Computerized Medical Imaging and Graphics 58 (2017) 23–32

Contents lists available at ScienceDirect

Computerized Medical Imaging and Graphics

journal homepage: www.elsevier.com/locate/compmedimag

Accurate 3D reconstruction of bony surfaces using ultrasonic synthetic aperture techniques for robotic knee arthroplasty

William Kerr[c,*], Philip Rowe[a], Stephen Gareth Pierce[b]

[a] Department of Biomedical Engineering, University of Strathclyde, Glasgow, United Kingdom
[b] Department of Electronic and Electrical Engineering, University of Strathclyde, Glasgow, United Kingdom
[c] Advanced Forming Research Centre, University of Strathclyde, 85 Inchinnan Drive, Renfrew, PA4 9LJ, United Kingdom

ARTICLE INFO

Article history:
Received 23 May 2016
Received in revised form
18 November 2016
Accepted 24 March 2017

Keywords:
Ultrasound
Total focussing method
Synthetic aperture focussing technique
Robotics
3D surface reconstruction
Computer-aided surgery

ABSTRACT

Robotically guided knee arthroplasty systems generally require an individualized, preoperative 3D model of the knee joint. This is typically measured using Computed Tomography (CT) which provides the required accuracy for preoperative surgical intervention planning. Ultrasound imaging presents an attractive alternative to CT, allowing for reductions in cost and the elimination of doses of ionizing radiation, whilst maintaining the accuracy of the 3D model reconstruction of the joint. Traditional phased array ultrasound imaging methods, however, are susceptible to poor resolution and signal to noise ratios (SNR). Alleviating these weaknesses by offering superior focusing power, synthetic aperture methods have been investigated extensively within ultrasonic non-destructive testing. Despite this, they have yet to be fully exploited in medical imaging.

In this paper, the ability of a robotic deployed ultrasound imaging system based on synthetic aperture methods to accurately reconstruct bony surfaces is investigated. Employing the Total Focussing Method (TFM) and the Synthetic Aperture Focussing Technique (SAFT), two samples were imaged which were representative of the bones of the knee joint: a human-shaped, composite distal femur and a bovine distal femur. Data were captured using a 5 MHz, 128 element 1D phased array, which was manipulated around the samples using a robotic positioning system. Three dimensional surface reconstructions were then produced and compared with reference models measured using a precision laser scanner. Mean errors of 0.82 mm and 0.88 mm were obtained for the composite and bovine samples, respectively, thus demonstrating the feasibility of the approach to deliver the sub-millimetre accuracy required for the application.

1. Introduction

Osteoarthritis (OA) is a degenerative disease affecting the articulating joints which causes pain and can eventually lead to a complete loss of mobility in the affected joint (Buckwalter and Mankin, 1997). When symptoms of OA of the knee become disabling, the prescribed solution is knee arthroplasty, which can alleviate pain and increase mobility (Woolhead et al., 2005). While traditional procedures are well established, the bone removal methods are outdated and inaccurate, which can lead to poor alignment and, as a result, poor clinical outcome (Huang et al., 2012). Robotically guided systems can improve cutting accuracy, leading to superior alignment and improved clinical outcome (Choong et al., 2009).

In order to plan and implement procedures using such systems, a preoperative 3D model of the joint is usually required (Roche et al., 2009; Siebert et al., 2002; Cobb et al., 2006; Hananouchi et al., 2008), which is currently achieved using Computed Tomography (CT). While this method is seen as the 'gold standard' in this capacity, it is costly (Fred, 2004) and applies significant doses of ionizing radiation which can contribute to and even cause malignancy (Albert, 2013). Ultrasound imaging could provide a viable alternative to CT in this capacity by offering comparable accuracies, reducing cost and eliminating the risk of ionizing radiation.

Considerable research efforts have been made in ultrasound bone surface imaging for intraoperative registration to preoperative CT models in computer-aided surgery – a brief review of which can be found in (Hacihaliloglu et al., 2014). A smaller number of groups have attempted to employ ultrasound for preopera-

* Corresponding author at: Advanced Forming Research Centre, University of Strathclyde, 85 Inchinnan Drive, Renfrew, PA4 9LJ, United Kingdom.
E-mail addresses: william.kerr@strath.ac.uk, williamkerr69@gmail.com (W. Kerr).

http://dx.doi.org/10.1016/j.compmedimag.2017.03.002

tive imaging, with most relying on Statistical Shape Models (SSM) to compensate for the sparsity of ultrasound-derived data points (Barratt et al., 2008; Talib et al., 2005; Chan et al., 2004). These models are entirely dependent on *a priori* data in the form of CT training sets. The accuracy of the reconstructions is, therefore, highly dependent on how representative the training set is of the variation of prospective samples (Heimann and Meinzer, 2009). Training sets are often small and include only healthy knee joints which lack the kind of variation in bone shape found in OA sufferers due to bone remodeling (Fleute and Lavall, 1998). These studies did not account for such localized variation, bringing into question the robustness of the use of SSMs in osteoarthritic patients. With this in mind, it would be preferable to reconstruct the surfaces using only patient-specific data.

In addition to employing SSMs, these studies had in common that manually positioned, optically tracked probes were employed. These would have imposed intra- and inter-operator variability in performance and possible line of sight problems (Hüfner et al., 2005). Further to this, traditional B-scan methods were employed universally, which are associated with low resolution and depth-dependent image quality (Guarneri et al., 2015).

Offering improved focusing capability and increased resolution and SNR over more traditional B-scan methods (Holmes et al., 2005), synthetic aperture imaging methods have become commonplace in research in nondestructive testing (NDT). These methods have been employed within biomedical imaging in attempts to increase frame rates through recursive imaging (Nikolov et al., 1999; Jensen et al., 2006) and through decreasing transmissions per image (Chiao et al., 1997).

Reductions in hardware size and cost have also been achieved by use of synthetic aperture techniques for simplification for small scale and hand-held systems (O'Donnell and Thomas, 1992; Karaman et al., 1995). Employment within medical imaging has, however, been relatively limited due to both the inability of most commercial systems to perform Full Matrix Capture (FMC) (a necessity for the Total Focussing Method (TFM)) and the requirement of real time imaging (Jensen et al., 2006). In particular, the techniques have yet to be applied to bone surface imaging, where only traditional B-scan techniques have been considered. Employment in this field could offer the level of accuracy required for preoperative imaging for robotic knee arthroplasty. However, NDT research into synthetic aperture methods has concentrated on the characterization of small defects (Zhang et al., 2013; Hunter et al., 2008; Wilcox et al., 2007), with surface imaging only being performed on simple, continuous surfaces for refractive correction in dual media applications (Mcgilp et al., 2014; Weston et al., 2012). A further challenge in using these methods is that a high number of images would be required to properly represent the entire sample surface, each with accurate positional data.

The ability of a synthetic aperture based ultrasound imaging system to accurately reconstruct complex, bony surfaces is assessed in order to determine the efficacy of using such methods in preoperative imaging for knee arthroplasty. Two samples were inspected: a composite, human-shaped distal femur (which provided accurate shape and scale) and a bovine distal femur (which provided representative acoustic properties and a test of system robustness). Using a 5 MHz, 128 element 1D array, FMCs were captured using an FPGA-based phased array controller, which allowed for high data throughput. The data were processed using both TFM and a form of the Synthetic Aperture Focussing Technique (SAFT) at high rates, achieved by use of a Graphics Processing Unit (GPU) implementation of both algorithms (Dziewierz et al., 2012). These processing methods were employed as they provided different characteristics found at opposing ends of synthetic aperture reconstructions, which allowed for optimisation of the synthetic aperture parameters for the application. Probe manipulation and positional recording were realized using a robotic positioning system. Finally, the 3D, ultrasound-derived surface data were compared with reference models, providing performance levels for the different methods and for the system in general.

