

Adaptive spatial compounding for improving ultrasound images of the epidural space

Denis Tran^a, Allaudin Kamani^b, Vickie Lessoway^c and Robert N. Rohling^a

^aDepartment of ECE, University of British Columbia, 2332 Main Mall, Vancouver, Canada;

^bDepartment of Anesthesia, University of British Columbia, 4500 Oak Street, Vancouver, Canada;

^cDepartment of Ultrasound, BC Women's Hospital and Health Centre, 4500 Oak Street, Vancouver, Canada

ABSTRACT

Epidural anesthesia can be a difficult procedure, especially for inexperienced physicians. The use of ultrasound imaging can help by depicting the location of the epidural space to choose the needle trajectory appropriately. Anatomical features in the lower back are not always clearly visible because of speckle poor reflection from structures at certain angles, and shadows from bony surfaces. Spatial compounding has the potential to reduce speckle and emphasize structures by averaging a number of images taken at different isonation angles. However, the beam-steered images are not perfectly aligned due to non-constant speed of sound causing refraction errors. This means compounding can blur features. A non-rigid registration method, called warping, shifts each block of pixels of the beam-steered images in order to find the best alignment to the reference image without beam-steering. By applying warping, the features become sharper after compounding. To emphasize features further, edge detection is also applied to the individual images in order to select the best features for compounding. The warping and edge detection parameters are calculated in real-time for each acquired image. In order to reduce computational complexity, linear prediction of the warping vectors is used. The algorithm is tested on a phantom of the lower back with a linear probe. Qualitative comparisons are made among the original plus combinations of compounding, warping, edge detection and linear prediction. The linear gradient and Laplacian of a Gaussian are used to quantitatively assess the visibility of the bone boundaries and ligamentum flavum on the processed images. The results show a significant improvement in quality.

Keywords: Contrast, image processing, image quality/display, registration, spatial compounding

1. INTRODUCTION

1.1. Epidural anesthesia in obstetrics

In obstetrics, a patient may require the use of anesthesia to ease the pain of labour and delivery. For these patients, the anesthesia is injected into the epidural space near the spine as seen on Fig. 1. The epidural space is at a depth of 20-90mm[1]. A failed epidural procedure in which the needle goes beyond the epidural space and perforates the dura mater may cause the patient to experience headaches and, in more severe cases, paralysis or death[2].

Anesthesia needs to be injected in several different doses in the epidural space of the patient. Rather than performing a separate epidural needle insertion procedure for every dose, the anesthesiologist inserts a catheter in the epidural space and anesthesia is then delivered at any time thereafter.

Depending on the form and angle of the spinous processes (Fig. 1) on each vertebra, the doctor has to choose the best needle angle that maximizes the chance of reaching the epidural space without touching bone. In general, the epidural needle is inserted perpendicularly to the skin surface because the space between the spinous processes of vertebrae lumbar 2 and lumbar 3 (L2-L3) and lumbar 3 and lumbar 4 (L3-L4) are usually best accessed at angle 0°, i.e. perpendicular to skin, as seen on Fig. 1. The anesthesiologist then slowly inserts

Further author information: (Send correspondence to D.T.)
D.T.: E-mail: denist@ece.ubc.ca




Figure 1. Epidural needle insertion, midline approach 1)interspinous ligament 2)ligamentum flavum 3)epidural space 4)spinal cord 5)intervertebral disk 6)spinous process

a saline-filled needle while applying continuous force on the plunger to obtain pressure feedback from the tissue around the needle tip. When the needle tip reaches the epidural space, which is a fluid-filled cavity (Fig. 1), the saline fluid is easily injected compared to the stiff surrounding tissue. The anesthesiologist then knows that the needle tip has reached the epidural space. The catheter is then inserted through the epidural needle shaft into the epidural space. Next, the needle is removed leaving the catheter in place. Anesthesia can finally be delivered via the catheter.

