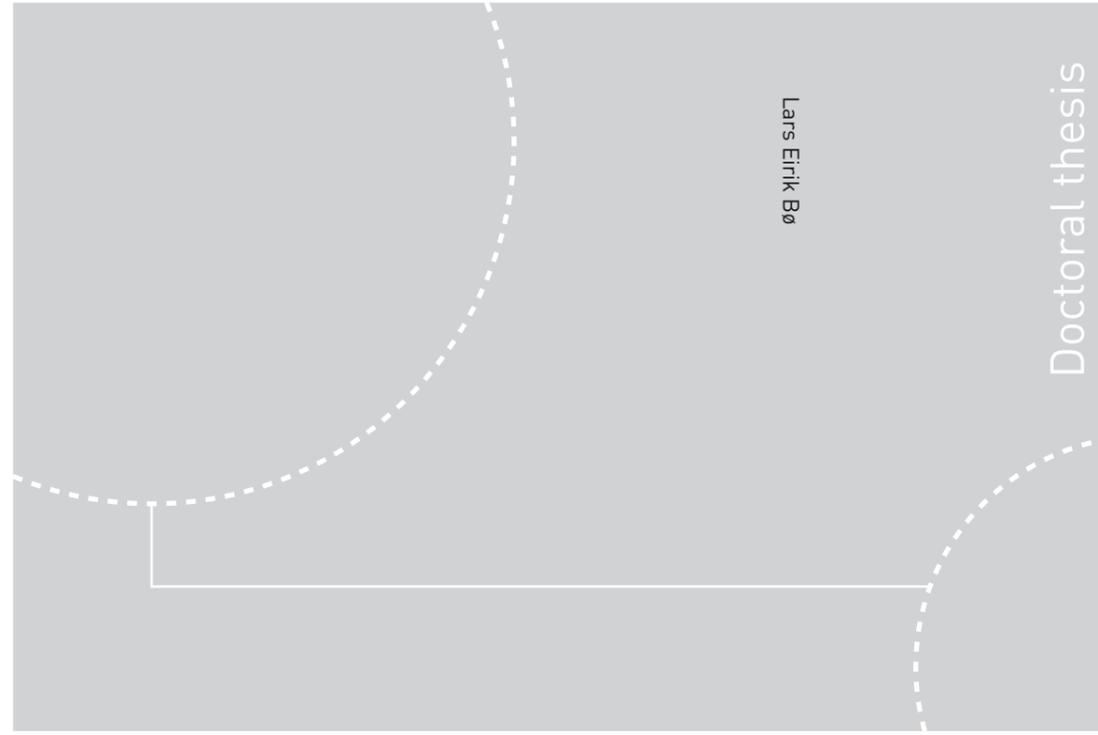


ISBN 978-82-326-3304-3 (printed ver.)  
ISBN 978-82-326-3305-0 (electronic ver.)  
ISSN 1503-8181



Doctoral theses at NTNU, 2018:257

Lars Eirik Bø

# Ultrasound in image-guided spine surgery

Enabling technologies and first steps

 **NTNU**  
Norwegian University of  
Science and Technology

Doctoral theses at NTNU, 2018:257

 NTNU

**NTNU**  
Norwegian University of Science and Technology  
Thesis for the Degree of  
Philosophiae Doctor  
Faculty of Medicine and Health Sciences  
Department of Circulation and Medical Imaging

 **NTNU**  
Norwegian University of  
Science and Technology

Lars Eirik Bø

# Ultrasound in image-guided spine surgery

Enabling technologies and first steps

Thesis for the Degree of Philosophiae Doctor

Trondheim, September 2018

Norwegian University of Science and Technology  
Faculty of Medicine and Health Sciences  
Department of Circulation and Medical Imaging



Norwegian University of  
Science and Technology

**NTNU**  
Norwegian University of Science and Technology

Thesis for the Degree of Philosophiae Doctor

Faculty of Medicine and Health Sciences  
Department of Circulation and Medical Imaging

© Lars Eirik Bø

ISBN 978-82-326-3304-3 (printed ver.)  
ISBN 978-82-326-3305-0 (electronic ver.)  
ISSN 1503-8181

Doctoral theses at NTNU, 2018:257

Printed by NTNU Grafisk senter

# Ultralyd i bildeveiledet ryggkirurgi

– *muliggjørende teknologi og første steg*

Kirurgiske navigasjonssystemer er i dag i utstrakt bruk innen mange ulike kliniske områder. På samme måte som GPS-systemer i biler hjelper førerne med å skjønne sammenhengen mellom kartet, bilens posisjon og det de ser utenfor bilen, hjelper disse systemene kirurgene ved å visualisere hvordan de medisinske bildene, som er deres kart, forholder seg til de kirurgiske instrumentene og pasienten. For å få til dette bruker man ulike typer posisjoneringsteknologi som i sanntid og med stor nøyaktighet kan følge med på posisjonen til de ulike instrumentene mens operasjonen pågår.

Det er utviklet flere navigasjonssystemer for ryggkirurgi basert på røntgenavbildning, og disse har vist seg å være nyttige, særlig ved plassering av skruer i ryggraden. Røntgen egner seg imidlertid dårlig for avbildning av bløtvev, og disse systemene kan derfor ikke brukes for en del andre vanlige inngrep, som for eksempel prolapskirurgi. I dette arbeidet har målet derfor vært å utvikle metoder og verktøy for å kunne veilede slike inngrep basert på ultralyd- og magnetresonansavbildning (MR).

Mange rygginngrep utføres i dag ved hjelp av mikroskop og bare små snitt i huden, såkalt mikrokirurgi. Dette gjør det vanskelig å komme til for avbildning med vanlige ultralydprober. Gruppen vår har

derfor tidligere utviklet en ny probe spesielt utformet for slike inngrep. I denne avhandlingen har vi utviklet og testet metoder både for å kunne bruke denne proben med eksisterende posisjoneringsteknologi og for å sette sammen bildene fra proben til tredimensjonale bildevolum som egner seg for navigasjon. I tillegg har vi studert metoder for å kunne tilpasse MR-bilder tatt før inngrepet til slike tredimensjonale ultralydbilder slik at disse også kan brukes til navigasjon. Disse siste metodene er imidlertid foreløpig ikke robuste nok til å kunne brukes på pasienter.

**Kandidat:** Lars Eirik Bø

**Institutt:** Institutt for sirkulasjon og bildediagnostikk

**Veiledere:** Toril A. Nagelhus Hernes, Ingerid Reinertsen, Frank Lindseth og Ole Solheim

**Finansieringskilde:** Samarbeidsorganet mellom Helse Midt-Norge RHF og NTNU

*Ovennevnte avhandling er funnet verdig til å forsvares offentlig for graden ph.d. i medisinsk teknologi. Disputas finner sted i auditorium MTA fredag 14. september 2018 kl. 12:15.*

# Abstract

Most people experience back pain at some point in their life. While most of these conditions are treated nonoperatively, surgical treatment has been shown to be both effective and cost effective compared to nonoperative care for both intervertebral disc herniation and spinal stenosis in selected patients. Surgical navigation systems enabling image guidance based on preoperative or intraoperative computed tomography (CT) images have found some use in spine surgery. Here, they are most frequently used for spinal fusion, and the benefits of image guidance in such procedures, under given conditions, have been documented in several studies. In spite of this, few spine surgeons use navigation routinely.

High cost is one of the most important barriers to a more widespread adoption of navigation systems in spine surgery. Extending the use of navigation to more than just fusion procedures, and to soft tissue procedures such as disk herniations in particular, could help the surgeon substantiate the cost of the equipment. A big step in this direction would be to enable navigation based on other imaging modalities than CT, such as magnetic resonance (MR) or ultrasound imaging. In this work, the goal has therefore been to enable the use of ultrasound imaging both for intraoperative imaging and for registration of preoperative MR images to the patient.

Many spine procedures are today performed with a microsurgical approach, and the small incisions used in these procedures prohibit the use of standard probes. Our group has therefore, in a previous project, developed a new probe specifically designed to enable ultra-

sound imaging through such small incisions. The main part of the work has been directed towards enabling tracking and navigation with this probe. In addition, we have studied methods for registering MR images to ultrasound images of the spine.

In Paper A, we looked at methods for reconstructing three-dimensional image volumes from tracked two-dimensional ultrasound images. Both different means of capturing the original ultrasound images and different reconstruction algorithms were thoroughly compared. We found that the differences were small, and while the various methods showed different strengths and weaknesses, the overall result was that they could not be separated.

In Paper B, we explored the feasibility of using electromagnetic (EM) tracking in an operating room setting, both alone and in combination with a robotic C-arm. We also compared the performance of the standard EM field generator with a new prototype designed specifically for use with fluoroscopic imaging equipment. We found that while the accuracy decreased considerably with the C-arm inside the operating field, the measurements were still stable. We thus concluded that by implementing a suitable static correction scheme, the tracking system and the C-arm could potentially be used together.