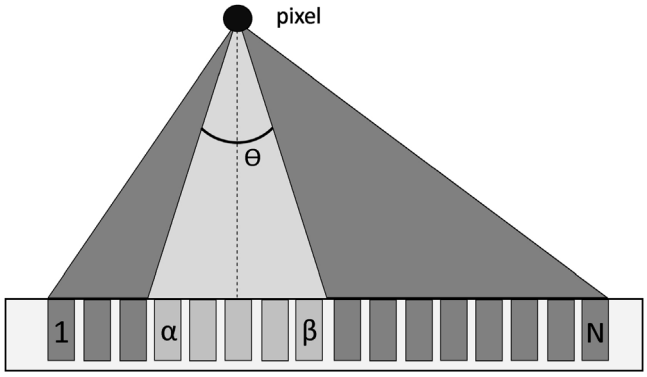

Fig. 1. Comparison of the TFM and SAFT synthetic aperture definitions, with the elements constituting the SAFT aperture shown in a lighter shade.

2. Materials and methods

2.1. Synthetic aperture methods

Both of the techniques employed herein required FMC, which begins by firing the first element of an array and receiving on all. The second element is then fired and reception is carried out on all elements once again. The process is then repeated over all N elements, resulting in an $N \times N$ matrix of A-scans which is collectively known as the full matrix.

It is possible to perform numerous processing techniques with the full matrix in post processing. One possible method is TFM, which synthetically focusses in reception for every pixel in the image using the all the elements in the array. The process begins by discretizing the region of interest into a grid, with each point being equivalent to a pixel in the final scalar image. The intensity of each pixel is calculated as described in (1), in which $S_{i,j}$ is the A-scan associated with a transmission at the ith element and reception at the jth element, while (x, y) are the coordinates of the pixel. The time of travel from the transmitting element to the pixel is $T_{i(x,y)}$, while that from the pixel to the receiving element is $T_{j(x,y)}$. In order to produce a full scalar image, the summation is repeated for every pixel.

$$I(x,y) = \sum_{i,j=1}^{N} S_{i,j}\left(T_{i(x,y)} + T_{j(x,y)}\right) \tag{1}$$

A form of SAFT was also employed in which the same process was performed, but focusing was not carried out on using the full aperture. Instead, the elements constituting the synthetic aperture were defined by the position of the pixel in question. This is shown in Fig. 1, where the SAFT aperture is restricted to those elements contained within an isosceles triangle defined by the angle θ. This is contrasted by the TFM aperture which employs the full aperture. The intensity of a pixel in a SAFT image is, then, given by (2), where α and β are the first and last elements of the aperture, as displayed in Fig. 1.

$$I(x,y) = \sum_{i=1}^{N}\sum_{j=\alpha}^{\beta} S_{i,j}\left(T_{i(x,y)} + T_{j(x,y)}\right) \tag{2}$$

These values vary depending with θ, with an increase in angle effectively increasing the aperture size.

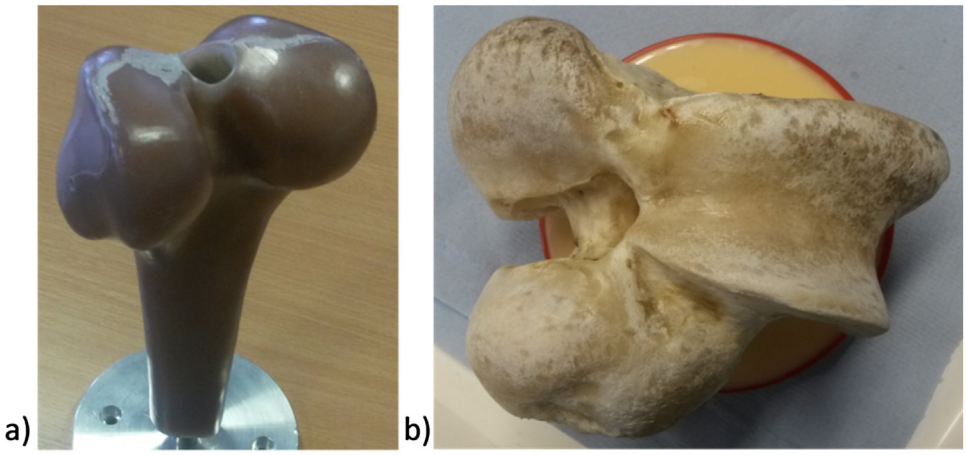

Fig. 2. The composite (left) and bovine (right) distal femur samples.

Employing SAFT, therefore, limited the width of the synthetic aperture relative to the full aperture offered by TFM. While this decreased the available resolving power, the likelihood of the formation of side lobes also decreased. Side lobes are regions of high intensity in the image which are do not represent physical reflectors (Dziewierz et al., 2012) and are a result of the reconstruction method, rather than any physical phenomenon. Referring to (1), it can be seen that the time of flight, rather than the position, determines the likelihood of a pixel receiving a non-zero contribution from the associated transmit-receive pairing. Therefore, regions which do not contain a true reflector, but which share a similar time of flight as a true reflector result in non-zero contributions from the A-scan and can form side lobes. Referring to (2), it can be seen that by employing SAFT, the likelihood of a pixel at position (x, y) receiving a non-zero A-scan sample value can be decreased by restriction of the number of transmit-receive pairings, β–α. Further, by limiting the receive aperture, the regions which could possibly attain a non-zero contribution are geometrically restricted, resulting in a decrease in the possibility of the formation of side lobes. A more detailed explanation of the effect on receive aperture width on the formation of side lobes can be found in (Kerr et al., 2016).

2.2. Experimental apparatus

Two samples were inspected, the first of which was a Sawbones® composite distal femur (Pacific Research Laboratories Inc., WA, USA). While the constituent materials of the sample were not specifically identified by the manufacturer, it was designed for mechanical testing and composed of fiber-strengthened epoxy. Therefore, the sample did not possess the acoustic properties of real bone, but did possess representative shape and scale. The second sample was a bovine distal femur which had undergone defleshing and boiling to remove soft tissues, leaving only the exposed cortical surface. As can be seen in Fig. 2, a base was attached to both samples to afford a level of stability whilst under inspection. Laser scanning was performed using a Faro Quantum laser line probe (Faro Technologies, FL, USA), which provided 3D reference models to an accuracy of 0.1 mm (Rajna et al., 2009).

For ultrasonic acquisition, a 128 element phased array (Vermon, Tours, France) with a center frequency of 5 MHz was employed. This offered both a large aperture width and a small theoretical resolving limit of 0.15 mm in water. The array was driven by a FlawInspecta phased array controller (Diagnostic Sonar Ltd., Livingstone, UK). This is a modular platform, allowing for parallel digitization using FlexRIO FPGA cards (National Instruments Corporation, TX, USA). The configuration used in this study employed two 32 channel digitizers, which allowed for parallel reception of 64 elements, with a maximum of four elements in simultaneous transmission. Two individual transmissions were, therefore, required for each full matrix transmission event when employing all 128 elements in reception for FMC.

Data transfer rates were limited by firmware configurations to an FMC frequency of 0.3 Hz (Brown et al., 2015). In principle, this can be increased with the next firmware update, such that the only limits on the Pulse Repetition Frequency (PRF) would be the number of individual transmissions during an FMC transmission event and the physical constraints of wave travel in the material (Lines et al., 2006). Therefore, the full scan for the human sample, consisting of 348 FMC acquisitions, could be reduced to as little as 24 s.

In order to reconstruct 3D surfaces using 2D images, accurate probe manipulation and positional recording were vital. To this end, a KUKA KR 5 arc HW (KUKA Robotics Corporation, Ausburg, Germany) was employed, providing six degrees of freedom and the ability to implement complex tool paths. Two bespoke probe holders were manufactured which allowed for line of sight access to all necessary surfaces by way of an increase in the variation of valid robot poses. A standard, two stage KUKA calibration method (KUKA, 2010) was employed to calibrate the Tool Centre Points (TCP) using bespoke calibration parts. These allowed for the TCP to be defined as the center of the probe face without the use of the probe itself during the calibration procedure. The accuracy of the calibration was dependent on a number of factors, including the accuracy of the calibration parts and the operator experience. Therefore, the accuracy of the calibration was assessed using six Vicon T160 cameras (Vicon Motion Systems, Oxford, UK), producing a worst-case mean error of 0.53 mm. This was comparable to commercial optical tracking employed in computer-aided surgery (Chassat and Lavallee, 1998), but the use of a robotic manipulator eliminated the intra- and inter-operator dependence of a manually positioned, optically tracked system, as well as the associated line of sight issues. During data capture, the position and orientation of the TCP was continuously relayed to the host computer via Ethernet, using custom written software employing the KUKA Robot Sensor Interface (RSI).