It has been observed that the use of ultrasound (US) to view the patient's lumbar vertebrae anatomy facilitates the localization of the epidural space[3]. Moreover, it is suggested that the use of ultrasound images to get a-priori estimates of a suitable puncture point and needle insertion angle greatly improves the learning curve of inexperienced doctors and failures happen less frequently[4]. A study shows two groups of residents performed epidural anesthesia, one with the help of ultrasound (ultrasound group) and the other without (control group)[4]. The ultrasound group achieved a success rate of 84% in the first 10 attempts whereas the control group had a success rate of 60%. Moreover, in the next 50 attempts, the ultrasound group achieved a success rate of 94% whereas the control group had a success rate of 86%. Epidural anesthesia is even more difficult when the patient has abnormal anatomical conditions such as scoliosis[1]. Ultrasound images were also used as real-time feedback for epidural procedures on young infants[5], using the image instead of the traditional loss-of-resistance technique to detect the epidural space. The study shows that the number of bone contacts was greatly reduced and the successful intraoperative analgesia increased from 90.5% to 100% when using ultrasound images as visual feedback. Although the study focuses on epidural depths of two month-old infants, which is around 7mm, and is obviously much smaller than the epidural depth of adults of 20-90mm, it does show the potential benefit of using ultrasound to visualize the lumbar region during an epidural needle insertion procedure.

1.2. Ultrasound imaging of lumbar region

Ultrasound is a very popular imaging tool because it is noninvasive, harmless at low power, portable, accurate and cost effective. However, ultrasound of the lumbar region shows an image filled with speckle and artifacts which can impede detection of important features such as the ligamentum flavum. Ultrasound uses the pulse-echo technique to generate images. The recorded echo is based on reflections from large-scale (relative to wavelength) structures, such as bone (i.e. specular reflection) and reflections from small-scale structures, such as cells (i.e. random scattering). If the specular reflection is strong enough, it casts shadows in the beam direction. Random




Figure 2. a) Ultrasound image of the lumbar spine with speckle noise and shadowing, sagittal paramedian approach is used b) A reference image (no beam-steering) and beam-steered image are shown. Spatial compounding visually uses positive and negative isonation angles to produce a symmetric set of beam-steered images that are subsequently averaged.

scattering creates speckle from constructive and destructive interference[6]. Although the texture of the speckle can be related to tissue type, in general it is considered as a noise present throughout the image. Other artifacts such as reverberation and refraction also affect image quality. As the epidural space is not visible, its location can only be estimated from more visible surrounding structures such as the ligamentum flavum and prior knowledge of the lumbar anatomy.

Ultrasound reflects best when the wavefront hits a structure perpendicular to the direction of propagation. In other words, since ultrasound beams can be steered, the propagation direction can be controlled to some extent. If the angle at which the beam is steered is far from perpendicular to the surface of the structure, the structure may not appear clearly in the image[7]. Structure visibility is also hindered by shadows cast by overlying structures that highly reflect ultrasound, such as bones. Some of these limitations are shown in Fig. 2(a).

1.3. Image processing techniques

Many post-processing methods employ filters such as the diffusion filter[8] and the adaptive weighted median filter[9] to reduce speckle. They all suffer to some degree from loss of fine detail resolution because of the removal of some high frequency content.

Averaging, also known as compounding, has been used previously in three ways: temporal compounding, spatial compounding [10][11][12] and frequency compounding[13]. All of these methods attempt to average a set of ultrasound images of the same region with uncorrelated speckle patterns. The amount of speckle reduction, according to the theory of compounding, on the signal-to-noise ratio (SNR)[14] can be as high as \sqrt{n} , where n is the number of images averaged.

Temporal compounding consists of averaging several images at different times with either a slight movement of the tissues or of the transducer, to try to decorrelate the speckle patterns among images. Unfortunately, the depiction of the tissue also changes so blurring is inevitable.

Frequency compounding consists of averaging several images at different ultrasound transmission frequencies while keeping the probe and tissues stationary. The speckle pattern also decorrelates at different frequencies but is limited by the range of frequencies used. Previous research has shown only a modest improvement in image quality so few commercial ultrasound machines use this approach[13].

