

promoting access to White Rose research papers



Universities of Leeds, Sheffield and York
<http://eprints.whiterose.ac.uk/>

This is the author's version of an article published in **Medical & Biological Engineering & Computing**

White Rose Research Online URL for this paper:

<http://eprints.whiterose.ac.uk/id/eprint/75786>

Published article:

Magee, D, Zhu, Y, Ratnalingam, R, Gardner, P and Kessel, D (2007) *An augmented reality simulator for ultrasound guided needle placement training*. Medical & Biological Engineering & Computing, 45 (10). 957 - 967 . ISSN 0140-0118

<http://dx.doi.org/10.1007/s11517-007-0231-9>

An Augmented Reality Simulator for Ultrasound Guided Needle Placement Training

D. Magee¹, Y. Zhu¹, R. Ratnalingam², P. Gardner³, and D. Kessel²

Medical & Biological Engineering & Computing 45(10):957-967 Oct 2007

¹*School of Computing, University of Leeds, Leeds, UK*

²*School of Psychology, University of Leeds, Leeds, UK*

³*Leeds Teaching Hospitals NHS Trust, Leeds, UK*

Tel: +44 113 3436819

Fax: +44 113 3435468

d.r.magee@leeds.ac.uk

<http://www.comp.leeds.ac.uk/vision>

Abstract

Details of a low cost augmented-reality system for the simulation of ultrasound guided needle insertion procedures (tissue biopsy, abscess drainage, nephrostomy etc.) for interventional radiology education is presented. The system comprises physical elements; a mannequin, an ultrasound probe and a needle, and software elements; generating virtual ultrasound anatomy and allowing data collection. These two elements are linked by a pair of magnetic 3D position sensors. Virtual anatomy, for image generation, is generated based on full body CT scans of live humans. Details of the novel aspects of this system are presented including; image generation, registration and calibration.

Ultrasound, simulation, needle-placement, Augmented-reality

Introduction

Accurate ultrasound guided needle placement is a core element of many diagnostic and therapeutic procedures used in modern medicine. Ultrasound allows the tip of a needle to be seen during insertion permitting precise positioning even in small targets deep in the body. Evidence that targeted needle placement is safer has led the National Institute for Clinical Excellence (NICE) to recommend that ultrasound should be used during every elective central venous line insertion. Ultrasound is routinely used in every radiology department during deep tissue biopsy, abscess drainage and nephrostomy. Targeted needle placement requires distinct psychomotor abilities and good hand eye coordination. At present, the skills required for this technique are learnt by the classical method of apprenticeship based training [1]. Currently, there are no existing standards that must be met by a radiologist to practice this technique safely. In addition, there are no means by

which to measure interventional ultrasound performance objectively. As the number of radiology trainees increases the number of training opportunities reduces. [2] compares the training of pilots to surgeons; “It is frightening to speculate the level of mortality that might ensue if pilots received the same level of training that surgeons are currently subjected to in learning new surgical procedures.”

Simulators could help reduce both of these problems. One solution is the use of a physical phantom which can be scanned and allow needle placement (e.g. breast phantoms suitable for biopsy training from SIEL Ltd.). These require the use of an ultrasound machine which is a valuable and expensive clinical resource. An alternative approach is the use of virtual or augmented reality computer based simulations. These have the potential advantage of variable anatomy and automated feedback on performance. In augmented reality systems part of the simulation is physical this improves the correspondence between the simulation and the actual procedure. Our prototype system is based around a body (foam filled mannequin), a defunct ultrasound probe and a needle. These are the key physical objects with which the operator would interact during ultrasound guided procedures. The simulation is augmented with computer generated “virtual” ultrasound imagery which mimics the physical process of actually scanning the patient..

Computer based clinical simulators have been under development for over ten years. However, the realism and purpose of these simulators varies from one instance to another. Many simulators are entirely based around a conventional computer setup (CPU, keyboard, screen, mouse [3, 4]). These do not aim to recreate the entire procedure, rather certain visual aspects of the procedure, such as the radioactive seed implantation modelled by [5]. Such simulators have quite a wide scope for use in planning of needle insertion [6] and other procedures, but are of limited use for training beyond the early stages. More complex setups, based around virtual environments, haptic feedback devices and/or mannequins, can be used to provide more realistic experience for training. The principal success of simulation for invasive procedures has been in endoscopy related procedures (e.g. [7] [8], and various commercial products). This is primarily due to the original procedure using video display technology as the primary visual feedback, and the physical interface being reasonably simple. To quote [8]; “The benefits of endoscopic training through computer simulation, rather than the traditionally performed one-to-one apprenticeship schemes, are now well accepted in the community” (although supervised practice is still an essential part of training). Such training has been successfully validated [9, 10]. Similarly simulators for various teleoperated procedures, such as eye procedures [11], and arthroscopy (including haptic feedback elements) [12-14] are available. Simulation of more complex procedures (in terms of the perceptual environment) is still an active research topic. Holbrey *et al.* [15] present a system based on a haptic feedback device, stereo display, and Finite Element Models for vascular suturing (stitching) training. Similarly, Bro-Nielson *et al.* [16] use a horizontally orientated monitor, Finite Element Models, and a haptic feedback device for trauma surgery (removal of a damaged kidney). A variety of simulators exist for a range of minimally invasive procedures. Gorman *et al.* [17] present a system incorporating a mannequin and a haptic feedback device for the simulation of lumbar (back/spine) puncture procedure. This procedure