Given that the shape of the sample was unknown prior to inspection precluding a complex geometry inspection path, a simple cuboidal path was employed for both samples, as illustrated in Fig. 3(a). The path was implemented in KUKA Robot Language (KRL) – a BASIC-like, domain-specific language which allowed for elemental tool paths to be defined.

The samples were placed in a water bath, ensuring complete submersion, as can be seen in Fig. 3(b). The controlling KRL script required that an initial position be defined, as well as an approx-

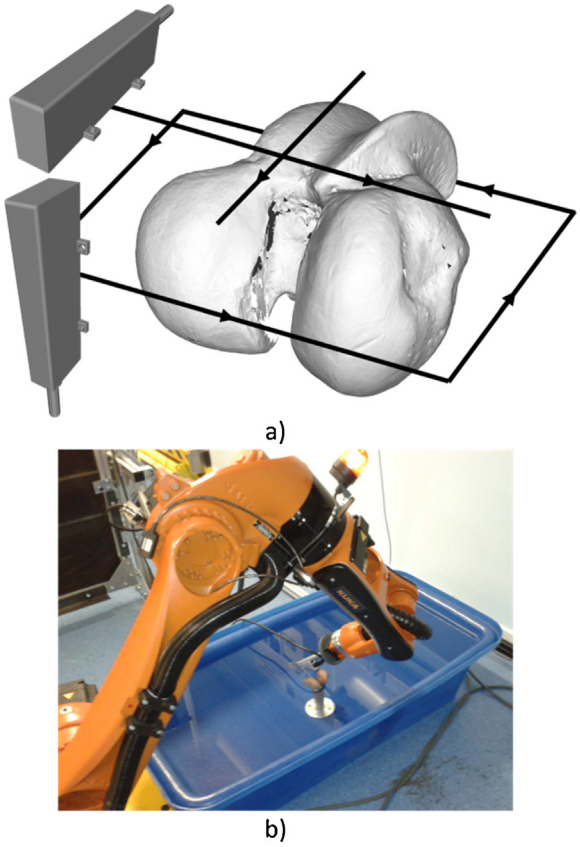

Fig. 3. Cuboidal tool path relative to the bovine sample (a) and experimental setup (b).

imate sample length, width and height. In doing so, a minimum standoff between the probe face and the sample surface was maintained at approximately 20 mm.

2.3. Surface reconstruction

Due to the time of flight calculations for every pixel and transmit-receive pair combination, image production using synthetic aperture methods is computationally expensive. However, due to the independent nature of each calculation, the techniques are highly parallelizable (Lambert et al., 2012; Romero et al., 2009). Therefore, cueART – a software platform which provides significant reductions in computation time by implementing the reconstruction algorithms on a GPU (Mcgilp et al., 2014; Dziewierz et al., 2012) – was employed for image reconstruction.

Based on previous studies (Kerr et al., 2016), it was anticipated that use of a narrow aperture SAFT would produce superior results. However, it was unknown which synthetic aperture width would provide the best results. Therefore, every FMC acquisition was processed using TFM and SAFT, employing 20°, 30° and 40° apertures, which offered both an extreme contrast between TFM and narrow aperture SAFT and a more detailed examination of which synthetic aperture width was optimal. Employing a speed of sound of 1480 ms^{-1} (i.e. the speed of sound in water (Halmshaw, 1991)) and a pixel size of 0.01 mm for both samples, images with a width of 13.44 cm and depths of 8.96 cm and 11.52 cm for the composite and bovine samples, respectively, were produced. This resulted in 1344×896 and a 1344×1152 pixel scalar images for the composite and bovine samples. Using an NVIDIA GeForce GTX 760 Ti (1344 CUDA cores), the reconstruction time for a single image from the associated FMC was approximately 30 s.

To extract the surface profile from each image, a column-wise thresholding approach was employed. For every column of pixels in the image, those pixels with intensity greater than a defined threshold limit were identified. The pixel with the smallest Z coordinate (i.e. the pixel closest to the probe face) was retained, while the rest were discarded. This was repeated for every column in the image, producing an array of points represented by 2D coordinates. This method was preferred over curve fitting techniques as it allowed for a high level of variation in surface profiles (including discontinuities) without any user intervention. This preserved automation and eliminated intra- and inter-operator dependence. The accuracy of this approach has previously been investigated by the authors in assessing surface reconstructions from known geometric primitives (Kerr et al., 2016).

To correctly locate the measured 2D points in 3D space relative to the coordinate system of the robot, the measured coordinates were rotated and then translated using the corresponding measured KUKA robot pose (position and orientation). This is presented in (3), where the subscripts f, o and T represent the final 3D coordinates, the original, image-based coordinates and the TCP coordinates. Additionally, R_{xyz} is the rotation matrix in the x, y, z order, respectively. In this fashion, the set of 2D measurements were converted into a 3D cloud of surface points.

$$\begin{pmatrix} X_f \\ Y_f \\ Z_f \end{pmatrix} = R_{xyz} \begin{pmatrix} X_o \\ Y_o \\ Z_o \end{pmatrix} + \begin{pmatrix} X_T \\ Y_T \\ Z_T \end{pmatrix} \quad (3)$$

The point clouds were spatially down sampled so that the minimum distance between any two points was 2 mm so that density inconstancies were minimized. The point clouds were then imported into Geomagic Wrap (Geomagic Inc., NC, USA), in which the wrapping algorithm (Fletcher et al., 2006) was applied, which reconstructed the point cloud data into a surface mesh. Surface filling and surface smoothing functions were then applied across the entire surface, resulting in the final surface meshes. The parameters of the smoothing and filling functions were set at the software default values. While it was possible to decrease the error in the resulting meshes through iterative parameter changes, doing so would produce a lack of system robustness and increase operator dependencies. By employing the same parameters for every data set, the level of automation and operator independence was maximized. This process is illustrated diagrammatically in Fig. 4.

In order to determine the accuracy of the reconstructed surfaces, the meshes were imported into CloudCompare (EDF/Telecom ParisTech, Paris, France), in which they were randomly subsampled down to 10^6 points. This was performed so that during comparison the faces of the mesh were represented, rather than only the vertices, providing a more accurate comparison. The representative point clouds were initially manually maneuvered so that an approximate positional match was found with the reference model. Fine matching was then achieved using Iterative Closest Point (ICP). The absolute error for each point was calculated using the Euclidean distance between the point and the nearest vertex on the surface of the reference mesh. From these values, mean error, maximum error and standard deviation were calculated.

3. Results

3.1. Composite distal femur

A feature common to each of the reconstruction methods was the ability to reconstruct surfaces lacking normality to the probe face. As demonstrated in Fig. 5, part of the surface was at an angle of approximately 45° to the probe face and was, despite this, recon-