In Paper C, we presented a new method for ultrasound probe calibration, which is the process of finding the spatial relationship between the coordinate system of the tracking sensor that is integrated in the ultrasound probe and the coordinate system of the ultrasound images generated by the probe. In a research setting, such as ours, new probes are tested regularly, and the method was therefore designed to be used with a large variety of probes without any adaptation. The method was tested on three very different probes demonstrating both great versatility and high accuracy.

In Paper D, we developed a method for registration of preoperative MR images to the patient by means of intraoperative ultrasound imaging using a tracked ultrasound probe. The method segmented the posterior bone surface from both the ultrasound images and the MR images and registered the two surfaces to each other using a modified version of the Iterative Closest Point algorithm. For this paper, the

method was only tested on one subject, but the accuracy of the registration on this subject was clinically relevant, and we concluded that the method was promising.

In conclusion, we have developed and tested technology that enables tracking of small, intraoperative ultrasound probes and allows the generation of three-dimensional volumes suitable for navigation from such images. We have also investigated the use of intraoperative ultrasound imaging for registration of preoperative CT and MR images to the spine. The latter is, however, a work in progress, as the methods that we have tested have so far have not been sufficiently robust for clinical use.



# Preface

This thesis has been submitted in partial fulfillment of the requirements for the degree philosophiae doctor (PhD) at the Faculty of Medicine and Health Sciences of the Norwegian University of Science and Technology (NTNU). The research was funded by the Liaison Committee between the Central Norway Regional Health Authority and NTNU, and was carried out under the supervision of professor Toril A. Nagelhus Hernes at NTNU. Co-supervisors were senior scientist Ingerid Reinertsen at SINTEF, professor Frank Lindseth at NTNU and professor and consultant neurosurgeon Ole Solheim at NTNU and St. Olavs hospital.

## Acknowledgments

First, I want to thank professor Toril Hernes for recruiting me into her research group at SINTEF and introducing me to the field of image-guided surgery. She was also the one who encouraged me to pursue a PhD and who supervised me through the lengthy process that it became. I am also very grateful to my co-supervisors Ingerid Reinertsen, Frank Lindseth and Ole Solheim who have all been very supportive and made valuable contributions to my work at different times. Special thanks to Ingerid, who saw me through the last couple of years when both results and motivation were in short supply.

Further, I want to thank my co-authors on the four papers that form part of this thesis. Here, Ole Vegard Solberg is worth a special

mention for the tremendous time and effort he invested as first author of the comprehensive first paper. I also want to thank everyone else who has helped me with big or small tasks or participated in discussions related to this work.

For the past ten years, seven of which were partly spent on my PhD project, I have been fortunate enough to be part of the medical technology research group at SINTEF. It has been a true pleasure to be surrounded by so many companionable, supportive and skilled people, and I want to take the opportunity to thank everyone who has formed part of the group or visited us for shorter or longer periods of time during these years.

Finally, I want to thank my wife, Mari. She is always there for me and has supported me generously in this work even, and especially, in times when other and more important matters required most of my attention and effort.

# Contents

|          |  |           |
|----------|--|-----------|
| <b>1</b> | <b>Background</b>  | <b>1</b>  |
| 1.1      | Surgical navigation . . . . .                                | 1         |
| 1.2      | Surgical treatment of spine disorders . . . . .              | 9         |
| 1.3      | Imaging of the spine . . . . .                               | 12        |
| 1.4      | Navigation in spine surgery . . . . .                        | 17        |
| 1.5      | Emerging technologies . . . . .                              | 19        |
| <b>2</b> | <b>Aims of study</b>   | <b>23</b> |
| 2.1      | Intraoperative imaging . . . . .                             | 24        |
| 2.2      | Registration of preoperative images to the patient . . . . . | 25        |
| <b>3</b> | <b>Equipment and infrastructure</b>                          | <b>27</b> |
| <b>4</b> | <b>Summary of papers</b>                                     | <b>31</b> |
| 4.1      | Paper A . . . . .  | 31        |
| 4.2      | Paper B . . . . .  | 32        |
| 4.3      | Paper C . . . . .  | 33        |
| 4.4      | Paper D . . . . .  | 34        |
| <b>5</b> | <b>Discussion and future work</b>                            | <b>37</b> |
| 5.1      | Intraoperative ultrasound imaging . . . . .                  | 38        |
| 5.2      | Reconstruction . . . . .                                     | 41        |
| 5.3      | Tracking accuracy . . . . .                                  | 42        |
| 5.4      | Probe calibration . . . . .                                  | 45        |
| 5.5      | Registration . . . . .                                       | 46        |

|  |            |
|--|------------|
| <b>6 Conclusion</b>  | <b>51</b>  |
| <b>References</b>  | <b>53</b>  |
| <b>A 3D ultrasound reconstruction algorithms from analog and digital data</b>                                  | <b>63</b>  |
| A.1 Introduction . . . . .   | 64         |
| A.2 Materials and methods . . . . .  | 66         |
| A.3 Results . . . . .  | 82         |
| A.4 Discussion . . . . .   | 94         |
| A.5 Conclusion . . . . .   | 105        |
| References . . . . .   | 106        |
| <b>B Accuracy of electromagnetic tracking with a prototype field generator in an interventional OR setting</b> | <b>113</b> |
| B.1 Introduction . . . . .   | 114        |
| B.2 Materials and methods . . . . .  | 116        |
| B.3 Results . . . . .  | 122        |
| B.4 Discussion . . . . .   | 127        |
| B.5 Conclusion . . . . .   | 129        |
| References . . . . .   | 130        |
| <b>C Versatile robotic probe calibration for position tracking in ultrasound imaging</b>                       | <b>133</b> |
| C.1 Introduction . . . . .   | 134        |
| C.2 Materials and methods . . . . .  | 139        |
| C.3 Results . . . . .  | 150        |
| C.4 Discussion . . . . .   | 150        |
| C.5 Conclusion . . . . .   | 154        |
| References . . . . .   | 155        |
| <b>D Registration of MR to percutaneous ultrasound of the spine for image-guided surgery</b>                   | <b>159</b> |
| D.1 Introduction . . . . .   | 160        |
| D.2 Methods and experiments . . . . .  | 162        |

|                          |     |
|--------------------------|-----|
| D.3 Results . . . . .    | 167 |
| D.4 Discussion . . . . . | 168 |
| D.5 Conclusion . . . . . | 170 |
| References . . . . .     | 171 |



# Abbreviations

|        |   |
|--------|---|
| 2D     | two-dimensional.  |
| 3D     | three-dimensional.  |
| CMUT   | capacitive micromachined ultrasound transducers.                |
| CT     | computed tomography.  |
| EM     | electromagnetic.  |
| GPS    | Global Positioning System.                                      |
| MR     | magnetic resonance.   |
| OR     | operating room.   |
| RCL    | Robotics and Control Laboratory.                                |
| RMS    | root mean square.   |
| VIRTUS | Vertebral Intraoperative Repair by Targeted Ultrasound imaging. |



# List of Figures

|      |  |    |
|------|--|----|
| 1.1  | Digital maps. . . . .  | 2  |
| 1.2  | Commercial navigation platforms. . . . .                         | 4  |
| 1.3  | The optical tracking system Polaris Spectra. . . . .             | 6  |
| 1.4  | The EM tracking system Aurora. . . . .                           | 7  |
| 1.5  | Navigated brain tumour surgery. . . . .                          | 10 |
| 1.6  | The augmented reality glasses HoloLens. . . . .                  | 11 |
| 1.7  | Projectional X-ray and CT of the spine. . . . .                  | 13 |
| 1.8  | The Artis zeego C-arm. . . . .                                   | 14 |
| 1.9  | MR volume of the lumbar spine. . . . .                           | 15 |
| 1.10 | Ultrasound images of the spine. . . . .                          | 16 |
| 1.11 | Ultrasound image of a spinal tumour. . . . .                     | 18 |
| 1.12 | The Airo intraoperative CT system. . . . .                       | 19 |
| 2.1  | The L13-7 ultrasound probe. . . . .                              | 24 |
| 3.1  | Prototype field generator. . . . .                               | 28 |
| 3.2  | The UR5 robotic arm and the SonixMDP ultrasound scanner. . . . . | 29 |
| 5.1  | Ultrasound image of a lumbar disc herniation. . . . .            | 39 |
| 5.2  | Pituitary tumour surgery. . . . .                                | 40 |
| 5.3  | Prototype ultrasound probe. . . . .                              | 41 |
| 5.4  | Statistical shape model of the lumbar spine. . . . .             | 48 |
| A.1  | System setup. . . . .  | 68 |
| A.2  | The imaging sector of an ultrasound probe. . . . .               | 71 |
| A.3  | Theoretical lateral and elevation resolution. . . . .            | 73 |