relies on tactile feedback and offline CT data, and so has no real-time visual feedback (as part of the realistic part of the simulation). Simulators for fluoroscopy (real-time, or near real-time, X-ray based 2D imaging) guided procedures [1] [18] (see also PERC Mentor: www.simbionix.com) are closer in nature to the ultrasound guided simulation we present. However, these differ fundamentally in the way the imaging is steered (or not) physically. Direction of imaging is far less important with fluoroscopy than with ultrasound due to the absorbency/reflection properties of ultrasound. However, ultrasound is inherently safe, where as exposure to X-ray should be minimised. The commercial product UltraSim (www.medsim.com) simulates ultrasound imaging without any interventional procedure. This simplifies the implementation somewhat. UltraSim is used for diagnosis training (rather than interventional training). DiMaio and Salcudean [19] present a fast Finite Element model for the haptic modelling of needle insertion that could be added to our system (see discussion). We do not know of any simulator for ultrasound guided needle insertion other than the 2D computer only simulation mentioned previously [5], which does not model probe motion or any 3D aspect of the task.

Simulator Implementation

The hardware of our simulation system comprises three components, a standard PC (Dual Intel Xeon™ CPU 3.20 GHz, 2.0 GB of RAM, running Microsoft Windows™ Professional), a full scale penetrable model made of latex plastic and foam, and a pair of Ascension PCIBirds magnetic 3D position/orientation sensors. One of the sensors is rigidly attached to an ultrasound probe, and the other attached to a standard biopsy needle. Figure 1 illustrates the hardware setup of our simulator.



Figure 1: Hardware Setup for Simulator

The production of a synthetic ultrasound image involves 2 fundamental processes each comprising a number of specific tasks.

1. Firstly, anatomically correct data must be mapped to the mannequin
 - a) A volumetric CT data set is obtained from a human subject.

- b) The CT volume is manually segmented, with each voxel assigned a label that corresponds to the physical and anatomical properties of tissues within the voxel. For example liver, blood vessel, bowel gas and bone.
- c) The volume is automatically registered to the surface of the physical model.
- d) A bank of volumetric textures is constructed from real ultrasound images to represent the typical appearance of the objects in the segmented volumetric data.

2. Secondly, this data is used to construct the virtual ultrasound image

- a) The motion sensors are calibrated so that their position and orientation map exactly to the physical mannequin and thus to the virtual dataset.
- b) The PCI Bird sensor attached to the “ultrasound” probe gives precise spatial information which is used to derive a virtual scan plane correctly oriented to the probe.
- c) The pixels in the ultrasound plane are raster scanned, and each point projected to the corresponding point in 3D space
- d) A grey-level is assigned to each pixel according to the property of the tissue determined in 1b.
- e) In addition, a 2D ray-casting method is applied to resemble the acoustic shadowing effect of reflective materials such as bone and bowel gas.
- f) The synthetic image is then rendered using an alpha blending routine to simulate the radial blur effect.

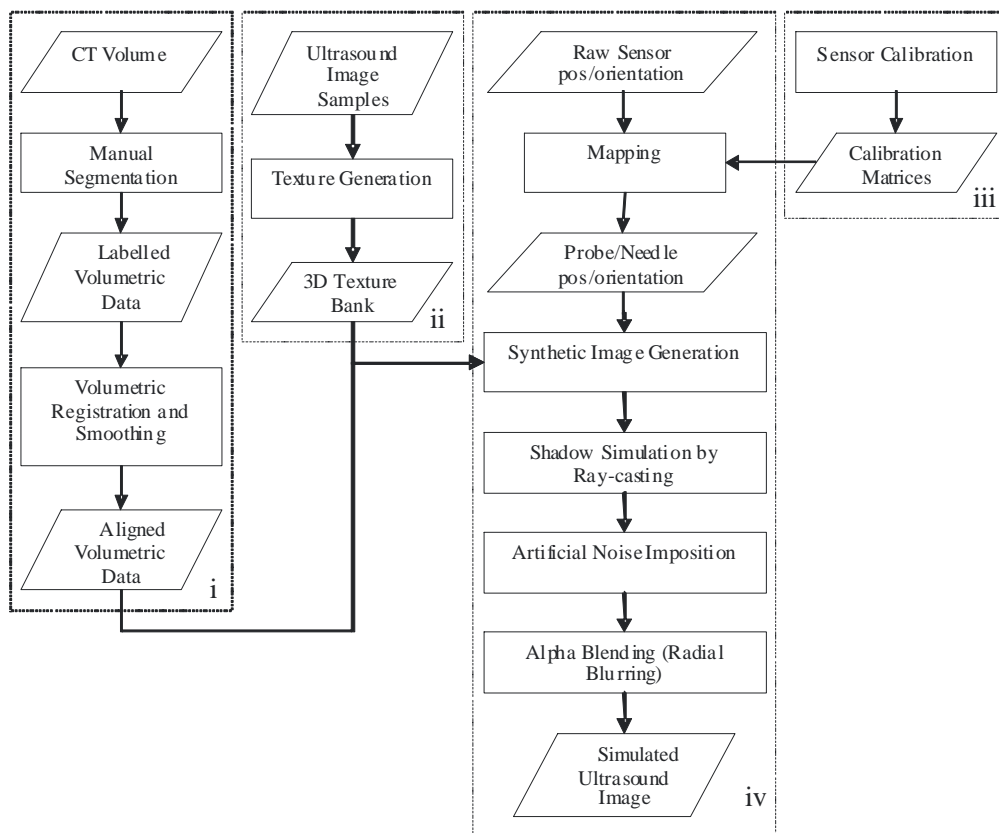


Figure 2: A Diagrammatic Overview of the System

The following sections describe this process in more detail.