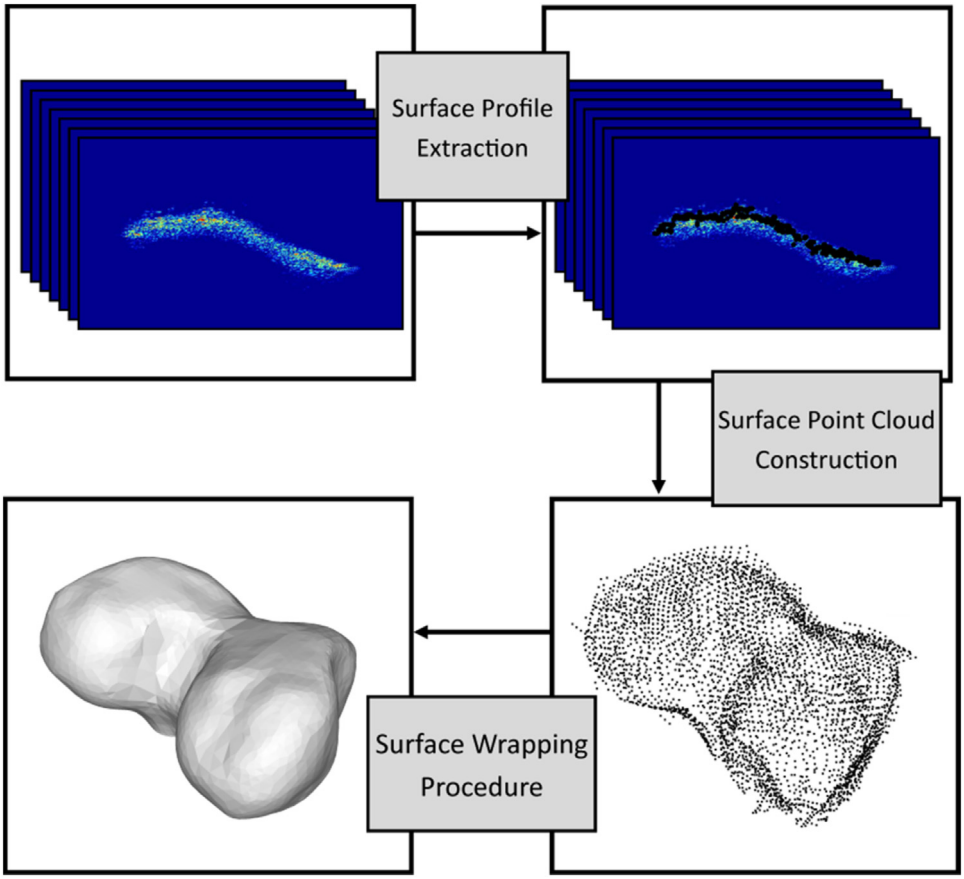

Fig. 4. The surface reconstruction method, beginning with input images. The surface profile was then extracted in the form of 2D coordinates. These were then rotated and translated according to the associated KUKA data to produce a surface point cloud. The point cloud data was then wrapped in Geomagic, yielding the final surface representation.

structed well. In fact, a small region sat at an angle of less than 45° and still formed part of the resulting image. This effect was made possible by the focusing capability of the reconstruction methods.

As illustrated in Fig. 5, the TFM images contained significant unwanted side lobe content, which resulted in 'ghost' surfaces either side of the true surface, parts of which were of intensities comparable to the true surface. This was caused by the large synthetic aperture employed for pixel reconstruction. In contrast, only weak side lobes were produced in the SAFT images, due to the limited synthetic aperture width. In particular, the side lobe contribution decreased as the synthetic aperture decreased, resulting in almost no side lobes in the 20° SAFT images. In fact, the maximum side lobe intensities for a typical image were −15 dB, −9 dB and −6 dB for the 20°, 30° and 40° SAFT apertures respectively, while that for TFM was −3 dB.

Because of this, while the global threshold limit for surface profile extraction was set to −8 dB for the SAFT images, it was increased to −4 dB for TFM, to reduce the likelihood of side lobe extraction. Despite this, side lobes were occasionally extracted in the TFM images, which resulted in a small erroneous surface, as can be seen anterior to the medial condyle in part (d) of Fig. 6.

Despite the unwanted effects described above, an increase in the synthetic aperture width offered superior resolution, with the thickness of the surface profiles decreasing as the aperture width increased. Additionally, TFM allowed for surface reconstructing past the edge of the probe, which was not possible with the limited synthetic aperture of the SAFT methods.

This effect is contrasted further with the B-scan image shown in part (e) of Fig. 5, which did not allow for any reconstruction beyond the probe face and formed a poorly defined representation of the surface. Further, while no side lobes were present, the resolution of the image was poorer in the horizontal axis due to the image reconstruction method.

This scalar image was produced by stacking A-scans, such that for each transmission event, the A-scan which corresponded to both transmission and reception was plotted in the vertical direction from the appropriate element position. This method of image reconstruction does not correspond to B-scan imaging that would typically be employed in biomedical ultrasound imaging, where focused or sectorial imaging methods would be employed. However, these imaging methods could not be replicated using the FMC acquisition technique employed herein. With this in mind, while the image presented in part (e) might not be representative of the capability of typical B-mode imaging methods, it does allow for a comparison with the ability of the synthetic aperture methods in terms of reconstructing non-perpendicular surfaces with good surface definition.

The errors found in the comparisons with the reference model and the final reconstructed surfaces meshes are presented in Fig. 6. While all the surfaces were of similar shape, the TFM-derived mesh appeared to have less definition than the SAFT results. The reason for this was that the surface point cloud resulting from TFM processing was sparser than those from SAFT. This was caused by the higher threshold limit, which yielded smaller surface profile extractions. Therefore, the surface wrapping procedure employed a higher level