Spatial compounding is now becoming more popular among commercial manufacturers and is presented in the next section.

2. SPATIAL COMPOUNDING WITH WARPING, EDGE DETECTION AND LINEAR PREDICTION

2.1. Spatial compounding

Spatial compounding uses beam steering[15][16][17], which captures several frames by sending the ultrasound pulses at different angles of incidence (see Fig. 2(b) for an illustration of the principle). Since speckle noise is dependent on the distribution of reflectors along the ultrasound path [18], changing the beam angle will also change the speckle noise pattern. Then, since the images depict different noise patterns but similar anatomical features, averaging these images will reduce the noise and improve image quality.

Spatial compounding can also produce other benefits, such as the possibility of enhancing structures that are only visible at certain beam angles. For example, it is observed that certain weak but important features, such as a biopsy needle[7][19] only appear at certain beam steered angles. This is because weak features may be situated under a stronger feature, such as a bone, or the surface of the weak feature is at an oblique angle to the ultrasound. The structure of main interest in this work is the epidural space, which is immediately under the ligamentum flavum and initial experience suggests it is only clear at certain beam angles.

Spatial compounding has been investigated as a way to improve other diagnostic tasks. Spatial compounding was applied to ultrasound images of atherosclerotic plaques and breast cancer[11][12][20]. The application of spatial compounding to these images reduces speckle noise and improves the boundary continuity.

2.2. Registration techniques

Despite the ability to rapidly steer the ultrasound beam and rapidly acquire several images in succession of the same region, spatial compounding still suffers from blurring due to misalignment of frames. The speed of sound varies by as much as 14% in soft tissue[21] and this, in turn, causes the apparent positions of structures to be slightly different under different angles of incidence. Re-alignment of the features using an additional non-rigid registration (warping) was previously proposed to properly align the structures of each image[10]. The result was a sharper ultrasound image. This method was tested on artificial phantoms and on a human forearm. Building on those results, the warping/compounding method is extended here to improve visibility of the ligamentum flavum, and therefore the epidural space.

2.2.1. Similarity measures

The success of registration is highly dependent on the similarity measure used as a cost function for finding the best alignment. A successful similarity measure will yield a single strong peak upon best alignment.

Previous literature use several methods such as sum of absolute differences (SAD), mean squared error (MSE)[22], normalized covariance (NCOV), normalized crosscorrelation (NCC), entropy of the difference image and mutual information[23]. Mutual information[24][25] is a very popular similarity measure for registration of multimodality images but is too easily affected by artifacts. NCC and NCOV are very similar measures, the difference being the mean removed term in the covariance, making it slightly more costly. Since the two images to be registered are quite similar, the means are assumed to be small enough so that the NCC and NCOV will yield similar results. The NCC is used for this work.

2.2.2. Interpolation and mapping

The beam-steered images are divided into blocks and each block registered to the reference image. Once the individual warping vectors have been found for each block, each pixel is assigned a warping vector. By smooth interpolation such that a smooth warping transition occurs from one block to its neighbours.

Many interpolation techniques are well known and have been compared on operations such as resizing and rotation[26][27][28]. Popular interpolation techniques are placed in order of performance as follows: nearest neighbour, linear, cubic and cubic B-spline. When testing the different methods, we found that performance was very similar so cubic interpolation was chosen.

Forward or reverse mapping is another choice when it comes to interpolation and compounding[28]. Forward mapping associates a vector to each pixel in the original image and the pixels are copied to the designated position in the destination image. Forward mapping encounters two problems, holes and overlapping pixels in

the destination image. Holes occur when no pixels in the original image were mapped to a pixel in the destination image. The problem can be fixed by interpolating the missing values after mapping (hole filling). The second problem is overlapping, where two pixels in the original image are mapped to the same pixel in the destination image. Overlapping values are usually averaged.

Inverse mapping associates a vector to each pixel in the destination image, corresponding to a pixel in the original image. Inverse mapping does not encounter the problem of holes and overlapping as each pixel has only one associated value. Inverse mapping can encounter problems with complicated warpings, but the speed of sound errors are generally small and smoothly varying. We will use inverse mapping.