|      |   |     |
|------|---|-----|
| A.4  | Illustrations of ultrasound sectors. . . . .                          | 81  |
| A.5  | RMS values for the three different reconstruction algorithms. . . . . | 83  |
| A.6  | Identification of 6 structures placed increasingly closer. . . . .    | 85  |
| A.7  | Thread size as a measurement of spatial resolution. . . . .           | 86  |
| A.8  | Anyplane images of small structures. . . . .                          | 91  |
| A.9  | Number of identified structures. . . . .                              | 91  |
| A.10 | Ultrasound images of the underside of the forearm. . . . .            | 93  |
|      |   |     |
| B.1  | The Aurora electromagnetic measurement system. . . . .                | 118 |
| B.2  | Accuracy phantom. . . . .   | 120 |
| B.3  | Variation in measured distance during imaging. . . . .                | 124 |
| B.4  | Robustness measurements. . . . .                                      | 126 |
|      |   |     |
| C.1  | Coordinate systems in probe calibration. . . . .                      | 137 |
| C.2  | Probe calibration setup. . . . .                                      | 141 |
| C.3  | Imaging target for probe calibration. . . . .                         | 143 |
| C.4  | Ultrasound images of the imaging target. . . . .                      | 144 |
| C.5  | Imaging target segmentation. . . . .                                  | 145 |
| C.6  | Ultrasound probes used for evaluation. . . . .                        | 147 |
| C.7  | Accuracy phantom. . . . .   | 148 |
|      |   |     |
| D.1  | Ultrasound bone segmentation. . . . .                                 | 165 |
| D.2  | Segmented bone surfaces. . . . .                                      | 166 |
| D.3  | Bone surface registration. . . . .                                    | 168 |
| D.4  | Ultrasound images overlaid on MR images of the spine. . . . .         | 168 |

# List of Tables

- A.1 RMS values obtained by removing data from input images. 84
- A.2 Number of identified small structures. . . . . 87
- A.3 Image measurements of spatial resolutions. . . . . 88
- A.4 Comparisons based on the number of identified structures. . 92
- A.5 Comparisons of reconstruction algorithms. . . . . 94
- A.6 Comparisons of input ultrasound sources . . . . . 95
- A.7 Summary of results. . . . . 96
  
- B.1 Accuracy measurements. . . . . 125
  
- C.1 Ultrasound probes used for evaluation. . . . . 146
- C.2 Probe calibration results. . . . . 151
  
- D.1 Registration accuracy. . . . . 169



# Chapter 1

## Background

### 1.1 Surgical navigation

Since the invention of X-ray imaging in 1896 [1], surgeons have used medical images to plan their procedures. One of the big challenges for an operator performing a medical intervention guided by image data, is to relate what is seen in the images to the actual patient. The problem is similar to trying to determine your position in a map by looking at the surroundings and deciding on where to go based on the information in the map. How hard this is depends both on the quality of the map and on the characteristics of the surroundings: some places have distinct landmarks that are easily identified both in the map and in real life, while in other places, it is hard to tell one hill or house from the next.

In digital maps, this is now solved using Global Positioning System (GPS) tracking to show your position in the map. The automotive navigation systems that are common in modern cars can also generate a virtual view of the landscape ahead based on the information in the maps, making it easier for the driver to compare this information to the actual surroundings. An example of this is shown in Figures 1.1(a)–(b).

Surgical navigation systems are very analogous to this. The med-

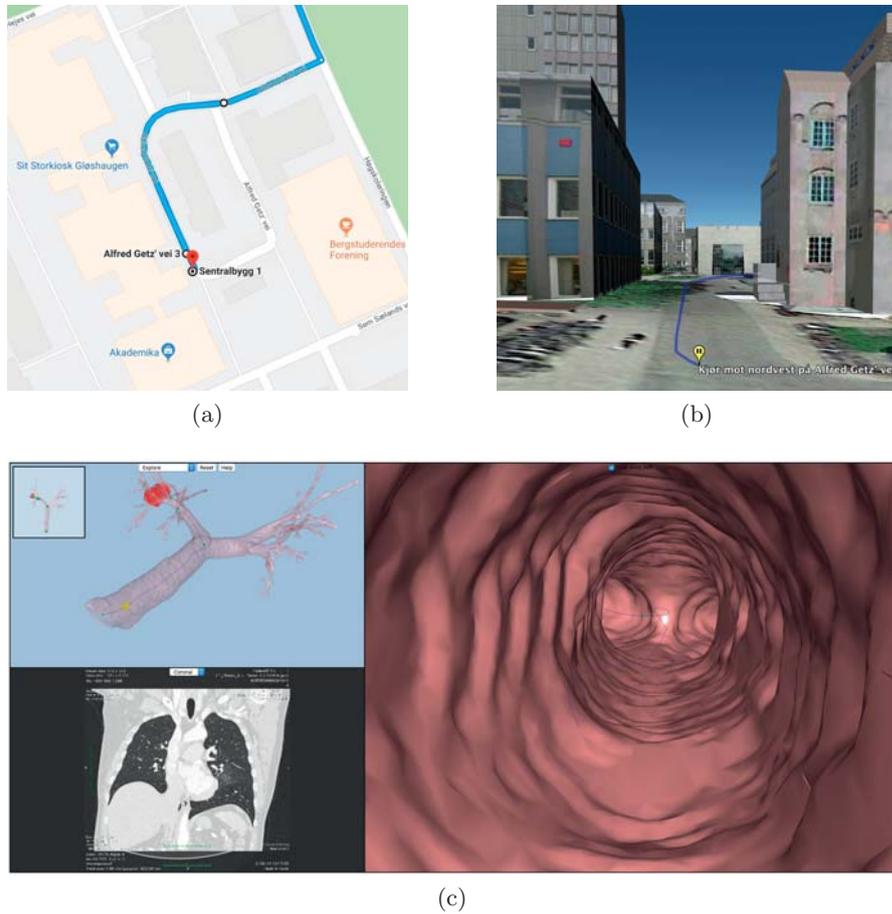


Figure 1.1: (a) Digital maps, such as this one from Google Inc. (<https://www.google.no/maps/@63.4177685,10.4051033,18.2z>), can use GPS tracking to indicate the users' location within the map and show them where to go. (b) They can also generate a virtual view of the surroundings to ease the interpretation of the information in the map. (c) Surgical navigation systems do the same using medical images and tracking systems. This example is from the open-source navigation system Fraxinus (SINTEF, Trondheim, Norway).

ical images are the maps, and instead of GPS they use various kinds of tracking systems to measure the positions orientations of the surgical instruments. These can then be visualised in the images, giving the surgeon information about where the instruments are in relation to anatomy that is not directly visible to the eye. During the planning of the procedure, the surgeon can also add information to the system, such as where to make the incision, which trajectory to follow to the target anatomy and where various organ boundaries are. This information can then be visualised together with the images and the instruments, and the visualisations can be adapted to each step of the procedure in order to make the interpretation as intuitive as possible for the surgeon. An example of this can be seen in Figure 1.1(c). There are today many commercial navigation systems on the market ranging from the very specialised ones to more general systems that can be used in a wide range of clinical settings. The major navigation platforms are today the Medtronic StealthStation (Medtronic, Dublin, Ireland), the Brainlab Curve (Brainlab AG, Munich, Germany) and the Stryker NAV3i (Stryker Corporation, Kalamazoo, MI, U.S.A), which all three are adapted to a number of different medical procedures. They are shown in Figure 1.2.

## **Tracking**

The most common tracking technology for surgical navigation is optical tracking using infrared light. [2] Here, each of the tools that are to be tracked is equipped with three or more light sources, which are arranged in a certain geometrical pattern. The sources can either generate light themselves, or reflect light from an external light source. Infrared light pulses are then emitted at a high frequency. For each pulse, one or more infrared cameras record the distance and angle to each of the sources, thus determining their position in space, and by comparing these positions to the known geometrical patterns of the various tools, the position and orientation of each tool is calculated. In medical applications, the most common system of this kind is the Polaris Spectra from Northern Digital Incorporated (NDI, Waterloo,



Figure 1.2: The three major, commercial navigation platforms: (a) The Medtronic StealthStation (Medtronic, Dublin, Ireland), (b) the Brainlab Curve (Brainlab AG, Munich, Germany) and (c) the Stryker NAV3i (Stryker Corporation, Kalamazoo, MI, U.S.A.).

