurement with no need for averaging. The asymptotic accuracy is more than one is ever likely to need, given the dominance of other sources of error in the 3D reconstruction process. Everyday accuracy, achievable with relatively little effort, is

Frequency (depth)		mean error (mm) 3D	95% confidence limits (mm) x	y	z	3D
7 MHz (6cm)	centre	0.34	±0.53	±0.25	±0.47	< 0.60
	top left	0.33	±0.47	±0.26	±0.47	< 0.58
	top right	0.31	±0.47	±0.24	±0.40	< 0.54
	bottom left	0.40	±0.59	±0.26	±0.57	< 0.69
	bottom right	0.38	±0.59	±0.24	±0.51	< 0.66
3.5 MHz (12cm)	centre	0.37	±0.57	±0.25	±0.50	< 0.65
	top left	0.45	±0.46	±0.33	±0.82	< 0.79
	top right	0.37	±0.45	±0.29	±0.61	< 0.65
	bottom left	0.55	±0.70	±0.33	±0.93	< 0.97
	bottom right	0.50	±0.70	±0.29	±0.79	< 0.87

Table 7: **Everyday accuracy.** The tabulated values account for re-mounting the sensor on the probe holder, as well as uncertainty in the measurement of $\mathbf{T}_{S \leftarrow G}$ and $\mathbf{T}_{G \leftarrow W}$.

sufficient for the vast majority of applications.

While the new technique's accuracy and precision compare favourably with previously documented alternatives, they should not be allowed to distract from its other qualities, in particular robustness and ease of use. Importantly, it is difficult for even an inexperienced user to arrive at a totally incorrect calibration. This is because the adjustment mechanism, used in conjunction with the image processing tools, largely deskills the alignment process[10]. In contrast, inexperienced users struggle with techniques requiring multiple views in non-degenerate configurations, and with unguided, hand-held alignment techniques.

The decoupling of the calibration procedure from the position sensor has ergonomic implications too. The phantom can be placed anywhere convenient, without the constraint that it must be somewhere within the position sensor's operating zone. Thus, it could act as a probe rest, sitting next to the ultrasound machine. If the user places the probe in the phantom when changing the scanner's depth, pan or zoom settings, instant and automatic recalibration is feasible. To achieve this, the wire frame is festooned with an array of small, reflective beads. The idea is that any reasonably sized portion of the frame can be identified by the pattern of beads it contains. Thus, with the aid of an automatic segmentation algorithm to detect the beads in the B-scan, the system can infer any changes in depth, pan or zoom — see Figure 12. Note that no realignment is required: this is just a matter of re-estimating $\mathbf{T}_{W \leftarrow B}$. The design of suitable bead patterns and corresponding segmentation algorithms is a matter for future work.

Acknowledgements

The authors thank Alastair Palmer for his pilot study on depth, pan and zoom recalibration, and for providing the images in Figure 12.

References

[1] C. D. Barry, C. P. Allot, 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.

[2] L. G. Bouchet, S. L. Meeks, G. Goodchild, F. J. Bova, J. M. Buatti, and W. A. Friedman. Calibration of three-dimensional ultrasound images for image-guided radiation therapy. *Physics in Medicine and Biology*, 46:559–577, 2001.

[10] There is even the potential for complete automation, with motorised adjustment stages driven under image-guided control.

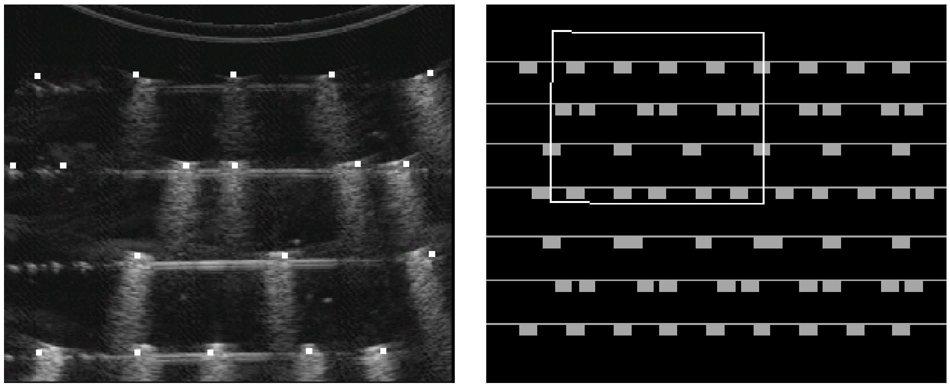

(a) B-scan (b) frame

Figure 12: **Depth, pan and zoom recalibration**. The user has adjusted the pan and zoom settings to focus on a region of interest near the top left corner of the frame, as shown in (b). This portion of the frame is identifiable from the pattern of beads visible in the B-scan (a). With automatic detection of the beads (the white dots in (a)), the probe can be recalibrated instantly and automatically.

[3] P. R. Detmer, G. Bashein, T. Hodges, K. W. Beach, E. P. Filer, D. H. Burns, and D. E. Strandness Jr. 3D ultrasonic image feature localization based on magnetic scanhead tracking: *in vitro* calibration and validation. *Ultrasound in Medicine and Biology*, 20(9):923–936, 1994.

[4] A. Fenster, D. B. Downey, and H. N. Cardinal. Three-dimensional ultrasound imaging. *Physics in Medicine and Biology*, 46:R67–R99, 2001.

[5] A. Gee, R. Prager, G. Treece, and L. Berman. Engineering a freehand 3D ultrasound system. *Pattern Recognition Letters*, 24(4–5):757–777, February 2003.

[6] D. F. Leotta, P. R. Detmer, and R. W. Martin. Performance of a miniature magnetic position sensor for three-dimensional ultrasound imaging. *Ultrasound in Medicine and Biology*, 23(4):597–609, 1997.

[7] F. Lindseth, G. A. Tangen, T. Langø, and J. Bang. Probe calibration for freehand 3-D ultrasound. *Ultrasound in Medicine and Biology*, 29(11):1607–1623, November 2003.

[8] N. Pagoulatos, D. R. Haynor, and Y. Kim. A fast calibration method for 3-D tracking of ultrasound images using a spatial localizer. *Ultrasound in Medicine and Biology*, 27(9):1219–1229, September 2001.

[9] R. W. Prager, A. H. Gee, and L. Berman. Stradx: real-time acquisition and visualization of freehand three-dimensional ultrasound. *Medical Image Analysis*, 3(2):129–140, 1999.

[10] R. W. Prager, R. N. Rohling, A. H. Gee, and L. Berman. Rapid calibration for 3-D freehand ultrasound. *Ultrasound in Medicine and Biology*, 24(6):855–869, 1998.

[11] Y. Sato, M. Nakamoto, Y. Tamaki, T. Sasama, I. Sakita, Y. Nakajima, M. Monden, and S. Tamura. Image guidance of breast cancer surgery using 3-D ultrasound images and augmented reality visualization. *IEEE Transactions on Medical Imaging*, 17(5):681–693, October 1998.

[12] G. M. Treece, A. H. Gee, R. W. Prager, C. J. C. Cash, and L. Berman. High definition freehand 3D ultrasound. *Ultrasound in Medicine and Biology*, 29(4):529–546, April 2003.