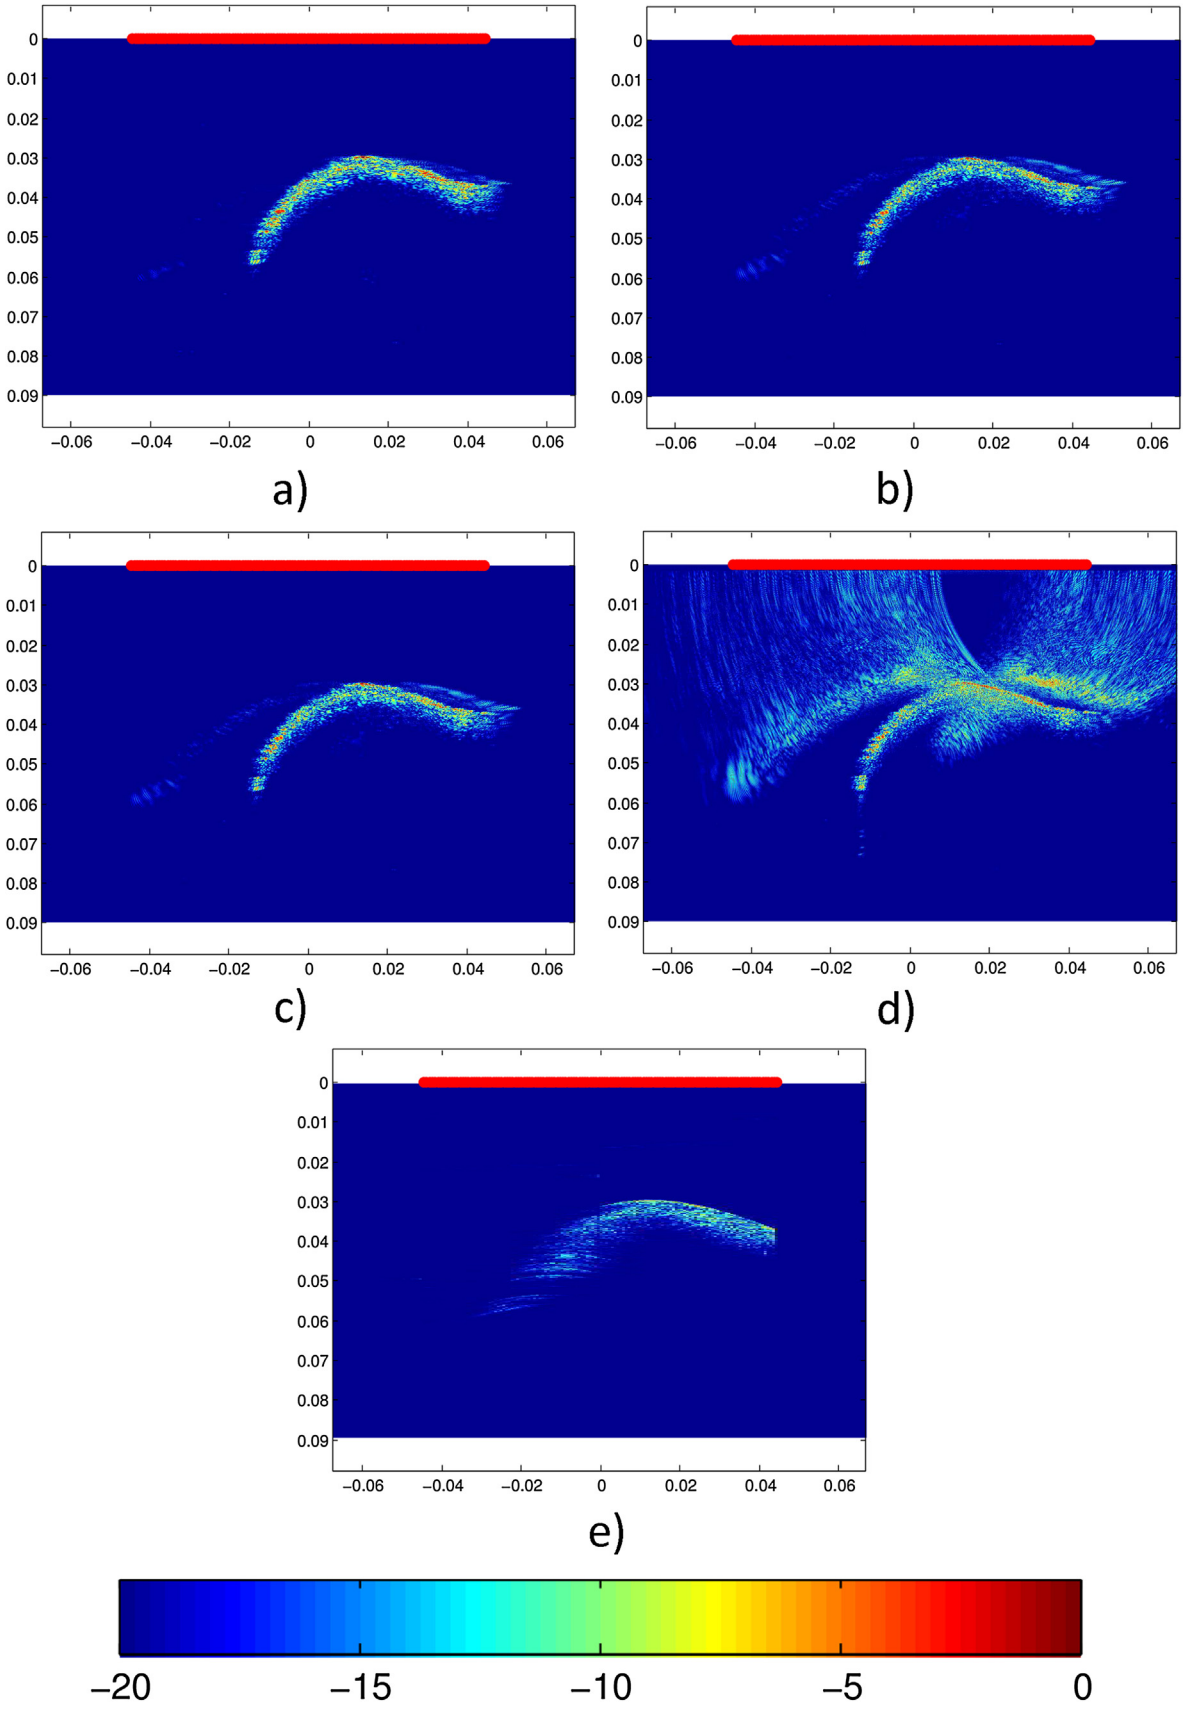

Fig. 5. Typical scalar images from the composite sample resulting from SAFT with 20° (a), 30° (b) and 40° (c) apertures and from TFM (d). Part (e) shows a B-scan image reconstructed using the same FMC by stacking the A-scans of each transceiver. The red bar represents the position of the probe face. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

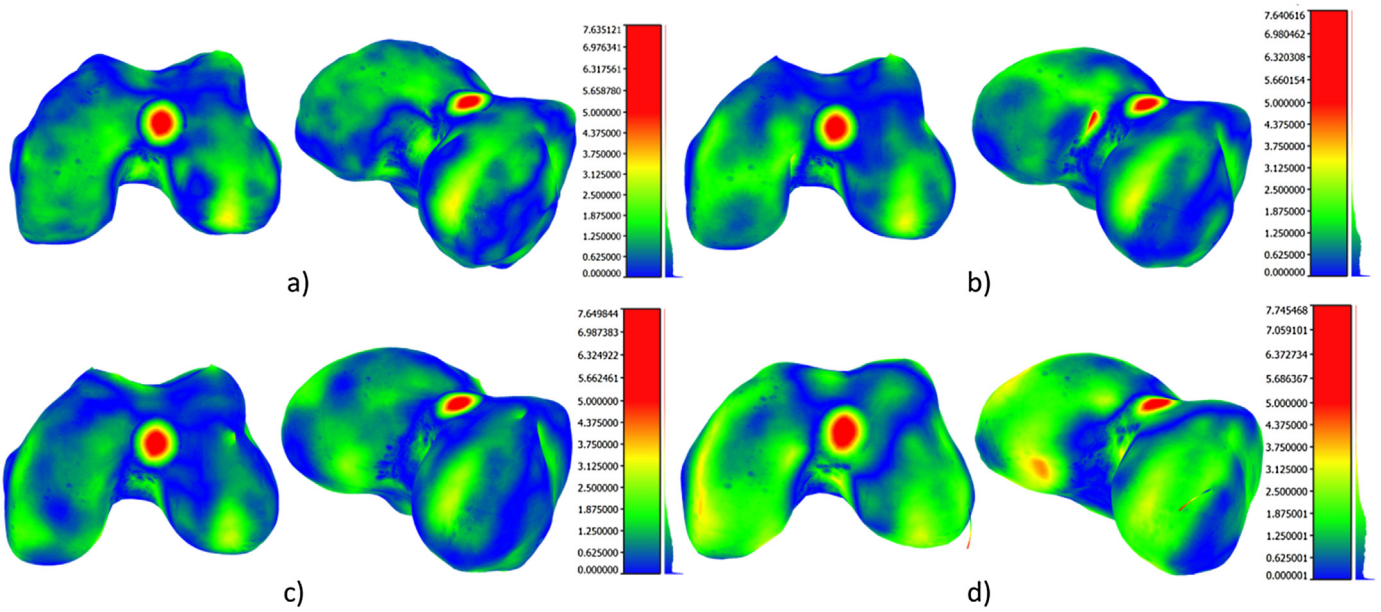

Fig. 6. Color-coded absolute error in millimeters for surfaces representing the composite, human-shaped sample, created using SAFT with 20° (a), 30° (b), 40° (c) apertures and with TFM (d).

Table 1
Results of Comparison with Reference Models.

Sample	Processing Method	Mean Error (mm)	Standard Deviation (mm)	Maximum Error (mm)
Human	SAFT (20°)	0.82	0.63	7.64
Human	SAFT (30°)	0.94	0.71	7.65
Human	SAFT (40°)	0.85	0.65	7.65
Human	TFM	1.40	0.96	7.75
Bovine	SAFT (20°)	0.88	1.02	10.47
Bovine	SAFT (30°)	0.91	1.03	10.20
Bovine	SAFT (40°)	0.91	1.01	10.28
Bovine	TFM	1.25	1.23	12.61

of shape estimation and was less dependent on the data. As a result, there was a higher level of error generally. This was reflected in the high mean error and standard deviation, as can be seen in Table 1.

While the regions in which low error existed differed for each processing method, the areas containing the higher errors were maintained. A region on the front of the medial condyle, for example, was of high error in each example. This was caused by the transmitting waves being reflected away from the probe face, as the surface lacked normality relative to the transmitting waves. More obviously, a large circular region of high error was present in the center of every example. This corresponded to the drilled hole in the physical sample, as can be seen in Fig. 2. While the hole was represented in the surface point clouds, the wrapping function is designed to create continuous surfaces and so filled the hole. The hole was implemented for mounting purposes and such features would be unlikely in real distal femurs. Despite this, it was included in calculations of the mean and maximum error and standard deviation, as it was a true feature which was inaccurately recreated.

As can be seen in Table 1, the TFM mesh was the only example to yield a mean error greater than 1 mm. This, in addition to producing a significantly higher standard deviation, indicated that TFM was the least suitable for complex bone surface profile imaging. All of the SAFT examples, on the other hand, achieved mean errors below 1 mm. In particular, the 20° aperture SAFT mesh was significantly below the 1 mm threshold, with the lowest mean error and standard deviation.