ON, Canada), shown in Figure 1.3, which is incorporated in both the StealthStation and the Brainlab system. Under good conditions, with well calibrated tools and a no visual obstructions, these tracking systems can be very accurate, typically with an overall root mean square (RMS) distance error below 0.3 mm. [3] The biggest disadvantage with such systems is that they require a clear line of sight between the camera and the tool, which can often be challenging to achieve in a crowded operating room. [2]

An alternative to the optical systems is the electromagnetic (EM) tracking systems. These generate an EM field with a known geometry and measure this field with small EM field sensors. [2] An example of such a system can be seen in Figure 1.4. Based on these measurements, the positions and orientations of the sensors within the field are calculated. The main advantage of the electromagnetic systems is that there is no line-of-sight requirement, which is particularly useful in the case of flexible instruments such as flexible endoscopes, catheters or guide wires. The sensors can also be made very small, down to a few tenths of millimetre in diameter, and they are therefore very easily integrated in the various surgical tools. On the downside, the sensors must be connected to the system, usually with cables, which can be a problem in certain settings. Still, the most important challenge with the electromagnetic systems is that they are very susceptible to disturbances from ferromagnetic interference sources in the surroundings. [4] It is therefore very important to avoid large, ferromagnetic objects within the magnetic field, which requires a very careful setup. Assuming that this can be achieved, the accuracy of the electromagnetic systems is only slightly worse than that of the optical ones, often with an overall RMS distance error just below 1 mm. [4, 5]

## **Registration**

In order to visualise the tracked tools and the medical images together, they must have a common frame of reference. To achieve this, the spatial relationships between the coordinate systems of the tracking system and of the various images must be found. The process of finding

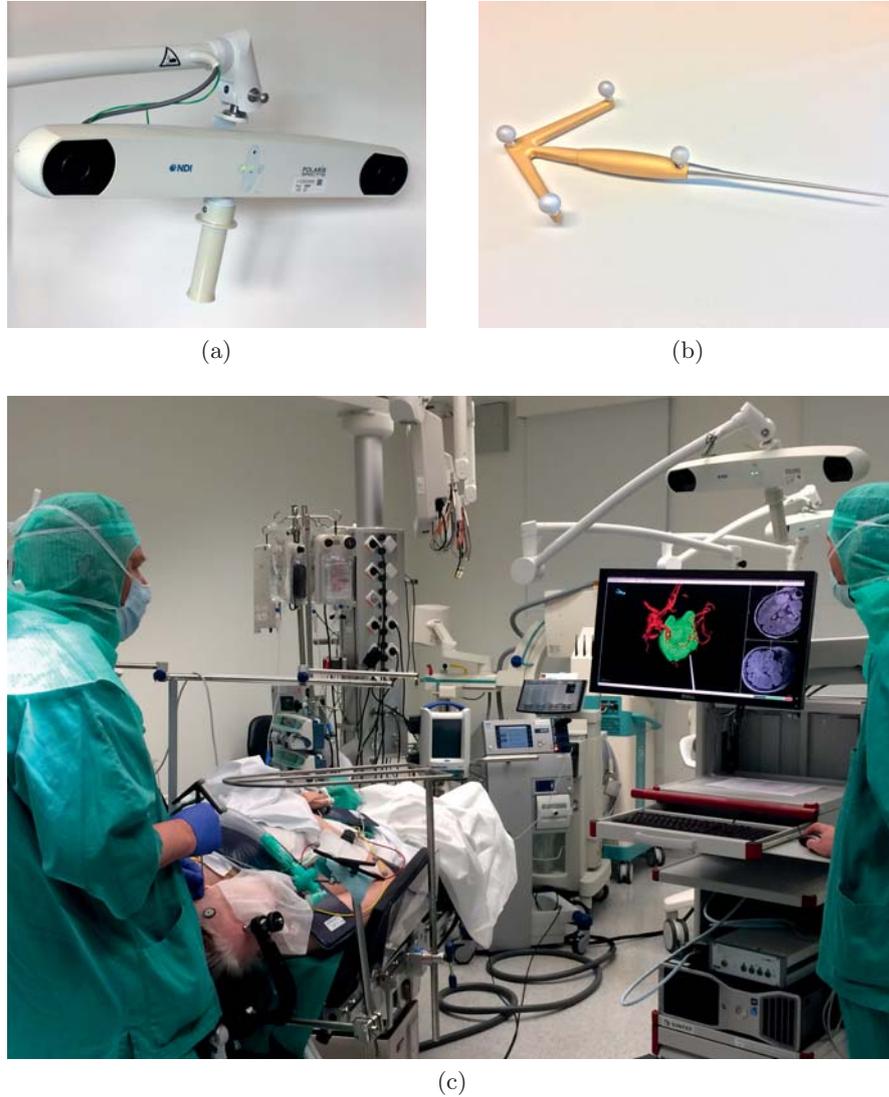


Figure 1.3: The optical tracking system Polaris Spectra from NDI (Waterloo, ON, Canada). The tracking camera in (a) incorporates two infrared cameras and a flash consisting of multiple infrared light-emitting diodes and the navigation pointer in (b) has four reflective spheres that reflect the light emitted by the flash. (c) shows a neurosurgeon using the pointer for navigation during an operation.

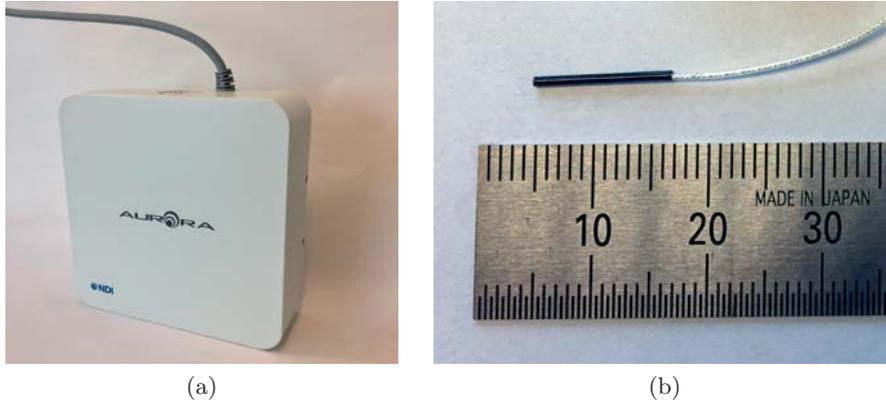


Figure 1.4: The EM tracking system Aurora (NDI, Waterloo, ON, Canada), with the field generator to the left and a typical tracking sensor with six degrees of freedom to the right.

these relationships is called registration, and this is very important for the overall accuracy of the system. There are several ways of doing this, and the various methods can generally be classified as landmark based (also called feature or fiducial based), surface based or voxel based (also called intensity based). [6]

In landmark-based methods, a number of visible landmarks are identified in the different coordinate systems, creating pairs of corresponding coordinates, and the spatial relationship is computed based on these pairs. Such methods are typically used to register medical images to the patient. The landmarks can either be natural anatomical features or artificial markers, so-called fiducials, that are attached to the patient before imaging. These are often identified manually by the surgeon, first by marking them in the images and then, when the patient is fixed to the operating table, by pointing at them with a tracked pointer.

In surface-based methods, the shape of an exposed surface is extracted from all of the various coordinate systems by sampling a large number of points lying on the surface. The recorded shapes can then

be fitted to each other. Here, there is no one-to-one correspondence between the sampled points. Extracting such surfaces from medical images is a task that usually can be automated, and on the patient the surface can be sampled either by moving a tracked pointer across it or by using a three-dimensional (3D) camera or other kinds of automatic 3D scanning devices.

Finally, to register different images to each other, voxel-based methods are common. These methods are based on some image similarity metric, which quantifies how similar two images are. Assuming that the two images are most similar when they are perfectly aligned with each other, the spatial relationship can be found by investigating a range of possible relationships and searching for the one that optimises this similarity. The more different the images are, the more difficult it is to design an appropriate metric, and images from the same image modality will in general be easier to register to each other than images from very different modalities.

Often, a registration method will assume that the images or objects that are to be registered only differ in position and orientation and not in shape or size, as they are just different representations of the same object. This is referred to as rigid registration. However, since most parts of the human body are not rigid and may have different shapes at different times or in different poses, this assumption is often not accurate. Also, the geometric accuracy varies between image modalities and image sequences. In order to account for this, the registration method can also try to adapt the shape of one of the images to fit the other or to fit the patient on the operating table. This is often done by allowing other affine transformations, such as scaling and shearing, in addition to translation and rotation. However, since these are all linear transformations, they can only model global differences between the images or objects. In order to account also for local deformations, a completely elastic or nonrigid registration can be performed, but to do this deformation in a biomechanically correct manner is challenging.