Registration of Volumetric Data to Mannequin Surface

As stated previously, the simulator renders “virtual ultrasound” images based on volumetric data acquired from real humans. This data is registered to the surface of the mannequin in a 3 stage process. Firstly, surfaces are extracted from the volumetric data and mannequin as point clouds. The surface on the volumetric data is extracted by 2D hand annotation of slices, and the surface of the mannequin is extracted by passing a (calibrated) motion sensor over the surface of the mannequin.

The second stage of registration is a rigid alignment of the two data sets. To do this we use the Iterative Closest Point Algorithm (ICP) [20] with multiple (randomly selected) starting offsets and rotations. The result with the lowest error between surfaces is selected.

The ICP based rigid registration algorithm described aligns the volumetric and mannequin surfaces reasonably well. However, for the surfaces to be completely aligned a non-rigid registration process is also necessary due to differences in the shape of the surfaces (particularly around the belly). We formulate this non-rigid registration as the estimation of a quadratic mapping function that maps the volumetric surface on to the mannequin surface as in equation 3.

$$\begin{pmatrix} x_{new} \\ y_{new} \\ z_{new} \end{pmatrix} = K_c + K_x \begin{pmatrix} x \\ y \\ z \end{pmatrix} + K_{xx} \begin{pmatrix} x^2 \\ y^2 \\ z^2 \end{pmatrix} + K_{xy} \begin{pmatrix} xy \\ xz \\ yz \end{pmatrix} \quad (3)$$

Where K_c is a 3×1 matrix, and K_x , K_{xx} and K_{xy} are 3×3 matrices. Equation 3 defines a 3D warp that may be applied to volumetric data. If a set of point correspondences ($\gg 10$ correspondences) are known, the values in K_c , K_x , K_{xx} and K_{xy} may be estimated using a least-squares approach (applied separately for each dimension (x,y,z)). However, we desire a smooth mapping from a (relatively) small number of correspondences. To achieve this we add some smoothness constraints to the least squares set of equations. For the x dimension mapping:

$$\frac{dx_{new}}{dx} = 1, \frac{dx_{new}}{dy} = 0, \frac{dx_{new}}{dz} = 0 \quad (4)$$

and for the y dimension mapping:

$$\frac{dy_{new}}{dx} = 0, \frac{dy_{new}}{dy} = 1, \frac{dy_{new}}{dz} = 0 \quad (5)$$

and for the z dimension mapping:

$$\frac{dz_{new}}{dx} = 0, \frac{dz_{new}}{dy} = 0, \frac{dz_{new}}{dz} = 1 \quad (6)$$

These constraints are selected to force the transform to be as close as possible to an unrotated rigid mapping. The constraints are implemented as soft constraints by including 3 equations relating to each volumetric data point in the set of equations presented to the least-squares solver (i.e. 4 simultaneous equations per point are used, one for the mapping and 3 for the constraints).

To obtain the pairs of correspondences required for the non-rigid estimation process described we use an approach based on the RanSaC algorithm [21]. This commonly used algorithm in model fitting and 3D geometry is based on a three part sample, estimate and verify approach. A small set of correspondences is sampled from the set of possible correspondences based on some probabilistic matching criteria. This set is then used to estimate the model parameters (the quadratic function in our case). The model is then verified by application to the entire data set using a ‘‘Consensus function’’. The process is repeated many times, and the model with the highest ‘‘consensus’’ selected. The model may then be re-estimated based on a larger set of correspondences calculated using the selected model. Our RanSaC correspondence sampling probability density is based on the Euclidean distance between potential correspondences normalised by the minimum distance between the volume point and any point on the mannequin surface, as in equation 7.

$$D_{match}(v, m) = \frac{E_d(v, m)}{E_d(v, C_m(v))} \quad (7)$$

Where $E_d(x, y)$ is the euclidean distance between points x and y . A $N_v \times N_m$ surface point similarity matrix is formed based on equation 7, which is binarised using a threshold (T_d , typically 1.3), as in equation 8.

$$P_{match}(v, m) = \begin{cases} 1 & \text{if } D_{match}(v, m) < T_d \\ 0 & \text{if } D_{match}(v, m) \geq T_d \end{cases} \quad (8)$$

This excludes a large proportion of possible poorly matching correspondences. Around 20-30 correspondences per trial are selected randomly from the non-excluded (true) matches, in order to calculate the mapping (equation 3). Typically 100 trials are performed in order to select the best mapping.

The consensus function (C) used to select the final mapping from the set of candidate mappings is based on the proportion of transformed volume surface points that are within a fixed threshold of a mannequin point (P_t).

$$C = \frac{\# \text{points where } |X_{nr_i} - C_m(X_{nr_i})| < T_c}{N_v} \quad (9)$$

Where X_{nr_i} is the i 'th non-rigidly registered volume point, and T_c is a threshold (typically 0.05 times the mean distance of the mannequin points from the mannequin centroid). The solution with the highest value for C is selected, and the mapping re-calculated based on a larger set of correspondences made up of the closest mannequin surface matches to the transformed volume surface points where the distance between correspondences is below the threshold T_c . In practice this (generally) has the effect of increasing the consensus

value (C) for the final solution (if not the original mapping is used). Typical results of this procedure are given in figure 3.