2.3. Edge detection techniques

Since the features we are looking for in this work are the ligamentum flavum and the bone boundaries, edge information can be used for several purposes.

First, for reducing computational cost, only blocks with significant edge information are to be warped. In other words, a block needs to depict a significant amount of anatomical structure. This not only reduces computational cost but also prevents useless warping of speckled featureless blocks that might produce an incorrect warping vector, and thus affect the final interpolated warping vector field. Second, the edges can be emphasized in the compounding process to make them more obvious to users which are looking for them.

There are many edge detection methods grouped in two categories[29]: gradient-based and Laplacian-based. Gradient-based methods include filters such as the Prewitt, Sobel, Roberts and the Laplacian-based methods include filters such as the Marr-Hildreth and zero-crossings of the second-order derivative. The Canny edge detection technique[30] is a gradient-based method and is currently the most commonly used method as no methods have shown to consistently yield better performance[29][31][32][33]. The Canny edge detector optimizes 3 criteria:

- the edge detector only has an output when there is an edge and misses no edges
- the distance between the detected edge and the actual edge is as small as possible
- the edge detector only outputs one edge point for each edge, i.e. no multiple detections for a single edge.

The Canny edge detector is used for the purpose of detecting and counting edges and thus assess if a block is suitable for registration.

2.4. Linear prediction techniques

In order to further reduce computational cost, several approaches can be taken. A popular method is a coarse to fine, or multi-resolution approach where lower resolution blocks are registered and then, high resolution blocks are registered using a smaller search region[10]. This method, however requires considerable memory overhead and does not make use of the trends in the warping vectors of neighbouring blocks.

Linear prediction[34] is a standard technique in audio coding in which one would give an initial guess of the current value based on a function of the previous values. The error signal, typically much smaller than the signal itself, requires fewer bits to be coded. The same idea can be used in making an initial guess at the value of the warping vector of the current block based on surrounding blocks, under the assumption that the warping vector field is smoothly varying. The search region for the current block can then be reduced.

When using linear prediction, one needs to define the first block for which one will find the warping vector, and the other blocks will follow. There are three sensible block positions where one can start. The first position is at the top of the image, where the refraction errors are smallest, therefore expecting the smallest warping required. However, the top of the image is mostly noise and is not suitable as the basis for finding warping vectors. A second candidate is at the middle of the image, where a user is most likely to put the anatomical region of interest, and where the resolution is highest. However, it is not guaranteed to be the location where the user places the epidural space. The last location is where there are most edges (features). In this case, the Canny edge detector is used to detect areas containing edges, and then the block with the highest count is assumed to be the best starting candidate. This is a suitable choice for the first warping vector as it should produce the most accurate registration result.




Figure 3. Comparison of human and phantom subjects a) ultrasound image of the L3-L4 of human; b) ultrasound image of L3-L4 of the phantom

The algorithm used for the course of this work builds on the work found in[10]. Spatial compounding is normally performed with the average of images taken at different angles because, in the ideal case, it is said that equal-weight averaging gives best performance[35]. Images are acquired from the Ultrasonix RP-500 ultrasound machine on the phantom, and different beam-steered images are acquired and stored for offline processing. Images are taken using the 4-9MHz linear probe centred at 6.6MHz, a transducer width of 38mm, and 192 scan lines of length 304 pixels.

2.5. Phantom

In order to validate the use of spatial compounding, test images are first taken of a known structure. For this purpose, a phantom is built to closely emulate the structure of the region of interest in the spine, in particular L2-L3 and L3-L4.

A plastic container is used as the base for the phantom to hold the simulated tissue and avoid reflections from the walls and bottom. To emulate the various tissues in the lumbar region, different materials are used. For the vertebra, plastic models are used since the hardness of plastic-like bone is sufficient to fully reflect ultrasound. To emulate fat and muscle, a mixture of agar gelatin, cellulose, glycerol and water is used[36].