## Visualisation

The medical images contain large amounts of information, but they are often subtle and difficult to interpret. Also, visualising the three-dimensional structures of the inside of the body in one two-dimensional (2D) or 3D scene is not straight forward: the structures are all packed together, some are hidden behind or inside others and some are big and some are small. It is therefore not possible nor desirable to visualise all the information in one view. However, there are, an endless amount of visualisation techniques that can be used to present as much of the relevant information as possible while attempting to make the interpretation of the images as intuitive as possible. For example, structures can be made transparent, or one can cut through them to show what is on the inside; the point of view can be moved around; information about the planned trajectory and critical structures to avoid can be added; and very small and detailed images can be overlaid on larger and coarser volumes showing the surrounding anatomy to facilitate the interpretation. An example of a typical visualisation used in navigated brain tumour surgery can be seen in Figure 1.5.

So far, navigation systems typically rely on ordinary 2D monitors for display of the navigation data. [8] However, new display technology can also aid in the visualisation and interpretation of the image data. Stereo displays, such as 3D monitors or head-mounted displays, have existed for a long time and can potentially aid the user with the depth perception. So far, the quality and usability of these solutions has hindered widespread adoption, but this may change. Augmented reality displays, which are partly transparent and allow for the images to be overlaid on what the user is actually seeing, is also an emerging technology that could be useful in a navigation setting. An example of this can be seen in Figure 1.6.

## 1.2 Surgical treatment of spine disorders

While surgical navigation has been widely used within many surgical disciplines for decades, its use in spine surgery has been limited. In this

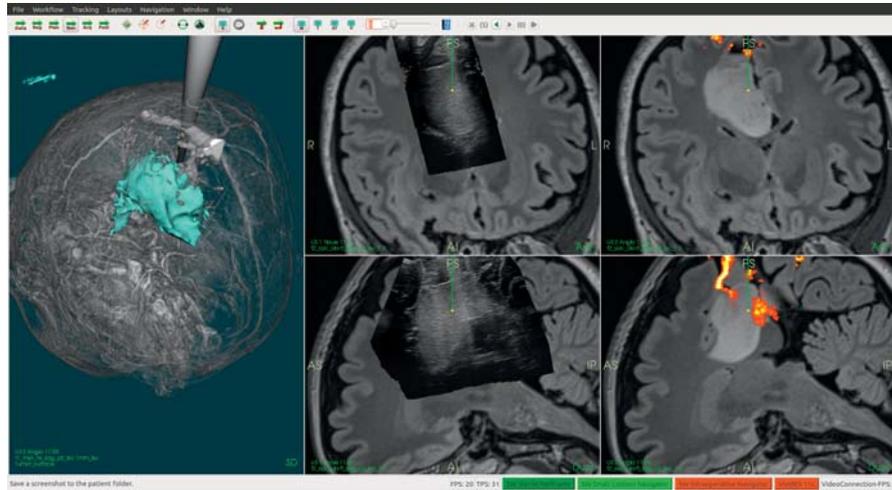


Figure 1.5: Visualisation of medical images during navigated brain tumour surgery using the open-source navigation system CustusX (SINTEF, Trondheim, Norway) [7]. The left-hand frame shows a 3D rendering of the blood vessels imaged with both magnetic resonance (MR) and ultrasound together with the tumour surface, which has been extracted from the MR images preoperatively. A model of the tracked surgical pointer that the surgeon is holding is also included in the scene. The rest of the frames show 2D slices through the MR and ultrasound volumes (B-mode in the middle and power doppler to the right) aligned with the tracked pointer. The shaft and tip of the pointer are here indicated by the green line and the yellow dot respectively.

thesis, the aim has been to investigate technology that could expand this use.

Most people experience back pain at some point in their life. Internationally, it has been found that the lifetime prevalence is between 49% and 84% [9–11], and a Norwegian study reported that 50% of all adults have experienced low back pain during the last year. [9] Back pain is also one of the most common reasons to visit primary care providers [10, 11], in Norway accounting for 10% of all visits. [9] As such conditions often affect the capacity for work, the societal costs



Figure 1.6: Doctor at the Oslo University Hospital using the augmented reality glasses HoloLens (Microsoft, Redmond, WA, USA) to examine a 3D model of a heart. (Image courtesy of Teknisk Ukeblad, <https://www.tu.no/>.)

of back pain are also substantial, mainly due to lost productivity and early retirement. In Norway, it is the nonlethal health issue that costs the most, amounting to between 13 and 15 billion kroner per year. [9]

Less than 5% of Norwegian back patients are treated surgically. [9] The most common indications for spine surgery in Norway today are disc herniation (49.5%), central spinal stenosis (29.1%) and lateral spinal stenosis (30.2%). [12] Here, it should be noted that each patient may have more than one diagnosis. Spinal surgery is performed both by orthopedic surgeons and by neurosurgeons, and the division of labour between these two groups varies greatly between hospitals. The most common intervention is discectomies for the treatment of disc herniations. [11] Surgical treatment obviously has a higher risk than nonoperative care, with the most common complications being dural tear, and postoperative wound infection. [11] Still, surgery has been shown to be both effective and cost effective compared to nonoperative care for both intervertebral disc herniation and spinal stenosis

in selected patients. [13]

### 1.3 Imaging of the spine

Imaging plays a crucial role in both diagnosis and surgical treatment of spine disorders. The two most important modalities for spine imaging today are computed tomography (CT) and magnetic resonance (MR) imaging. However, ultrasound imaging has more recently also found its use, particularly for intraoperative imaging.

#### Radiography and CT

2D projectional radiography, as seen in Figure 1.7(a), has traditionally been a useful tool for imaging of the spine. While it can be used for diagnosing both skeletal pain in the back and arthritic disorders [14], it is today mainly used for the assessment of scoliosis and spondylolisthesis. By injecting a contrast medium in the cerebrospinal fluid after lumbar puncture, radiography can also be used to evaluate lesions such as spinal stenosis or nerve root impingement due to disc herniations. [14] Now, this modality, called myelography, is mainly used when MR imaging is not feasible.

While conventional radiography is still in use, CT offers superior image quality with more anatomical and diagnostic information. [15] Therefore, as CT scanners have become more generally available, it has replaced conventional radiography for most diagnostic applications within spine imaging. A typical example of a 2D slice from a 3D CT volume can be seen in Figure 1.7(b). CT images bone very well, particularly hard cortical bone. [14, 16] It is thus well suited for evaluating spinal fractures, and is often used to identify traumatic injury. It is also used to assess sacroiliac joint abnormalities and to image spinal bone tumours. [16] The 3D nature of the CT images, which enable both visualisation of arbitrary image planes and 3D renderings, is also very useful for a number of applications, not least for surgical navigation.

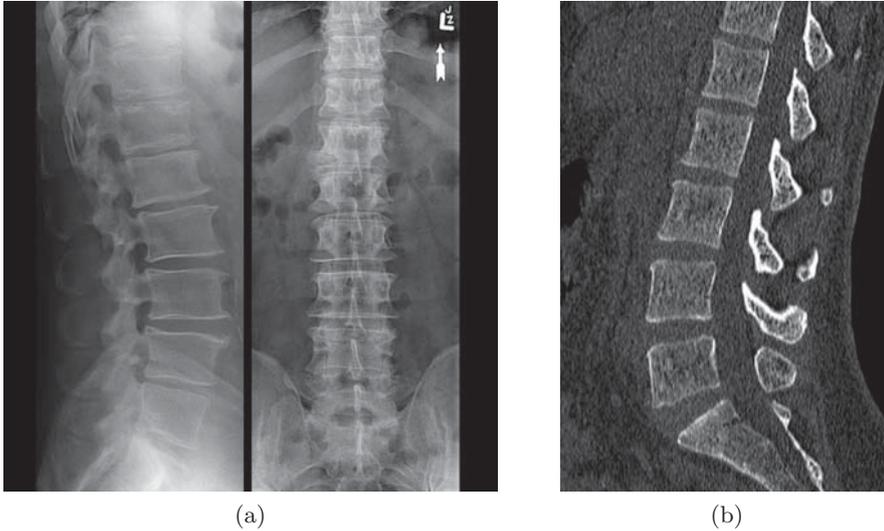


Figure 1.7: (a) 2D projectional X-ray images of the spine and (b) a 2D slice from a CT volume of the spine. (Images courtesy of (a) Associate Professor Frank Gaillard and (b) Dr. Ian Bickle, <https://radiopaedia.org/>.)

In spine surgery, intraoperative real-time 2D imaging, referred to as fluoroscopy, is widely used, and most image-guided spine interventions are still performed with only fluoroscopic guidance. [17] Modern fluoroscopic systems are usually configured as a C-arm, where the X-ray source and the detector are mounted on each end of a C-shaped arm. These can either be mobile or fixed to a rail system or a robotic arm in the operating room as seen in Figure 1.8. Such systems are very flexible and can produce live X-ray images at virtually any angle. In addition to real-time imaging, some C-arms can also produce CT-like 3D volumes, which also can be used for surgical navigation.