Figure 3: a) Rigid, and b) Non-rigid registration results

The quadratic function is used to warp the volumetric data offline to a new volume (with arbitrarily sized cubic voxels). The warp is performed by calculating the corresponding position in the original volume to each voxel in the new volume using i) a fixed scaling and offset derived from the choice of voxel size/range in the new volume, ii) The inverse of the quadratic function, and iii) the inverse rigid mapping. To calculate ii) we initially calculate a quadratic approximate inverse by swapping the (rigidly-aligned) volume and mannequin surface point sets in the last stage of the non-rigid estimation process. The output of this function is then used as an initial value for the calculation of an accurate inverse of the forward function using Newtons method (typically only 2 or 3 iterations are required). Once the warped volume has been formed on-line calibration is simply matter of a rigid scale and offset between the volume co-ordinate system and the world co-ordinate system.

Volume Smoothing

The nature of CT (and MRI) volumes is that the between plane resolution is rather lower than the in plane resolution. When such data (or a segmentation based on this data) is warped to a cubic voxel volume a step-like effect can sometimes be observed if the new volume is sliced in a non-axial direction (figure 4a). The solution to this is to smooth the segmentation using a Level-Set method [22]. Level-set methods implicitly model a surface as the zero valued iso-surface (“Level-set”) of a signed distance transform defined on a uniform (i.e cubic in 3D) grid. Grid values inside the surface are given negative values, and grid values outside the surface are given positive values. The evolution of this curve is controlled by a speed function defined at each grid point. We use a speed function consisting of a smoothness term, and a term to ensure similarity to the original segmentation, as in equation 10.

$$\varphi_{t+1}(x, y, z) = (K_0(\varphi_0(x, y, z) - \varphi_t(x, y, z)) + C(x', y', z')) \nabla_t \quad (10)$$

Where $\varphi_t(x, y, z)$ is the grid value at position (x, y, z) at timestep t , $\varphi_0(x, y, z)$ is the original value of the level set (initialised from the volume segmentation using the method described in [23]), $C(x', y', z')$ is the Mean curvature [22] at the closest level set point to (x, y, z) , and K_0 is a constant used to set how closely the final solution relates to the initial segmentation. This process is carried out for each segmentation layer (anatomical structure), and a new binary segmentation generated for each layer by thresholding the final value of $\varphi_t(x, y, z)$ with a threshold of 0 (and inverting). Figure 4 illustrates this process.




Figure 4: A non-axial slice rendered a) Without Level-set segmentation smoothing, b) With Level-set segmentation smoothing (Note: different anatomical layers are mapped with different textures for clarity. No other processing is applied)

Calibration of Probe and Needle

The motion sensors used output 3D position, orientation and rotation information. Due to the practical constraints on sensor placement, it is not always possible to place the sensor at the position that we want to measure (e.g. the centre of the end of the probe), or orient it in the direction we want to measure. However, it is possible to rigidly attach the sensors and estimate a calibration function to map the sensor values to the desired values. The raw output of the sensors may be visualised as a point and two vectors as shown in figure 5.




Figure 5: Raw output of 3D position sensor

Probe position calibration is performed as the estimation of an offset from S_{pos} in a frame of reference defined by S_{dir} and S_{ori} . This is in fact a 3 degree of freedom estimation problem, thus we introduce the variable θ , an angle used to rotate the S_{ori} vector about S_{dir} (thus allowing the calibration point to be off the plane defined by S_{ori} and S_{dir}). Equation 11 relates the sensor outputs to the calibration point (C_p).

$$C_p = S_{pos} + d_1 S_{dir} + d_2 R_\theta(S_{ori}) \quad (11)$$

Where $R_\theta()$ is a rotation of θ , about S_{dir} , as in equation 12 (effectively defining a third axis using the cross product).

$$R_\theta \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} x \cos \theta + (S_{dir_y} z - S_{dir_z} y) \sin \theta \\ y \cos \theta + (S_{dir_z} x - S_{dir_x} z) \sin \theta \\ z \cos \theta + (S_{dir_x} y - S_{dir_y} x) \sin \theta \end{pmatrix} \quad (12)$$

Substituting equation 12 into equation 11, and rearranging gives:

$$\begin{pmatrix} 1 & 0 & 0 & -S_{pos_x} & -S_{ori_x} & -\begin{pmatrix} S_{dir_y} S_{ori_z} - S_{dir_z} S_{ori_y} \end{pmatrix} \\ 0 & 1 & 0 & -S_{pos_y} & -S_{ori_y} & -\begin{pmatrix} S_{dir_z} S_{ori_x} - S_{dir_x} S_{ori_z} \end{pmatrix} \\ 0 & 0 & 1 & -S_{pos_z} & -S_{ori_z} & -\begin{pmatrix} S_{dir_x} S_{ori_y} - S_{dir_y} S_{ori_x} \end{pmatrix} \end{pmatrix} \begin{pmatrix} C_{p_x} \\ C_{p_y} \\ C_{p_z} \\ d_1 \\ d_2 \cos \theta \\ d_2 \sin \theta \end{pmatrix} = \begin{pmatrix} S_{pos_x} \\ S_{pos_y} \\ S_{pos_z} \end{pmatrix} \quad (13)$$

Equation 13 is in a form where the unknowns ($C_p, d_1, d_2 \cos \theta$, and $d_2 \sin \theta$) may be estimated from multiple unique examples using a standard least-squares solver (e.g. the *dgels* routine in the LAPACK library). To obtain a set of values for the least squares solver a set of (around 30) samples are taken with the calibration point resting on a fixed point, as illustrated in figure 6a.