First, the muscle layer is made and cooled in a tub of cold water to expedite the cooling process, then the thin fat layer is poured on top of the solidified muscle layer. The purpose is to have a significant difference between the two layers so that realistic refraction errors will occur. Finally, a water-saturated coagulated mixture of rice flour is used to emulate the ligamentum flavum. This material permits ultrasound to propagate through and yet, at some angles, a clear doublet can be seen. The doublet is typically seen in human subjects. The epidural space is right underneath the bright ligamentum flavum as seen on Fig. 3.

Fig. 3 is used to compare the phantom at L3-L4 to the image of a human subject at L3-L4. As can be seen, the features are comparable, so the phantom is used for initial testing.

2.6. Finding the right parameters

It is crucial to find the right parameters for the algorithm to yield best performance at a reasonable computational cost. The first choice is how many beam-steered images and how many degrees between each image. Too small an angle between each image and the speckle noise pattern is still too correlated, too large and there are less images in the range of angles giving a good quality image and therefore, less noise averaging. We will choose images from -8° to 8° with a step size of 2° as it is a good tradeoff between individual image quality and number of images to average[10].




Figure 4. a) with a search region of $\pm 8 \times \pm 4$ b) with a search region of $\pm 16 \times \pm 8$. Little difference can be seen between the images.

2.6.1. Warping parameters

The second parameter to choose is the size of the blocks. The blocks have to be large enough so that a block contains significant anatomical features, therefore making the registration more accurate, and small enough to allow different warping vector for each block as each block is associated with a different refraction error.

The third parameter to choose is the size of the search region for aligning the blocks. A large region increases the chance of having misregistration, where a structure in the beam-steered image is matched with the wrong structure in the reference image. Moreover, a large region causes increases the computational cost. However, a small search region size means that the correct registration may not be found.

In order to find the correct parameters, a checkered image formed by interleaving blocks of the reference image and the warped beam-steered image. A qualitative visual inspection is performed in order to see which parameters give most continuous lines and shades. Fig. 4 compares the effect of changing the search region from $\pm 8 \times \pm 4$ to $\pm 16 \times \pm 8$ on alignment of features. The checkered images show that there is no improvement from increasing the search region. Therefore, the search region is set at $\pm 8 \times \pm 4$, which will reduce computational cost.

2.6.2. Edge detection parameters

The Canny edge detector also contains a set of parameters which need to be tuned for each set of images. The parameters to be tuned are the level of Gaussian smoothing prior to gradient filtering, the low threshold and the high threshold[30]. Here, the low threshold is defined as the value below which all edge pixels are set to zero. The high threshold is the threshold above which all edge pixels are set to mark an edge. An edge pixel is set to mark an edge if its associated value is between the two thresholds and there is a pixel marking an edge in the neighbourhood.

For our purposes of counting edge pixels, the most important feature of the Canny edge detector is to output only one edge pixel per edge. Therefore, conservative thresholds are set so relatively few edge pixels are produced but edges associated with noise are not detected. Fig. 5 compares three different thresholds. From these edges maps, one can see that below 150, noise is still highly present in the edge map, therefore, the high threshold is set to 150, and the low threshold to 90% of the high threshold value as we try to filter out noise as much as possible.

3. RESULTS AND DISCUSSION

3.1. Compounded images

We now compare the images after compounding, which is the image that the operator sees. As can be seen on Fig. 6, simple compounding (Fig. 6(a)) successfully averages out speckle noise, however, the ligamentum flavum and the bone boundaries are blurred. An important point to observe is how the doublet has been blurred to




Figure 5. edge map with a high threshold of a) 50, b) 100 and c) 150. The low threshold is set at 90% of the high threshold.

the point that it appears as a single line. Using warping (Fig. 6(b)) sharpens the compounded image and the ligamentum flavum is seen as a doublet again. Feature detection (Fig. 6(c)) further brightens the ligamentum flavum and the bone boundaries. And finally, using linear prediction (Fig. 6(d)) yields results that are faster but comparable to warping alone; the ligamentum flavum is still seen as a clear doublet.