## MR

In spite of being both slower and more expensive than CT, MR imaging has come to be the most important imaging tool for spine pathology.



Figure 1.8: The C-arm Artis zeego from Siemens Healthineers (Erlangen, Germany) during imaging of an accuracy phantom. The C-arm is integrated with a floor-mounted robotic arm, which makes it very flexible, and can produce CT-like 3D volumes in addition to ordinary real-time fluoroscopic images.

[14] This is mainly due to its superior soft tissue imaging capabilities: MR imaging can produce detailed images of the spinal cord, the dural sac and the nerve roots within the neural foramina. [16] Extradural structures, such as the intervertebral discs, can also be imaged well. [14] An example of a 3D MR volume is shown in Figure 1.9.

The most common use of MR imaging is for the evaluation of local pain or radicular symptoms. It is also used in the diagnosis of degenerative disc disease, infection and tumours of the spinal cord, nerve roots and the vertebral column. In trauma, MR imaging is used to look for disc rupture, injury to ligaments or the spinal cord and

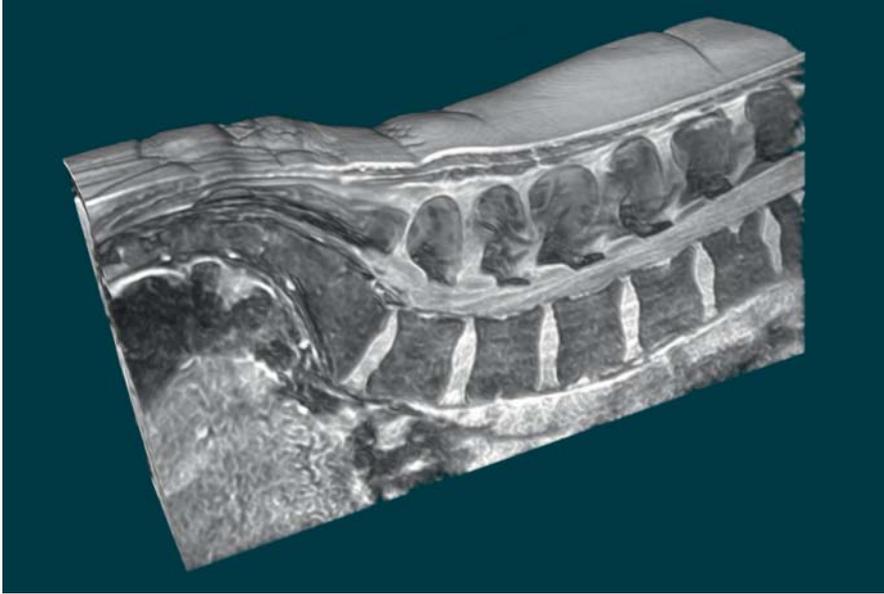


Figure 1.9: 3D rendering of a MR volume of the lumbar spine.

intraspinal hematoma. [16] It is also better than CT in determining the acuity of vertebral fractures. [14] Finally, MR imaging is used to evaluate inflammatory disorders such as multiple sclerosis, sarcoidosis and transverse myelitis. [16]

Due to its higher cost, large size and long acquisition times, MR imaging is not well suited for intraoperative use. The strong magnetic fields also pose challenges to the logistics in the operating room (OR). Although there are several intraoperative MR imaging systems on the market, in spine surgery, there is still disagreement as to their advantages over other techniques, and their use is therefore still very limited. [18]

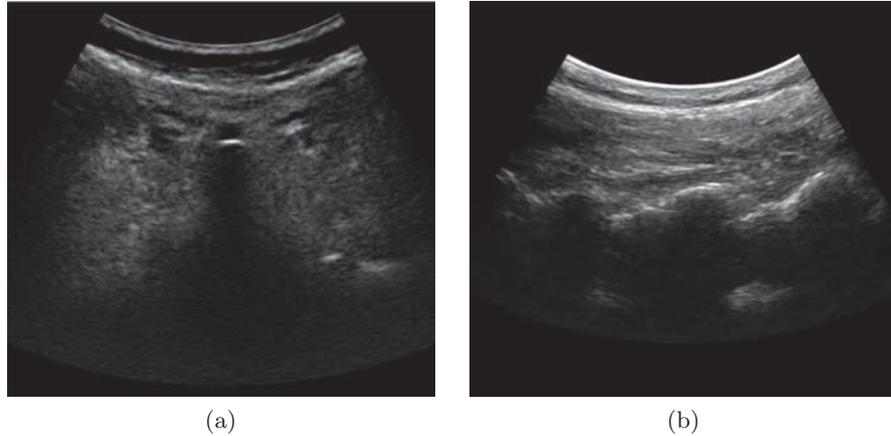


Figure 1.10: (a) Axial ultrasound image of the spine showing a clear shadow where the vertebral body is and a distinct reflection from the top of the spinous process, and (b) a sagittal ultrasound image showing reflections from the facet joints.

## Ultrasound

Due to the large difference in acoustical impedance between soft tissue and bone, ultrasound waves are unable to propagate properly across the interface between the two. [19] Because of this, it is in general not possible to image structures located inside or behind bone surface with ultrasound. This makes it unsuitable for evaluating most musculoskeletal conditions in the back. [14] It is, however, possible to image the bone surface and many of the structures surrounding the spine, and it can therefore be used for guiding various procedures. Examples of such images can be seen in Figure 1.10. This is particularly useful for percutaneous procedures, as the few natural anatomical landmarks on the back otherwise usually are identified through manual palpation, and this can in some cases be difficult.

The area where ultrasound guidance for spinal procedures has achieved the most widespread adoption in clinical practice is anaes-

thetia. Here, ultrasound can be used to guide both lumbar puncture, epidural procedures, nerve root blocks and facet joint injections. [20, 21]. The advantages of such ultrasound guidance over the traditional techniques guided by surface landmarks has not yet been established, but results from recent research are promising. [20]

While percutaneous ultrasound imaging of the spine has many limitations, intraoperative imaging of the spinal cord after removal of the bone has a lot of potential. This was first investigated as early as 1982 [22]. While its clinical adoption so far has been variable, advances in ultrasound technology has increased its utility as the scanners are becoming both cheaper and smaller and thus easier to integrate in the clinic. The most important development, however, is of course the improved imaging quality of modern ultrasound scanners, and intraoperative ultrasound has now been shown to accurately image a variety of different spinal lesions, as well as surrounding neural and vascular structures. [23–28] An example of this can be seen in Figure 1.11.

## 1.4 Navigation in spine surgery

Navigation systems for image-guided spine surgery have been available since the late 1990s, and today, all of the three major navigation platforms StealthStation, Brainlab and Stryker NAV3i have modules for spine surgery. [29–31] These systems are based on traditional optical tracking of the instruments in combination with preoperative CT images or intraoperative 3D images. While all three systems offer integration with various third-party C-arm systems capable of 3D imaging, both Medtronic and, more recently, Brainlab now have their own mobile, intraoperative CT systems; Medtronic's O-arm was introduced as early as 2006 and is a traditional cone beam CT system, while Brainlab's Airo, which was launched in 2014, is a full 32-detector-row CT system. [29, 32] The latter can be seen in Figure 1.12.

In spine surgery, navigation systems are still most frequently used in fusion procedures. The benefits of image guidance in such procedures, under given conditions, have been documented in several stud-



Figure 1.11: Ultrasound image of a spinal tumour as presented by Selbekk et al. [24].

ies, of which Moses et al. [18] give a good overview. These include improved pedicle screw placement accuracy, reduced radiation exposure and reduced risk of injury to neurovascular structures. In addition, the use of navigation systems may shorten the learning curve for less experienced surgeons compared to conventional landmark-based screw placement. However, while these improvements all may be beneficial, no study has so far demonstrated improved overall patient outcomes. [18, 33]

In 2013, Härtl et al. [33] performed a worldwide survey among spine surgeons on their use of navigation systems. They found that even though there is a lack of scientific evidence supporting its use, 80% of surgeons held positive opinions about image-guidance for spine surgery. The factors that were most commonly perceived as benefits among the respondents were increased accuracy, increased safety and



Figure 1.12: The 32-detector-row intraoperative CT system Airo from Brainlab (Munich, Germany). [32]

reduced radiation. In spite of this, only 9% of the surgeons used navigation routinely for fusion procedures, and 66% never used it. Neurosurgeons reported to use navigation slightly more often than orthopaedic surgeons. Finally, they found that the main reasons that navigation systems are not more widely adopted are high costs, increased OR time and lack of equipment and training. [33]