Figure 6: a) Probe calibration setup, b) Needle Calibration setup

The two probe directions C_{dir} and C_{ori} (analogous to S_{dir} and S_{ori}) are estimated by moving the probe in these directions, taking samples at approximately uniform spatial intervals. Each pair of samples gives an estimate of C_{dir} (or C_{ori}) in world co-ordinates

(V_n , a unit direction vector). This may be projected into the co-ordinate system defined by S_{dir} , S_{ori} , and S_3 (the cross product of S_{dir} and S_{ori}) for the first of these samples. We perform this projection using the least-squares solution to the system of equations given by:

$$V_n = (K_d, K_p, K_3) (S_{dir_n}, S_{ori_n}, S_{3_n})^T \quad (14)$$

Where K_d , K_p and K_3 are scalar unknowns that define C_{dir} or C_{ori} w.r.t. the co-ordinate system defined by vectors S_{dir_n} , S_{ori_n} and S_{3_n} .

Needle calibration is initially based on the same mathematical approach as for the probe position calibration. Figure 6b illustrates the data acquisition in this case, using a container with a fixed pinhole entry point. The procedure outlined previously is performed multiple times (typically 5) with different values of D_n . These calibrations are projected into the 3D space defined by S_{dir_n} , S_{ori_n} and S_{3_n} , and Principal Components Analysis performed to extract the eigenvector with the highest variance. This relates to the needle direction (analogous to C_{dir} in the probe case). Each calibration is projected onto this principal eigenvector (using the dot product of the calibration and the vector), and the resulting vector (projection \times eigenvector) subtracted from each calibration. The mean of these resultant calibrations is taken as the final needle calibration. N.B. The rotation of the needle is not relevant for this application.

Generation of Synthetic Ultrasound Images

Generation of the ultrasound images is a 3 stage process. Firstly, a raw textured image is generated by raster scanning through the ultrasound image portion of the display image and projecting each pixel into 3D world co-ordinates (using the probe sensor output and calibration). The “layer” (anatomical structure / material property) to which that pixel belongs is determined as a nearest-neighbour look-up into the registered segmentation. The greyscale value of the pixel is determined based on a look-up into one of a set of tessellating “solid textures” (one for each “layer”). As these solid textures tessellate in all directions, they may be assumed to be infinite textures. The look-up location is simply a matter of Position MOD Texture size, for each of the 3 dimensions. The second step is to render the needle (if appropriate). The third step is to apply image transforms to the rendered image to simulate ultrasound-specific artefacts, such as noise, radial distortion, shadowing and attenuation. Details of the image formation process are given in the following sections.

Construction of Texture Database

A set of volumetric “solid textures” [24, 25] (one for each visually distinct tissue type) is learned from sets of hand selected 2D ultrasound examples. Solid textures are volumes for which the surfaces appear similarly textured if cut by any plane (or other surface). There is a large body of literature on the generation of synthetic 2D textures, but much less on solid texture generation. This is probably due to the fact that they are not currently as widely used for a number of reasons including; Consumer graphics hardware architectures, higher memory requirements than 2D textures, and the comparative

difficulty in making volume textures anisotropic. However, they are ideally suited to our application as 2D textures cannot model the effect of sensor motion in three dimensions. 2D synthetic texture methods fall broadly into two categories; i) Cut-and-paste based methods (e.g. [26]), in which patches of training images are pasted together to form a new texture, and ii) Image statistics based methods (e.g. [27]). We have found approach i) does not easily extend to learning 3D solid textures from 2D examples. Our approach is to use a Laplacian image pyramid [27] (with 3 or 4 levels) to represent the image and volume statistics, and to match the statistics of the volume to the training images. A Laplacian image pyramid is a set of images (or volumes) generated from a single image by iteratively applying a low pass filter and subtracting the result from the previous image generated. This is a form of frequency decomposition. We use a 256 bin greylevel histogram at each level of the Laplacian pyramid to capture the training data statistics ($P_l(g)$). We also form two dimensional histograms to capture the joint greylevel probability distribution of each level and the level below ($P_{l,l-1}(g_l, g_{l-1})$). To generate the solid texture we start at the lowest (most coarse) level of the Laplacian pyramid and sample greylevel values from $P_l(g)$. Subsequent levels of the pyramid are generated based on a weighted combination of $P_l(g)$ and $P_{l,l-1}(g_l, g_{l-1})$:

$$P(g, l) = K_l P_l(g) + (1 - K_l) P_{l,l-1}(g_l, g_{l-1}) \quad (15)$$

Where K_l is typically 0.6 for all levels except the top (highest frequency) level, where it is 0. The use of $P_l(g)$ compensates for the data sparsity at the lower levels which can lead to textures that are too regular. This is unnecessary at the top (highest frequency) level as sufficient data is available. A tessellating solid texture (typically 32x32x32) is generated from the Laplacian volume pyramid generated by the sampling process by use of a low-pass filter that assumes 3D wrapping at the edges of the volume. This is a very fast process as it does not rely on searching through the training set (e.g. [28]) or iterative optimisation (e.g. [27]).