3.2. Bovine distal femur

The images resulting from the data captured during the inspection of the bovine sample presented similar characteristics to those seen in the composite sample. As such, the global thresholding limits of −8 dB for SAFT and −4 dB for TFM images were maintained for the bovine sample.

As with the composite sample, the shape of the mesh was similar for all the processing methods. However, the TFM-derived mesh suffered from a lack of definition. This was again caused by the associated sparse surface point cloud.

As can be seen in Fig. 7, protrusions were present on the back of the lateral condyle in every SAFT-derived mesh, producing a region of high error. This was caused by reflections from a region of high curvature. Generally, specular reflection dominates when imaging bony surfaces (Brendal et al., 2002), but in regions of high curvature, diffuse reflection can begin to dominate. In such cases, off-axis waves can be reflected and recorded by the transducer. The signals representing these reflections were, therefore, inaccurately placed, as they were wrongly assigned to the on-axis image plane. The size of the protrusion grew as the aperture width increased, as more elements receiving the offending reflections were employed in pixel reconstruction. The feature was not found in the TFM mesh, however, as the higher image threshold did not allow for extraction of the feature.

The regions in which high error was found were common to all the resulting meshes. This was due to incomplete regions in the reference model, as illustrated in Fig. 8. As there were no reference

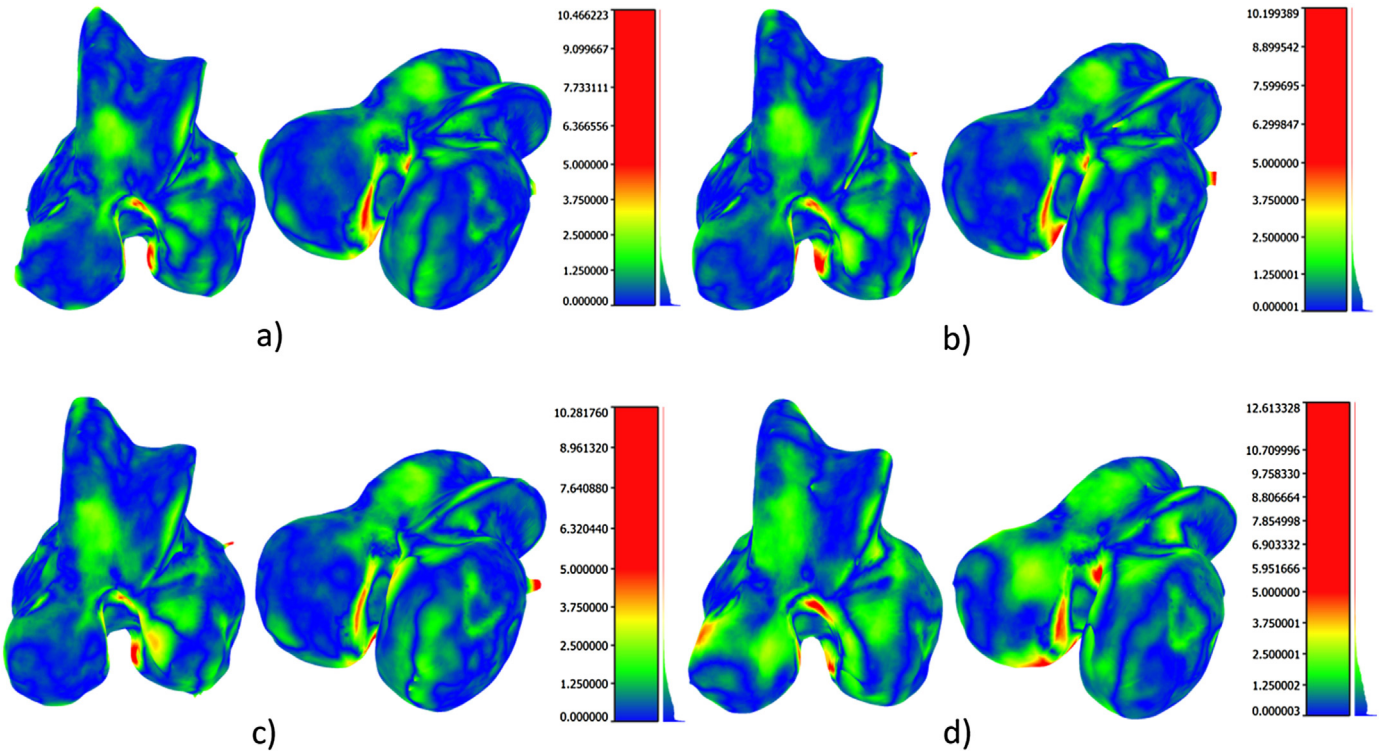

Fig. 7. Color-coded absolute error in millimeters for surfaces representing the bovine sample, created using SAFT with 20° (a), 30° (b), 40° (c) apertures and with TFM (d).

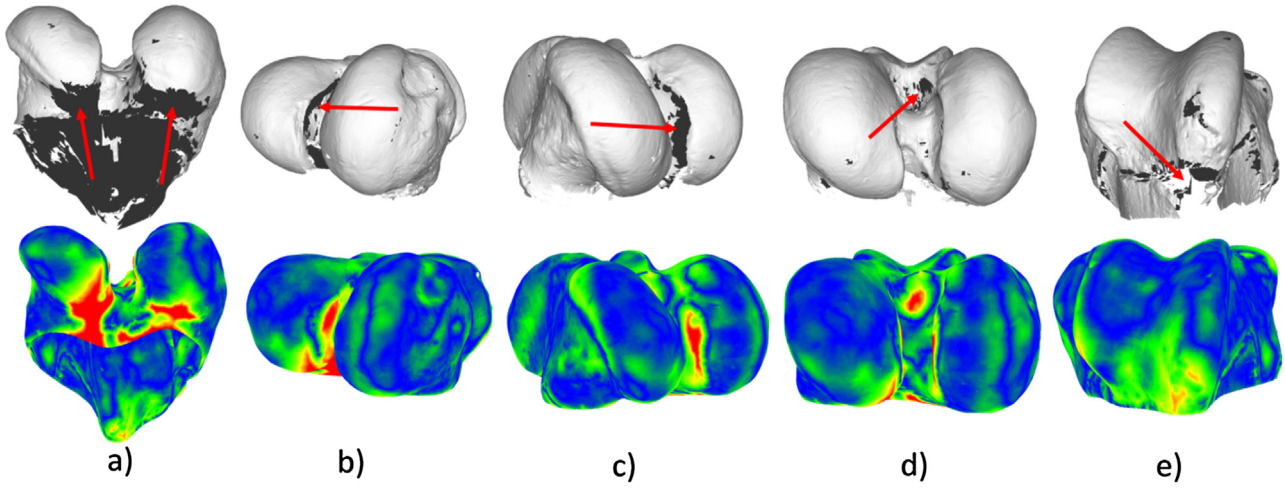

Fig. 8. Incomplete regions in the reference model (top row) caused by laser line of sight issues (a–c), remnants of tendon attachment (d) and bone degradation due to boiling (e). Also shown are the corresponding regions of high error (bottom row).

surfaces to compare with in these regions, the comparison yielded high errors.

As can be seen in Table 1, the mean error, maximum error and standard deviation for the bovine sample showed an increase relative to the composite sample. This was due to the incomplete regions in the reference model, which resulted in significant regions in which the error was large. If those regions were complete, the error therein would have been significantly decreased, which would have decreased the bovine error shown in Table 1 to similar levels as found with the composite sample.

Despite this, all the SAFT meshes achieved a mean error below 1 mm. The TFM mesh produced the worst mean error, maximum error and standard deviation and did not achieve similar levels of accuracy. For this reason, TFM was found to be the least suitable imaging method for reconstructing complex, bony surfaces. As with the composite sample, the 20° aperture SAFT mesh was found to be the most accurate in terms of mean error.