## 1.5 Emerging technologies

High cost is one of the most important barriers to a more widespread adoption of navigation systems in spine surgery. [33] Extending the use of navigation to more than just fusion procedures, and to soft tissue procedures such as disk herniations in particular, could help the surgeon substantiate the cost of the equipment. A big step in

this direction would be to enable navigation based on other imaging modalities than CT, such as MR or ultrasound. [34] Navigation has been used for spinal tumour surgery, but mainly for resection of primary bone tumours, which are readily visualised in the CT images used in current navigation systems. [34] By substituting or enhancing the CT with MR images, navigation could also be very useful for the resection of complex soft tissue tumours. To date, only D'Andrea et al. [35] have reported on the use of CT/MR coregistration for use in navigated spine surgery. [34]

As mentioned, another option for imaging soft tissue tumours is intraoperative ultrasound. Several groups have reported on the use of ultrasound for guiding spinal tumour surgery, and its usefulness both for the planning of the surgical approach, confirmation of the location and extent of the tumour and resection control has been demonstrated. [23–28] In some cases, intraoperative ultrasound even outperforms pre-operative MR imaging in terms of image quality. [24]

Recent advances in ultrasound technology could make these tools even more useful. In particular, silicon-based ultrasound arrays, such as the capacitive micromachined ultrasound transducers (CMUT), have a great potential as they can be made much smaller, cheaper and more robust than traditional arrays based on piezoelectric materials; they are thus perfect for intraoperative use. Due to the low production costs they could potentially even be made disposable. They also have a much higher bandwidth than other transducers, which means that the same probe can use a larger range of frequencies, thus enabling optimal image quality at all depths. [36]

Most of the current navigation systems for spine surgery use optical tracking systems to track the position and orientation of the surgical instruments and the patient. However, the interest in EM tracking systems is slowly increasing. While this is old technology, its use in surgical navigation is still quite limited. Both Medtronic and Brainlab currently have navigation systems that incorporate EM tracking, but these are so far only adapted to cranial surgery and mainly for guiding the placement of ventricle drains. [37, 38]

The main advantages of EM tracking over optical tracking are much

smaller position sensors that can be integrated in the surgical tools much more easily, and no line-of-sight requirement. The latter is particularly crucial in the case of flexible, endoluminal instruments. Both of these factors could also be a great advantage in spine surgery, as it is very difficult to attach a big, optical position sensor to the back in a sufficiently rigid manner and with a free line of sight to the tracking camera before the spine itself has been exposed. A small, electromagnetic sensor, on the other hand, could easily be attached directly to the back before the patient is draped and cleaned for surgery, making the logistics much better. [18]

More novel technologies include the use of 3D cameras for registration as in the new 7D Surgical System (7D Surgical, Toronto, ON, Canada), which was launched in January 2017. [39] Their navigation system is very similar to others on the market, but the technology they have implemented potentially simplifies and improves the registration step, which is crucial to achieve high accuracy throughout the procedure.

In July 2017, Philips released their new concept with high-resolution optical cameras integrated with their AlluraClarity interventional X-ray system. The cameras are used to create a 3D optical image of the patient that is registered to the 3D X-ray image. Live video images from the cameras can then subsequently be fused with X-ray images and used to guide the instruments during the following procedure without the need for external tracking systems. [40, 41]



## Chapter 2

# Aims of study

A major limitation of navigated spine surgery today is its reliance on X-rays, in the form of C-arms or mobile CT systems, for intraoperative imaging. As mentioned, CT is excellent for imaging bone, and thus suitable for orthopaedic surgery. In most other applications, however, MR is the modality of choice for spine imaging. [14] One way of bringing the MR images into the OR is obviously intraoperative MR imaging. Another option, which is both faster and much less costly, is to use intraoperative CT to register preoperative MR images to the patient, as suggested by D'Andrea et al. [35]. However, our proposal, and the ultimate goal for this work, has been to get rid of the CT altogether and use ultrasound imaging both for intraoperative imaging and for registration of the preoperative MR images to the patient.

In order to realise this goal, we have set out to study both methods for creating 3D volumes from 2D ultrasound images, the accuracy of EM tracking in an OR setting, ultrasound probe calibration methods and, finally, methods for registering MR images to ultrasound images of the spine. All of these topics of study are described in more detail in the following sections.

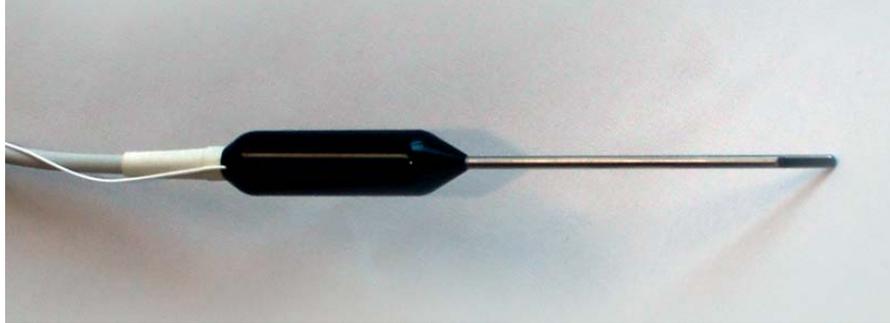


Figure 2.1: The custom-made L13-7 ultrasound probe for intraoperative spine imaging. The probe has a long, thin shaft, a small footprint of only  $13\text{ mm} \times 1.5\text{ mm}$  and a centre frequency of 10.3 MHz.

## 2.1 Intraoperative imaging

Previous efforts at intraoperative imaging of the spine using navigated ultrasound have made use of ordinary linear or phased array ultrasound probes with optical tracking. [23, 24, 26, 42] However, many spine procedures are today performed with a microsurgical approach, and the small incisions used in these procedures prohibit the use of standard probes. Our group has therefore, in a previous project, developed a new probe specifically designed to enable ultrasound imaging through such small incisions. The probe, which can be seen in Figure 2.1, has a long, thin shaft, a small footprint of only  $13\text{ mm} \times 1.5\text{ mm}$  and a centre frequency of 10.3 MHz. In order to use this probe for navigation purposes, it had to be integrated with our navigation system CustusX (SINTEF, Trondheim, Norway). [7] This posed several challenges.

First of all, in order to use the tracked 2D ultrasound images for navigation, it is often useful to reconstruct them into a 3D volume. The high resolution and the small image sector of the new ultrasound probe poses high demands on both the quality of the images, the accuracy of the tracking and the reconstruction algorithm. In most commercial ultrasound scanners, the raw ultrasound images are not readily available, and navigation systems therefore used to rely on digitised analogue

video for image transfer. However, the new probe interfaces with the ultrasound scanners from Ultrasonix Medical Corporation (Richmond, Canada), which were among the first to offer a research interface with access to all of the raw data. We therefore started by comparing both different 3D ultrasound reconstruction algorithms and the difference between reconstructions based on analogue video and digital raw data.

Secondly, the imaging array of the probe is situated on the side of the shaft. In order to image different parts of the anatomy and potentially create a 3D volume, the probe thus has to be rotated around the longitudinal axis of the shaft. This, in combination with the small size of the probe, makes it unsuitable for optical tracking. An EM tracking sensor was therefore integrated in the probe. EM tracking is, as mentioned, vulnerable to disturbances from ferromagnetic interference sources in the surroundings, and the small size and high resolution of the ultrasound image posed high demands on the tracking accuracy. One aim was therefore to study the accuracy of the tracking system in the OR setting.

A final challenge was what is referred to as probe calibration. This is the process of finding the spatial relationship between the coordinate system of the tracking sensor that is integrated in the ultrasound probe and the coordinate system of the ultrasound images generated by the probe. To achieve this, one typically employs some kind of imaging phantom, which is an object that can be easily imaged by ultrasound and whose position can be accurately determined in both coordinate systems. It quickly turned out that the existing phantoms could not easily be imaged by the new probe, due both to the placement of the imaging array and to the small size of the image sector. We therefore decided to develop a more versatile calibration method.

## **2.2 Registration of preoperative images to the patient**

The problem with registration of preoperative images to the patient in spine surgery is that there are no easily available natural landmarks

on the back. Often, this is solved by starting the procedure without navigation. Then, after the spine has been exposed, a reference is attached to one of the spinous processes and the registration is carried out using landmarks on the surface of the exposed vertebrae.

Another approach is to use intraoperative imaging for registration. The spatial relationship between the imaging equipment, such as the CT scanner or the ultrasound probe, and the images it produces is usually fixed. By attaching a tracking sensor to the equipment, the spatial relationship between the coordinate system of the tracking system and the image can also be found, e.g. through a calibration procedure such as the probe calibration mentioned in the previous section. Thus, by registering the preoperative images to the intraoperative ones, they are automatically also registered to the patient. This was what was proposed by D'Andrea et al. [35], who used images from an intraoperative CT scanner to register preoperative MR images to the patient. We aimed to do the same using ultrasound instead of CT.