Simulation of Ultrasound-Specific Artefacts

With the generated textures, an initial ultrasound image can be generated by raster scanning through the pixels in the ultrasound portion of the scan plane, and projecting each pixel to the corresponding voxel in the 3D volume space so that its label can be determined. The pixel's label and its 3D coordinates are then used to sample the corresponding texture to determine the greylevel of the pixel. An example of the intermediate result at this stage is shown in Figure 7(a). However, this image is not exactly like real ultrasound output due to the lack of ultrasound-specific artefacts, such as speckle noise, shadows, radial motion blurring, and attenuation. The speckle effect can be simply simulated by adding Gaussian distributed artificial noise to the image pixels (We currently use $\sigma=15$ for the Gaussian noise). Since ultrasound pulses do not penetrate bones or air, a shadow will be cast behind these objects if they present in the scan plane. We simulate the shadowing effect by a 2D ray-casting approach. Regarding the centre of the end of the ultrasound probe as a spot light, we compute the greylevel decrement of each image pixel according to the number of bone pixels (or air pixels) along the straight line segment that connects the spot light and the pixel; the more bone pixels appear along the line segment, the darker the pixel is. In fact, note that the ultrasound image is

generated by a top-bottom raster scanning process, both the speckle and shadowing effects can be produced in one scan of the full image, and hence maintain a satisfactory frame rate in real-time interactive mode.

The radial blur effect creates blurs around a specific point in an image, simulating the effects of a rotating camera. This effect can be used to simulate the radial scanning motion of a real ultrasound transducer. However, the traditional method that computes the blurred image by a convolution operation is too computationally expensive to be applied in real-time. Therefore, we have made use of the power of modern graphics cards to achieve the same effect. The traditional blurring convolution can be viewed as a weighted sum process that adds up the original image and a number of its weighted copies with certain offsets (or rotations). This can be performed by the alpha blending function that is efficiently implemented by most modern graphics cards. We draw the generated image many times to create the radial blur effect, varying the angle of rotation about the blurring centre, and increasing the alpha level while the angle increases. Figure 7 (b) is an illustration of the alpha blending process, and an example of the final ultrasound image generated by our simulator is given in Figure 7(c). Intensity attenuation is simply modelled as a brightness scaling, the value of which is a decreasing function of the distance from the probe. It is worth emphasising that, although consisting of a series of tasks, the image generation and rendering process is still highly computationally efficient. The average frame rate of our ultrasound simulator is around 12-13fps, which is very close to the 16fps rate that is commonly used in clinical radiology.




Figure 7: a) Image before ultrasound specific artefacts added, b) alpha blending process, c) Image with ultrasound specific artefacts added

Needle Rendering

After the textured ultrasound image is constructed (but before the ultrasound specific artefacts are applied), the needle is rendered onto the image as a bright line. This is done by raster scanning through the ultrasound portion of the image and rendering with a high intensity pixels falling within the geometry of the needle (modelled as a cylinder). To produce realistic rendering of the needle, the ultrasound plane is assigned a thickness, t , so that the transparency of the rendered needle is determined by the intersection volume of the needle with the thickened ultrasound plane. This models partial volume effects present in ultrasound. A value of t that is equal to the needle diameter produces satisfactory needle appearance in our simulator. However, the thickness can be adjusted to control the difficulty level during training and practicing (e.g., a thinner plane

ultrasound plane requires more proficient skills to keep the needle visible). Other effects, such as varying the transparency away from the needle tip have also been implemented.

It is natural that, in real clinics, radiologist may tend to bend the needle to a small reasonable extent to correct the needle direction in case that it is placed slightly off the target. In our simulation system, however, since the sensor is attached to the needle handle, and the needle is modelled as a rigid object, such attempt to bend the needle may result in a large displacement of the needle in the generated image. To tackle this problem, a constraining strategy is applied such that the constrained needle direction is a linear combination of the measured needle direction, and the direction at the time when the needle breaks the skin, where the combination coefficients are determined by the inserted depth of the needle tip. Thus, the deeper the needle is inserted, the less displacement of the needle caused by bending the actual needle may be noticed in the image (equation 16).

$$d_c = wd_e + (1 - w)d_m \quad (16)$$

Where d_m is the measured needle direction vector, d_e is the entry direction vector, and w is defined as:

$$w = w_{\max} \frac{1}{1 + e^{-\alpha(l-\delta)}} \quad (17)$$

Where l is the depth to which the needle is inserted, and α and δ are constants determined experimentally by a consultant radiologist to give maximum realism (typically we use $\alpha=60$ and $\delta=0.12$).

Evaluation

A simulator is only of value if it can be shown to fulfil its proposed function, this requires a series of experiments to validate the simulator. To date our simulator has undergone the preliminary stages of the validation process. In the initial stage 60 consultant interventional radiologists and specialist registrars in radiology trialled the simulator performing at least two needle passes. Participants completed a questionnaire detailing their opinion on the role of this simulator in skills training and testing for ultrasound guided procedures. Overall opinion about the simulator is favourable especially on the critical question of whether the simulator accurately replicates the technique of needle placement under ultrasound guidance (Face Validity). High positive scores were also achieved on items such as the anatomical position of organs relative to each other, overall hardware reality and ergonomics. The visual display realism was less strongly positive, although it should be noted that we have significantly improved the realism since this evaluation (see the section on Simulation of Ultrasound-Specific Artefacts). The reality of tactile feedback is considered not very realistic but this is not surprising as there has been no attempt to include tactile feedback into the simulation. Respondees indicated that the simulator was likely to be a successful and cost effective tool for all stages of medical training including trainee radiologists, junior doctors, medical students and other staff in the technique of ultrasound guided needle placement. The simulator was also considered to have potential to reduce error and enhance skills.