3.2. Quantitative measures

There are two important features in the image, the ligamentum flavum and the bone boundaries. As the bone boundaries are a transition from a bright region to a dark region (shadow), visibility is related to its gradient.

Therefore, the gradient of the one-dimensional profiles (Fig. 7(a)) are taken and used for quantitative comparison. The original image has the sharpest edge at a gradient of 112, however, the profile is very edgy. Spatial compounding has a smooth gradient, however, the value of the gradient at the edge is only 42. Spatial compounding with warping, with edge detection, and with linear prediction give gradients of 73, 87 and 81 respectively.

For the ligamentum flavum, one needs to see two peaks well above the background. To detect the strength of a peak, we use the Laplacian of a Gaussian (LOG), which is taken by performing a Gaussian smoothing on the one-dimensional profile, then using a Laplacian filter (Fig. 7(b)). The original image again gives the highest LOG of 34, but with substantial noise. The compounded image has less noise, but the LOG is only of 6. The spatial compounding with warping, with edge detection, and with linear prediction have less noise, but keep a high LOG of 25, 25 and 24 respectively.

3.3. Computational cost

The different methods presented above have different performances and computational cost. The different costs are summarized in Table 1. Spatial compounding alone adds very little extra cost. Standard warping is very costly (443ms), and brings the maximum frame rate down to two frames per second which makes it impractical for real-time implementation. Adding edge detection only adds a small amount of time and can cut the time required to obtain the warping vectors (163ms) as we are now only finding the warping vectors for blocks of interest. Then, when using linear prediction, reducing the size of the search region by 2 in each direction (LP2), the cost of finding the warping vectors becomes much smaller (122ms). In theory, with a factor of 2, which reduces the search region by 4, the time required to find the warping vectors should also reduce by a factor of 4. When combining Canny edge detection and LP2, the time is only 35.4ms, which makes it possible for the algorithm run at more than 20 frames per second.

4. CONCLUSION

The choice of parameters will not only affect the intelligibility of the compounded image but also affects the computational cost. For each set of images, a set of parameters was found to yield the best compounded images for a reasonable cost. Compounding with warping using an exhaustive search to align blocks of angled images to the reference frames is still not feasible. By using linear prediction, one can reduce the search region while




Figure 6. images after compounding a) simple compounding, b) compounding with warping, c) compounding with warping and feature detection, d) compounding with warping and linear prediction reducing the search space by a factor two.

Table 1. Computational cost of spatial compounding with warping for different linear prediction gain factors. CPU time is calculated using a P4 3.0GHz with 1GB RAM, an image of 192×304 and a warping search region of $\pm 8 \times \pm 4$ pixels. LP2 indicates registration with linear prediction and a reduction of the search space by factor two.

method	CPU time (ms)
image resizing and angle correction	1.5
averaging images	2.6
displaying image	2.4
interpolating warping vectors	1.5
remapping the image	1.3
Canny edge detection	1.5
finding the warping vectors	443
finding the warping vectors with LP2	122
finding the warping vectors with Canny	163
finding the warping vectors with Canny and LP2	35.4





Figure 7. one-dimensional slice of a) bone boundaries with gradient and b) ligamentum flavum with LOG, original (solid), compound (dashed), compound with warping (dotted), with edge detection (dot-dashed), and with linear prediction (dotted-dotted-dashed)

maintaining good performance on reducing speckle noise and restore continuous boundaries on features such as bones and the important ligamentum flavum. Real-time performance can be achieved on a standard processor.

Work has already begun on images taken *in vivo*. The algorithm's parameters will need to be optimized for the different cases and a global set of best parameters for all cases will be found. The parameters will depend on the actual size and appearance of the features.