4. Discussion

Sub millimetre mean errors were achieved for the full surface reconstruction of both samples. Providing context to this result requires that a target mean error is defined such that reconstructions below this threshold are considered accurate enough for the application. The target mean error in preoperative imaging for robotic knee arthroplasty is often quoted as 1 mm, but justification is rarely provided for this rule of thumb. However, mean errors obtained in placement of knee implants have been found to be in

excess of 1 mm in commercial robotic systems (Lonner et al., 2015; Smith et al., 2014). As this is considered to be an acceptable level of error, preoperative imaging for such systems should provide mean errors below this, in which case 1 mm would be a reasonable target. In addition to this, when preoperatively reconstructing bony surfaces with CT for robotic knee arthroplasty, standard scans employ 1 mm thickness slices (Roche et al., 2009; Conditt and Roche, 2009; Citak et al., 2013; Davies et al., 2006), limiting the resolution of the voxelised CT data to 1 mm. While the volumetric reconstruction accuracy of CT is not strictly defined by this resolution, it is strongly associated with it (Prionas et al., 2011). This provides further justification for the stated target mean error, given that slice thickness is usually adjustable and 1 mm has been deemed acceptably accurate by manufacturers of commercial systems.

The reconstructed surfaces of both samples surpassed the desired target, demonstrating, in principle, the efficacy of employing the presented system as an accurate form of preoperative knee joint imaging. These results compare well with similar studies, which failed to achieve sub millimetre mean errors (Barratt et al., 2008; Chan et al., 2004).

In saying this, the reconstruction of the medial condyle of the human-shaped sample contained a region with particularly high error (>3 mm). This level of error could lead to inaccuracy when matching the real bone to the model during the procedure, as that area corresponds to a palpation region. The consequences of this could be misplacement of the implant and, as a result, failure thereof. The error in this region was caused by the position of the probe face relative to the sample surface, as described in Section 3.1. As the positioning in this case was a result of the simple probe path employed, an improvement in the measurement procedure would be to employ a two stage scan, as proposed in (Kerr et al., 2016). This would provide a basic geometry with the first scan which could then be employed to generate a patient-specific probe path which would provide a greater level of normality between the probe face and the sample surface and, as a result, eliminate the described region of high error.

The 3D surface models presented herein were produced without the use of *a priori* data, using an almost fully automated system, which is in contrast to similar studies in which SSMs have been employed in combination with manual and semiautomatic stages (Barratt et al., 2008; Talib et al., 2005; Chan et al., 2004). This served as a demonstration of system robustness, given that the same system parameters (with the exception of the length of the cuboidal robot path) were employed for two samples possessing highly different shapes and sizes. This is of particular importance in the chosen application, in that osteoarthritic knees would be highly variable in shape and size, and would, therefore, require an imaging system which could account for such variability. This robustness was in part due to the robotic positioning system, which offered the same level of accuracy for each scan as a consequence the high level of repeatability offered by industrial robots (Brown et al., 2015) and the lack of inter- and intra-operator variability.

It should be mentioned, however, that other studies have employed cadaveric specimens. In this work, only the surface of the bone was interrogated, eliminating the effects of soft tissues and problems with line of sight access to the condylar surfaces associated with intact knee joints. Gaining access to the condylar surfaces of the knee could, however, be achieved by performing multiple scans. The first of these would be performed while the joint is in full extension and would include scanning a portion of the shaft of both the tibia and femur. A second scan would then be conducted with the joint at a level of flexion, with the same regions of the shafts included in the scan. This would allow completely different parts of the condylar surfaces to be exposed between scans. Positional matching of the two data sets would be achieved using the scans performed on the shafts, which would be fully accessible at any level of flexion.

The results of surface reconstructions using a number of synthetic aperture widths were presented, with the three narrow SAFT apertures clearly outperforming the full aperture of TFM. This was not surprising, as while TFM offers high resolving capability, this is optimised for acquiring information about small reflectors. The narrow aperture SAFT, on the other hand, provided more positional dependence during reconstruction, which allowed weaker, yet valid reflectors to have greater representation in the reconstructed image. This effect can be particularly useful when imaging surfaces, as it is beneficial to represent the entire surface, rather than those regions which reflect particularly well. Equally importantly, SAFT imaging significantly reduced side lobe levels, allowing for a lower threshold level, greater surface profile extraction and fewer spurious extractions due to side lobes.

Additionally, it was found that the most accurate reconstructions were found with the smallest SAFT synthetic aperture employed. Selection of synthetic aperture width when using this form of SAFT was an attempt to eliminate the negative aspects of TFM, whilst maintaining the positive characteristics as much as possible. It was found that the 20° SAFT offered the best balance in this capacity, with the larger synthetic apertures showing higher levels of side lobes.

5. Conclusions

The work presented herein has shown that complex bony surfaces can be reconstructed to the desired level of accuracy using ultrasonic synthetic aperture methods. Achieving mean errors of 0.82 mm and 0.88 mm for composite human and bovine distal femur samples respectively, SAFT with a narrow 20° synthetic aperture was found to be optimal for reconstruction with synthetic aperture methods. Employing robotic probe positioning, the surfaces of two samples of highly different shape and size were accurately reconstructed, illustrating robustness of the proposed system. This was achieved using sample-specific data only, eliminating dependence on *a priori* data.

The samples possessed characteristics shared with the bones of real human knee joints and were, therefore, representative thereof. As such, the results serve as a proof of concept for the use of ultrasound-based systems in preoperative imaging for robotic knee arthroplasty. Future studies will seek to further validate this proof of concept by incorporating soft tissues and line of sight issues associated with condylar surfaces in intact joints.

Acknowledgement

This work was supported by a studentship from the Engineering and Physical Sciences Research Council (grant number EP/F50036X/1).

References

Albert, J.M., 2013. Radiation risk from CT: implications for cancer screening. Am. J. Roentgenol. 201 (July), 81–87.

Barratt, D.C., Chan, C.S.K., Edwards, P.J., Penney, G.P., Slomczykowski, M., Carter, T.J., Hawkes, D.J., 2008. Instantiation and registration of statistical shape models of the femur and pelvis using 3D ultrasound imaging. Med. Image Anal. 12 (June (3)), 358–374.

Brendal, B., Winter, S., Rick, A., Stockheim, M., Ermert, H., 2002. Registration of 3D CT and ultrasound datasets of the spine using bone structures. Comput. Aided Surg. 7 (3), 146–155.

Brown, R.H., Pierce, S.G., Collison, I., Dutton, B., Jackson, J., Lardner, T., Macleod, C., Morozov, M., 2015. Automated full matrix capture for industrial processes. 41st Annu. Rev. Prog. Quant. Nondestruct. Eval. 34, 1967–1976.

Buckwalter, J.A., Mankin, H.J., 1997. Articular cartilage. Part II: degeneration and osteoarthrosis, repair, regeneration, and transplantation. J. Bone Joint Surg. Br. 79, 612–632.