Discussion and Future Work

We have presented various aspects of a system for training in ultrasound guided needle insertion. In addition to the functionality already discussed, the simulator has the ability to record and review sessions (using a number of visualisation techniques). Various task success metrics have been developed to give detailed feedback. Positional targets for needle placement (both realistic and highly visible) can be placed, and performance measured with respect to these. Our initial face validation, using a large group of experts, suggests that the simulator is sufficiently realistic to have potential in the teaching of interventional radiology. One can envisage trainees developing skills in a controlled and staged fashion in a stress free environment [29]. This can never be a complete replacement for supervised practice on real patients due to the remaining differences between the simulator and real life. However, it has been shown that simulators do not need to be accurate in every respect to be useful. This is well illustrated in validation studies of the MIST-VR which has been proven to improve subsequent performance of laparoscopic surgery in the operating room [9]. For these reasons the simulator has been developed pragmatically as a low cost system solely to teach the core skill of needle placement under ultrasound guidance. The simulator is designed to help the operator learn a subset of skills key to performing accurate needle placement such as initial pre-procedure safety assessment and route planning, needle-probe alignment with constant needle tip visualization and mental 3 dimensional interpretation of 2 dimensional data.

The realism of the synthetic ultrasound image was felt to be secondary to the positional accuracy and close resemblance to the feel of real life ultrasound scanning. Even with this constraint it is possible to provide a wide range of interchangeable abnormalities and anatomy for training. Similarly, it was felt that the incorporation of haptic (force) feedback was unnecessary and would unduly increase cost and possibly affect the ergonomics of the simulator in such a way as to detract from the experience of physically scanning the mannequin. It is therefore not surprising that the simulator scored poorly for tactile feedback. Despite this, both groups felt strongly about its potential as a teaching tool for trainee radiologists and other groups.

Since completion of the evaluation, the simulator has been further refined both visually and functionally. Ultrasound specific artefacts (noise, radial distortion, etc.), altered needle visibility, needle bending modelling, and more realistic targets are now all part of its standard features. The recording and playback functions have greatly improved, allowing the trainer and trainee to review the session from a multitude of different views with more detailed analysis and progress update.

Although simulators have been well received and are considered to have great potential, the uptake of such systems in medical training has been limited. There are 2 reasons for this. Firstly expense, the price and running costs of a large number of commercial simulators are beyond the reach of many hospitals and teaching departments. Secondly, most remain unproven as valid or reliable teaching tools. It is essential that we do not prematurely accept simulator based training solely on the basis of the positive subjective responses that they receive [30]. These new instruments must be properly evaluated

before they are incorporated into medical curricula [31]. Validation studies are crucial tests to determine whether a simulator fulfils its requirements. Validity can be categorized into different types: face, content, construct, criterion-related and concurrent validity [32, 33]. A new simulator should demonstrate multiple aspects of validity before becoming a reliable skills trainer and predictor of performance [33, 34]. Face validation establishes if experts rate the model's performance as accurate. It addresses the question, to what extent the instrument simulates what it is supposed to represent. In other words is performing targeted needle placement on the simulator like performing the actual task. This concentrates on whether or not the inputs and outputs appear realistic or representative [29, 34]. Content validation determines whether the simulator is judged to contain the necessary elements for the training it intends to impart? The value of the teaching instrument can only be assigned if it meets the expectations and requirements of expert operators who are beyond the learning curve of the procedures under study (expert opinion) as well as the potential trainee (referent opinion) [29]. The results of this study demonstrate a favourable and uniform opinion among experts and trainees regarding the face validity and content validity of the simulator. Both groups are very positive about its utility as a training tool particularly for Specialist Registrars in Radiology.

Simulators are not stand alone devices and should be used in the context of a recognized training curriculum to complement existing training methods. To fulfil these training objectives future work will include a structured program, incorporating scenarios of increasing complexity to facilitate stepwise improvement and understanding of the skills of ultrasound guidance. In addition to technical refinement of the hardware and software objectives, we are attempting to define measurable aspects of performance (metrics) which reliably discriminate between operators of different skill levels. This construct validity will help to identify when competence is achieved. Ultimately, for widespread acceptance of the simulator by the public, radiological community and regulating bodies, it must be demonstrated that skills developed on the simulator are transferable to the clinical setting (transfer of training) and will predict future clinical competence. Future work is intended to demonstrate its value in this process.

References

1. Wang, Y., Chui, C., Lim, H., Cai, Y., Mak, K. (1998) Real-Time Interactive Simulator for Percutaneous Coronary Revascularization Procedures. *Computer Aided Surgery* 3: 211-227
2. Tafra, L. (2001) The learning curve and sentinel node biopsy. *The American Journal of Surgery* 182: 347-350
3. Halvorsrud, R. (2000) The MATADOR project – a novel simulator in emergency medicine. *Proc. Medicine Meets Virtual Reality*
4. Kizakevich, P., McCartney, M., Nissman, D., Starko, K. (1998) Virtual medical trainer: patient assessment and trauma care simulator. *Proc. Medicine Meets Virtual Reality* 309–315
5. Alterovitz, R., Pouliot, J., Taschereau, R., Hsu, I., Goldberg, K. (2003) Simulating Needle Insertion and Radioactive Seed Implantation for Prostate Brachytherapy. *Proc. Medicine Meets Virtual Reality* 12 19-25