As 3D imaging becomes more useful and popular, it is a good idea to extend the previous concepts to 3D. 3D permits visualization of an entire volume and provides more insightful images to a user. Computational costs will become a major issue when it comes to imaging. Complexity tends to be of order n^2 in 2D and of the order n^3 for 3D. Efforts will be made to decrease computational complexity, following the techniques used to reduce computational complexity in the 2D case. For instance, linear prediction, with a gain factor of 4 in the 2D case theoretically decreases the cost by a factor 16, and in the 3D case the theoretical decrease is 64. Once again, when the tests on the phantom will show significant improvement, the algorithms will be tested *in vivo*.

ACKNOWLEDGMENTS

This work is supported by a collaborative Health Research Project grant funded by the Natural Sciences and Engineering Research Council and the Canadian Institute of Health Research.

REFERENCES

1. T. Grau, R. W. Leipold, R. Conradi, and E. Martin, "Ultrasound control for presumed difficult epidural puncture," *Acta Anaesthesiol Scand* **45**, pp. 766–771, 2001.
2. T. Grau, R. W. Leipold, S. Fatehi, E. Martin, and J. Motsch, "Real-time ultrasonic observation of combined spinal-epidural anaesthesia," *Eur J Anaesthesiol* **21**, pp. 25–31, 2004.
3. T. Grau, R. W. Leipold, R. Conradi, E. Martin, and J. Motsch, "Ultrasound imaging facilitates localization of the epidural space during combined spinal and epidural anesthesia," *Reg Anesth Pain Med* **26**, pp. 64–67, 2004.
4. T. Grau, E. Bartussek, R. Conradi, E. Martin, and J. Motsch, "Ultrasound imaging improves learning curves in obstetric epidural anesthesia: a preliminary study," *Can J Anaesth* **50**, pp. 1047–1050, 2003.
5. H. Willschke, P. Marhofer, A. Bösenberg, S. Johnston, O. Wanzel, C. Sitzwohl, S. Kettner, and S. Kapral, "Epidural catheter placement in children: comparing a novel approach using ultrasound guidance and a standard loss-of-resistance technique," *British Journal of Anaesthesia* **97**, pp. 200–207, 2006.
6. O. V. Michailovich and A. Tannenbaum, "Despeckling of medical ultrasound images," *IEEE Trans. on Ultrasonics, Ferroelectrics, and Frequency Control* **53**, pp. 64–78, 2006.
7. S. Cheung and R. Rohling, "Enhancement of needle visibility in ultrasound guided percutaneous procedures," *Ultrasound Med. Biol.* **30**, pp. 617–624, 2004.
8. P. Perona and J. Malik, "Scale space and edge detection using anisotropic diffusion," *IEEE Trans. Pattern and Anal. Machine Intell.* **12**, pp. 629–639, 1990.
9. T. Loupas, W. N. McDicken, and P. L. Allan, "An adaptive weighted median filter for speckle suppression in medical ultrasound images," *IEEE Trans. Circuits Syst.* **36**, pp. 129–135, 1989.
10. A. Groves and R. Rohling, "Two dimensional spatial compounding with warping," *Ultrasound Med. Biol.* **30**, pp. 929–942, 2004.
11. S. Huber, M. Wagner, M. Medl, and H. Czerbirek, "Real-time spatial compound imaging in breast ultrasound," *Ultrasound Med. Biol.* **28**, pp. 155–163, 2002.
12. S. K. Jespersen, J. E. Wilhjelm, and H. Sillesen, "In vitro spatial compound scanning for improved visualization of atherosclerosis," *Ultrasound Med. Biol.* **26**, pp. 1357–1362, 2000.
13. G. E. Trahey, S. W. Allison, J.W. and Smith, and O. T. Von Ramm, "A quantitative approach to speckle reduction via frequency compounding," *Ultrasonic Imaging* **8**, pp. 151–164, 1986.
14. R. C. Gonzales and R. E. Woods, *Digital image processing*, Prentice Hall, 2 ed., 2000.
15. M. Berson, A. Roncin, and L. Pourcelot, "Compound scanning with an electrically steered beam," *Ultrasonic Imaging* **3**, pp. 303–308, 1981.