Chan, C.S.K., Barratt, D.C., Edwards, P.J., Penney, G.P., 2004. Cadaver validation of the use of ultrasound for 3D model instantiation of bony anatomy in image guided orthopaedic surgery. Med. Image Comput. Comput. Interv. 3217, 397–404.
Chassat, F., Lavallee, S., 1998. Experimental protocol of accuracy evaluation of 6-D localizers for computer-integrated surgery: application to four optical localizers. Med. Image Comput. Comput. Interv. – Miccai'98 1496, 277–284.
Chiao, R.Y., Thomas, L.J., Silverstein, S.D., 1997. Sparse array imaging with spatially-encoded transmits. 1997 IEEE Ultrason. Symp. Proceedings. An Int. Symp. (Cat. No. 97CH36118) 2, 1679–1682.
Choong, P.F., Dowsey, M.M., Stoney, J.D., 2009. Does accurate anatomical alignment result in better function and quality of life? Comparing conventional and computer-assisted total knee arthroplasty. J. Arthroplasty 24 (June (4)), 560–569.
Citak, M., Suero, E.M., Citak, M., Dunbar, N.J., Branch, S.H., Conditt, M.A., Banks, S.A., Pearle, A.D., 2013. Unicompartmental knee arthroplasty: is robotic technology more accurate than conventional technique? Knee 20 (4), 268–271.
Cobb, J., Henckel, J., Gomes, P., Harris, S., Jakopec, M., Rodriguez, F., Barrett, A., Davies, B., 2006. Hands-on robotic unicompartmental knee replacement: a prospective, randomised controlled study of the acrobot system. J. Bone Joint Surg. Br. 88 (2), 188–197.
Conditt, M.A., Roche, M.W., 2009. Minimally invasive robotic-arm-guided unicompertmental knee arthroplasty. J. Bone Joint Surg. 91 (1), 63–68.
Davies, B., Jakopec, M., Harris, S.J., Rodriguez y Baena, F., Barrett, A., Evangelidis, A., Gomes, P., Henckel, J., Cobb, J., 2006. Active-constraint robotics for surgery. Proc. IEEE 94 (9), 1696–1704.
Dziewierz, J., Gachagan, A., Lord, N., Mullholland, A.J., 2012. An application-specific design approach for 2D ultrasonic arrays. 51st Annual Conference of the British Institute of Non-Destructive Testing 2012 (NDT 2012).
Fletcher, G., Gloth, T., Edelsbrumer, H., Fu, P., Methods, Aparatus and Computer Program Products That Reconstruct Surfaces From Data Point Sets, US 7,023,432 B2, 2006.
Fleute, M., Lavall, S., 1998. Building a complete surface model from sparse data using statistical shape models: application to computer assisted knee surgery. Lect. Notes Comput. Sci. 1496, 879–887.
Fred, H.L., 2004. Drawbacks and limitations of computed tomography. Tex. Heart Inst. J. 31 (4), 345–348.
Guarneri, G.A., Pipa, D.R., Junior, F.N., Valéria, L., De Arruda, R., Victor, M., Zibetti, W., 2015. A sparse reconstruction algorithm for ultrasonic images in nondestructive testing. Sensors 15, 9324–9343.
Hacihaliloglu, I., Rasoulian, A., Rohling, R.N., 2014. Local phase tensor features for 3-D ultrasound to statistical shape + pose spine model registration. IEEE Trans. Med. Imaging 33 (11), 2167–2179.
Halmshaw, R., 1991. Non-Destructive Testing, 2nd ed. Edward Arnold, London.
Hananouchi, T., Nakamura, N., Kakimoto, A., Yohsikawa, H., Sugano, N., 2008. CT-based planning of a single-radius femoral component in total knee arthroplasty using the ROBODOC system. Comput. Aided Surg. 13 (1), 23–29.
Heimann, T., Meinzer, H.-P., 2009. Statistical shape models for 3D medical image segmentation: a review. Med. Image Anal. 13 (August (4)), 543–563.
Holmes, C., Drinkwater, B.W., Wilcox, P.D., 2005. Post-processing of the full matrix of ultrasonic transmit–receive array data for non-destructive evaluation. NDT E Int. 38 (8), 701–711.
Hüfner, T., Meller, R., Kendoff, D., Zeichen, J., Zelle, B.A., Fu, F.H., Krettek, C., 2005. The role of navigation in knee surgery and evaluation of three-dimensional knee kinematics. Oper. Tech. Orthop. 15 (January (1)), 64–69.
Huang, N.F.R., Dowsey, M.M., Ee, E., Stoney, J.D., Babazadeh, S., Choong, P.F., 2012. Coronal alignment correlates with outcome after total knee arthroplasty: five-year follow-up of a randomized controlled trial. J. Arthroplasty 27 (October (9)), 1737–1741.
Hunter, A.J., Drinkwater, B.W., Wilcox, P.D., 2008. The wavenumber algorithm for full-matrix imaging using an ultrasonic array. IEEE Trans. Ultrason. Ferroelectr. Freq. Control 55 (November (11)), 2450–2462.
Jensen, J.A., Nikolov, S.I., Gammelmark, K.L., Pedersen, M.H., 2006. Synthetic aperture ultrasound imaging. Ultrasonics 44 (December (Suppl. 1)), e5–15.
Karaman, M., Li, P., O'Donnell, M., 1995. Synthetic aperture imaging for small scale systems. IEEE Trans. Ultrason. Ferroelectr. Freq. Control 42 (3), 429–442.
Kerr, W., Pierce, S.G., Rowe, P., 2016. Investigation of synthetic aperture methods in ultrasound surface imaging using elementary surface types. Ultrasonics 72, 165–176.
KUKA, 2010. KUKA System Software 5.
Lambert, J., Antoine, P., Gens, G., Bimbard, F., Lacassagne, L., Iakovleva, E., 2012. Performance evaluation of total focusing method on GPP and GPU. 2012 Conference on Design and Architectures for Signal and Image Processing (DASIP), 1–8.
Lines, D.I.A., Pettigrew, I.G., Kirk, K.J., Cochran, S., Skamstad, J.A., 2006. Rapid distributed data collection and processing with arrays – the next step beyond full waveform capture. 9th Jt. FAA/DoD/NASA Aging Aircr. Conf., 1–11.
Lonner, J.H., Smith, J.R., Picard, F., Hamlin, B., Rowe, P.J., Riches, P.E., 2015. High degree of accuracy of a novel image-free handheld robot for unicondylar knee arthroplasty in a cadaveric study. Clin. Orthop. Relat. Res. 473 (1), 206–212.
Mcgilp, A., Dziewierz, J., Lardner, T., Gachagan, A., Bird, C., 2014. Inspection of Complex Components using 2D Arrays and TFM. BINDT.
Nikolov, S., Gammelmark, K., Jensen, J.A., 1999. Recursive ultrasound imaging. 1999 IEEE Ultrason. Symp. Proceedings Int. Symp. (Cat. No. 99CH37027) 2, 1621–1625.
O'Donnell, M., Thomas, L.J., 1992. Efficient synthetic aperture imaging from a circular aperture with possible application to catheter based imaging. IEEE Trans Ultrason Ferroelectr Freq Control. 39 (3), 366–380.
Prionas, N.D., Ray, S., Boone, J.M., 2011. Volume assessment accuracy in computed tomography: a phantom study. J. Appl. Clin. Med. Phys. 11 (2).
Rajna, T., Herold, F., Baylard, C., 2009. Reverse engineering process of cryostat components of Wendelstein 7-X. Fusion Eng. Des. 84 (7–11), 1558–1561.
Roche, M., O'Loughlin, P.F., Kendoff, D., Musahl, V., Pearle, A.D., 2009. Robotic arm-assisted unicompartmental knee arthroplasty: preoperative planning and surgical technique. Am. J. Orthop. (Belle Mead, NJ) 38 (February (2 Suppl.)), 10–15.
Romero, D., Mart, O., Mart, C.J., 2009. Using GPUs for beamforming acceleration on SAFT imaging. 2009 IEEE International Ultrasonics Symposium Proceedings, 1334–1337.
Siebert, W., Mai, S., Kober, R., Heeckt, P.F., 2002. Technique and first clinical results of robot-assisted total knee replacement. Knee 9 (September (3)), 173–180.
Smith, J.R., Riches, P.E., Rowe, P.J., 2014. Accuracy of a freehand sculpting tool for unicondylar knee replacement. Int. J. Med. Rob. 10 (August), 162–169.
Talib, H., Rajamani, K., Kowal, J., Nolte, L.-P., Styner, M., Ballester, M.A.G., 2005. A comparison study assessing the feasibility of ultrasound-initialized deformable bone models. Comput. Aided Surg. 10 (5–6), 293–299.
Weston, M., Mudge, P., Davis, C., Peyton, A., 2012. Time efficient auto-focussing algorithms for ultrasonic inspection of dual-layered media using full matrix capture. NDT E Int. 47, 43–50.
Wilcox, P.D., Holmes, C., Drinkwater, B.W., 2007. Advanced reflector characterization with ultrasonic phased arrays in NDE applications. IEEE Trans. Ultrason. Ferroelectr. Freq. Control 54 (August (8)), 1541–1550.
Woolhead, G.M., Donovan, J.L., Dieppe, P.A., 2005. Outcomes of total knee replacement: a qualitative study. Rheumatology (Oxford) 44 (August (8)), 1032–1037.
Zhang, Jie, Drinkwater, B.W., Wilcox, P.D., 2013. Comparison of ultrasonic array imaging algorithms for nondestructive evaluation. IEEE Trans. Ultrason. Ferroelectr. Freq. Control 60 (August (8)), 1732–1745.