6. Zeng, J., C.Kaplan, Xuan, J., Sesterhenn, I., Lynch, J., Freedman, M., Mun, S. (1998) Optimizing prostate needle biopsy through 3D simulation. Proc. SPIE Conference on Medical Imaging 488-497
7. Gopalakrishnan, G., Devarajan, V. (2004) StrapSim: A Virtual Reality-based Stapling Simulator for Laparoscopic Herniorrhaphy. Proc. Medicine Meets Virtual Reality 12 111-113
8. ElHelw, M., Chung, A., Darzi, A., Yang, G. (2003) Image-Based Modelling of Soft Tissue Deformation. Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI) 83-90
9. Seymour, N., Gallagher, A., Roman, S., al., e. (2002) Virtual reality training improves operating room performance: results of a randomized, double-blinded study. Annals of Surgery 236: 458-464
10. Datta, V., Mandalia, M., Mackay, S., Darzi, A. (2002) The PreOp flexible sigmoidoscopy trainer – Validation and early evaluation of a virtual reality based system. Surgical Endoscopy 16: 1459 – 1463
11. Sagar, M., Bullivant, D., Mallinson, G., Hunter, P., Hunter, I. (1994) A Virtual Environment and Model of the Eye for Surgical Simulation. Proc. 21st Annual Conf. on Computer Graphics and Interactive Techniques 205-212
12. Moody, L., Waterworth, A., Arthur, J., Zivanovic, A., Dibble, E. (2004) Ensuring the usability of a knee arthroscopy simulator. Medicine Meets Virtual Reality 241-243
13. Smith, S., Wan, A., Taffinder, N., Read, S., Emery, R., Darzi, A. (1999) Early experience and validation work with Procedicus VA – the Prosolvia virtual reality shoulder arthroscopy trainer. Proc. Medicine Meets Virtual Reality 337-343
14. Mabrey, J., Cannon, W., Gillogly, S., Kasser, J., Sweeney, H., Zarins, B., Mevis, H., Garret, W., Poss, R. (2000) Development of a virtual reality arthroscopic knee simulator. Proc. Medicine meets virtual reality 192-194
15. Holbrey, R., Bulpitt, A., Brodlie, K., Walkley, M., Scott, J. (2004) A model for virtual suturing in vascular surgery. Proc. EurographicsUK: Theory and Practice of Computer Graphics 50-58
16. Bro-Nielsen, M., Helfrick, D., Glass, B., Zeng, X., Connacher, H. (1998) VR Simulation of Abdominal Trauma Surgery. Proc. Medicine Meets Virtual Reality 117-123
17. Gorman, P., T.Krummel, Webster, R., Smith, M., Hutchens, D. (2000) A Prototype Haptic Lumbar Puncture Simulator. Proc. Medicine Meets Virtual Reality 106-109
18. Chui, Z., Anderson, J., Ma, X., Peng, W., Cai, Y., Nowinski, Y. (2001) Computer environments for interventional neuroradiology procedures. Simulation and Gaming 32:3: 404-419
19. DiMaio, S., Salcudean, S. (2002) Needle Insertion Modelling for the Interactive Simulation of Percutaneous Procedures. Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI) 253-260
20. Besl, P., McKay, N. (1992) A Method for Registration of 3D Shapes. IEEE Transactions on Pattern Analysis and Machine Intelligence 14(2): 239-256
21. Fischler, M., Bolles, R. (1981) Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. Graphics and Image Processing 24(6): 381-395

22. Sethian, J.A.: Level Set Methods and Fast Marching Methods. Cambridge University Press (1999)
23. Magee, D., Bulpitt, A., Berry, E. (2001) Combining 3D Deformable models and Level Set Methods for the Segmentation of Abdominal Aortic Aneurisms. Proc. British Machine Vision Conference, Vol. 1 333-341
24. Perlin, K. (1985) An Image Synthesizer. Proc. International Conference on Computer Graphics and Interactive Techniques (SIGGRAPH) 287-296
25. Peachey, D. (1985) Solid Texturing of Complex Objects. Proc. International Conference on Computer Graphics and Interactive Techniques (SIGGRAPH) 279-286
26. Efros, A., Freeman, W. (2001) Image Quilting for Texture Synthesis and Transfer. Proc. International Conference on Computer Graphics and Interactive Techniques (SIGGRAPH) 341-346
27. Heeger, D., Bergen, J. (1995) Pyramid-Based Texture Analysis/Synthesis. Proc. International Conference on Computer Graphics and Interactive Techniques (SIGGRAPH) 229-238
28. DeBonet, J. (1997) Multiresolution sampling procedure for analysis and synthesis of texture images. Proc. International Conference on Computer Graphics and Interactive Techniques (SIGGRAPH) 361-368
29. Schijven, M., Jakimowicz, J. (2002) Face-, expert and referent validity of the Xitact LS500 Laparoscopy Simulator. Surgical Endoscopy 16: 1764 – 1770
30. Gould, D., Reekers, J., Kessel, D., Chalmers, N., Sapoval, M., Patel, A., Becker, G., Lee, M., Stockx, L. (2006) Simulation Devices in Interventional Radiology: Caveat Emptor. Cardiovascular Interventional Radiology 29: 4-6
31. Reznek, M., Harter, P., Krummel, T. (2002) Virtual Reality and Simulation: Training the Future Emergency Physician. Academic Emergency Medicine 9: 78-87
32. Gould, D., Kessel, D., Healey, A., Johnson, S., Lewandowski, W. (2006) Simulators in catheter-based interventional radiology: training or computer games. Clinical Radiology 61: 556-561
33. Dawson, S. (2006) Procedural Simulation: A Primer J Vasc Interv Radiol 17: 205-213
34. Schijven, M., Jakimowicz, J. (2003) Construct Validity: Experts and Novices Performing on the Xitact LS500 Laparoscopy Simulator. Surgical Endoscopy 17: 803-810