16. D. A. Carpenter, M. J. Dadd, and G. Kossoff, "A multimode real time scanner," *Ultrasound Med. and Biol.* **6**, pp. 279–284, 1980.
17. S. Jespersen, J. E. Wilhjelm, and H. Sillesen, "Multi-angle compound imaging," *Ultrasonic Imaging* **20**, pp. 81–102, 1998.
18. G. E. Trahey, S. W. Smith, and O. T. von Ramm, "Speckle patterns correlation with lateral aperture translation: Experimental results and implications for spatial compounding," *IEEE Trans. Ultrason. Ferro. Freq. Contr.* **33**, pp. 257–264, 1986.
19. A. Saleh, S. Ernst, A. Grust, G. Fürst, P. Dall, and U. Mödder, "Real-time compound imaging: Improved visibility of biopsy needles and localization wires as compared to single-line ultrasonography," *Fortschr Röntgenstr* **173**, pp. 368–372, 2001.
20. M. E. Anderson, M. S. Soo, R. C. Bentley, and G. E. Trahey, "The detection of breast microcalcifications with medical ultrasound," *Journal of the Acoustical Society of America* **101**, pp. 29–39, 1997.
21. D. Christensen, *Ultrasonic bioinstrumentation*, John Wiley and Sons, 1988.
22. B. Cohen and I. Dinstein, "New maximum likelihood motion estimation schemes for noisy ultrasound images," *Pattern Recognition* **35**, pp. 455–463, 2001.
23. G. P. Penney, J. Weese, J. A. Little, P. Desmedt, D. L. G. Hill, and D. J. Hawkes, "A comparison of similarity measures for use in 2-d–3-d medical image registration," *IEEE Trans. on Medical Imaging* **17**, pp. 586–595, 1998.
24. J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever, "Mutual-information-based registration of medical images: a survey," *IEEE Transactions on Medical Imaging* **22**, pp. 986–1004, 2003.
25. P. A. Viola and W. M. Wells, "Alignment by maximization of mutual information," *International Journal of Computer Vision* **24**(2), pp. 137–154, 1997.
26. P. Thévenaz, T. Blu, and M. Unser, "Interpolation revisited," *IEEE Trans. on Med. Imaging* **19**, pp. 739–758, 2000.
27. T. M. Lehmann, C. Gonner, and K. Spitzer, "Survey: interpolation methods in medical image processing," *IEEE Trans. on Med. Imaging* **18**, pp. 1049–1075, 1999.
28. G. Wolberg, *Digital image warping*, IEEE Computer Society Press, 1994.
29. M. Heath, S. Sarkar, T. Sanocki, and K. Bowyer, "Comparison of edge detectors: a methodology and initial study," *Computer and Image Understanding* **69**, pp. 38–54, 1998.
30. J. Canny, "A computational approach to edge detection," *IEEE Trans. Pattern Analysis and Machine Intelligence* **8**, pp. 679–714, 1986.
31. J. Shen and S. Castan, "An optimal linear operator for step edge detection, computer vision," *CVGIP: Graphical Models and Image Processing* **54**, pp. 112–133, 1992.
32. S. Alkaabi and F. Deravi, "Variations of ISEF edge detectors," *Electronics Letters* **39**, pp. 1174–1175, 2003.
33. M. Sharifi, M. Fathy, and M. T. Mahmoudi, "A classified and comparative study of edge detection algorithms," *Proc. of the Int. Conf. on Inf. Tech.: Coding and Computing*, pp. 117–120, 2002.
34. A. Gersho and R. A. Gray, *Vector quantization and signal compression*, Springer-Verlag, 1992.
35. J. E. Wilhjelm, M. S. Jensen, S. K. Jespersen, B. Sahl, and E. Falk, "Visual and quantitative evaluation of selected image combination schemes in ultrasound spatial compound scanning," *IEEE Trans. on Medical Imaging* **23**, pp. 181–190, 2004.
36. M. Burlew, E. Madsen, J. Zagzabski, R. Banjavic, and S. Sum, "A new ultrasound tissue-equivalent material," *Radiology* **134**, pp. 517–520, 1980.