## Chapter 3

# Equipment and infrastructure

The navigation system used in this work is the open-source system CustusX, which is a research platform focusing on intraoperative navigation and ultrasound imaging. Being a research platform, it is very flexible, incorporating several different reconstruction, registration and visualisation methods, a customisable user interface and support for a wide variety of ultrasound scanners. However, as the system has been developed in close collaboration with clinicians over many years, it also has a usability which makes it well-suited for clinical research. An example of the CustusX main window can be seen in Figure 1.5, and a more thorough description of the system can be found in Askeland et al. [7].

CustusX supports both the optical tracking system Polaris Spectra and the EM tracking system Aurora, both from NDI. These are among the most commonly used tracking systems for medical applications, and together they provide both high accuracy and great flexibility. In one of our studies, we used a new prototype field generator made by NDI for the Aurora system in addition to the standard field generator seen in Figure 1.4. The new field generator, which can be seen in Figure 3.1, had a torus shape allowing X-ray images to be taken through



Figure 3.1: The torus-shaped prototype field generator from NDI (Waterloo, ON, Canada) for the Aurora EM tracking system.

the centre opening. This made it suitable for use in combination with fluoroscopic imaging in an OR setting.

The navigation system also interfaces with the collaborative robot arm UR5 from Universal Robots (Odense, Denmark). This is a small, lightweight robot arm with six joints, a working radius of 850 mm and a lifting capability of 5 kg. The robot can be seen in Figure 3.2(a).

For ultrasound acquisition, we mainly used the ultrasound scanner SonixRP and its successor, the SonixMDP (Ultrasonix Medical Corporation, now part of Analogic Corporation, Peabody, MA, USA). These scanners were among the first clinically approved ultrasound scanners with a research interface, which grants the users access to and control over all of the internal settings and raw data that are used to produce the ultrasound images. For surgical navigation, the main advantage of

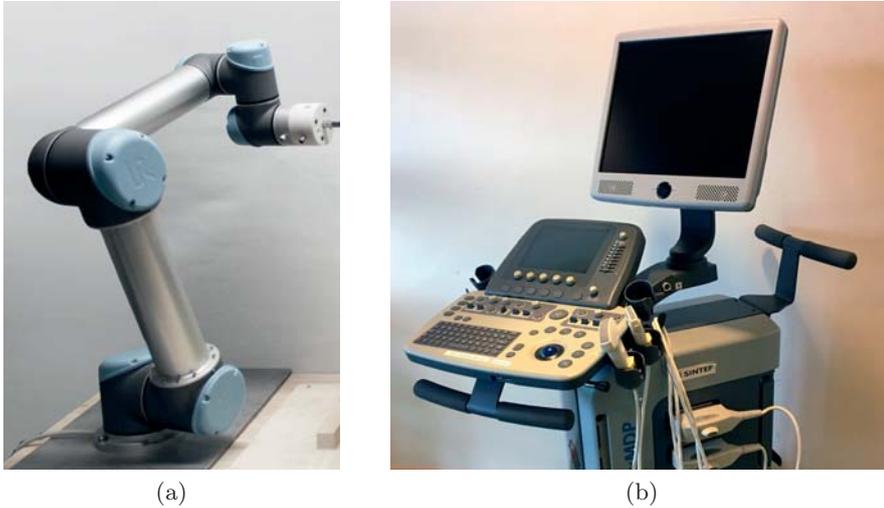


Figure 3.2: (a) The collaborative robot arm UR5 (Universal Robots, Odense, Denmark) and (b) the SonixMDP ultrasound scanner (Ultrasonix Medical Corporation, now part of Analogic Corporation, Peabody, MA)

this is the possibility to stream the digital raw images from the scanner with or without the postprocessing steps usually applied by ultrasound scanners. It also enables the navigation system to adjust automatically to changes in settings made on the scanner. The SonixMDP can be seen in Figure 3.2(b).

In addition to the Sonix scanners, we did in some cases use a Vivid E9 scanner from GE (Boston, MA, USA). This has a wider selection of probes and generally better image quality than the other two. Through an agreement with the manufacturer, we have gained access to a data interface on this scanner as well, and while this interface is very limited, it gives us the ability to stream raw image data and extract certain key parameters from the system.

For intraoperative imaging, we tested a customised spine probe designed by our group and manufactured by Vermon (Tours, France). The probe has been produced in two versions: one for the Vivid 7

scanner (GE, Boston, MA, USA) and one for the SonixMDP scanner. The latter, which can be seen in Figure 2.1, also has an EM tracking sensor for the Aurora tracking system (NDI, Waterloo, ON, Canada) integrated into the handle. The probe has a thin shaft, which is 12 cm long and 4 mm in diameter. The transducer array, which consists of 64 elements, measures only  $1.5 \text{ mm} \times 13 \text{ mm}$  and is situated on the side of the shaft. The centre frequency of the transducer is 10.3 MHz, and it has an axial and lateral resolution of 0.19 mm and 0.22 mm respectively. While originally designed for spine surgery, the probe has also been tested in transsphenoidal surgery of pituitary tumours. [43, 44]

The accuracy testing of the EM tracking system and the clinical testing of the spine probe were both performed in the facilities of the Operating Room of the Future at St. Olavs hospital (Trondheim, Norway). [45] This is a research infrastructure comprising several ORs in various operating clinics throughout the hospital. These are equipped with state-of-the-art equipment, such as imaging equipment, navigation systems and display technology. The research activity in the Operating Room of the Future is also supported by a number of research coordinators and scientific advisors.

## Chapter 4

# Summary of papers

### 4.1 Paper A: 3D ultrasound reconstruction algorithms from analog and digital data

Algorithms for the reconstruction of 3D image volumes from tracked 2D ultrasound images are still an important part of ultrasound-based navigation systems. The 2D images used to be acquired as analogue video, but with the appearance of new scanners with open research interfaces, it was suddenly possible to acquire digital data directly from the scanner. In this paper, we aimed at analysing how this affected the quality of the reconstructed 3D volumes. We also presented a new reconstruction method based on acoustical considerations, and compared this to an existing method.

The experiments were performed using the SonixRP ultrasound scanner with the L14-5/38 linear probe. The comparison of volumes from the various image sources and reconstruction algorithms was based on a number of different criteria: both the algorithms ability to correctly recreate removed data, the reconstructed volumes ability to display existing structures and the correctness of the reconstructed volumes geometry were measured.

The results showed that image volumes reconstructed from digital data were, overall, slightly better than volumes reconstructed from

analogue video. However, the differences were so small that we concluded that for many applications they would be negligible. The performance of the reconstruction algorithms was also very similar, and while there were differences between them in some of the tests indicating that they have different strengths and weaknesses, the overall result was that they could not be separated.

This paper was published in *Ultrasonics*, vol. 51, no. 4, pp. 405–419, 2011.

## 4.2 Paper B: Accuracy of electromagnetic tracking with a prototype field generator in an interventional OR setting

In this study, we wanted to explore the feasibility of using EM tracking in an OR setting, both alone and in combination with a robotic C-arm. More specifically, we wanted to find out how far away from the tracking volume large equipment, such as a C-arm, had to be in order to achieve an accuracy sufficient for navigation. We also compared the performance of the standard Aurora field generator with a new prototype designed specifically for use with fluoroscopic imaging equipment.

The experiments were carried out in one of the ORs of the Operating Room of the Future, which is equipped with the robotic C-arm Artis zeego (Siemens Healthcare, Forchheim, Germany). For the accuracy evaluation, we followed a protocol developed and presented by another group. [5, 46] The motivation for using this standardised protocol was to enable easy comparison with other studies. In line with Yaniv et al. [5], we also measured the robustness of the system, which was defined as its resilience to distortions due to tools and equipment in its surroundings.

We found that, as long as the C-arm was kept outside of the operating field, the tracking system performed similarly to the specifications presented by the manufacturer, and also to what had been reported by other groups. In this setting, the difference between the two field

generators was also very small. When the C-arm was moved into the operating field, the accuracy decreased considerably, with the standard field generator being more affected than the prototype field generator. However, the measurements were still stable as long as the C-arm did not move, and they were also not affected by fluoroscopic imaging. We thus concluded that the system was robust for a given configuration of the C-arm, and that by implementing a suitable static correction scheme, the tracking system and the C-arm could potentially be used together.

This paper was published in *Medical Physics*, vol. 39, no. 1, pp. 399–406, 2012.

### **4.3 Paper C: Versatile robotic probe calibration for position tracking in ultrasound imaging**

Probe calibration has been extensively studied, also by our group. [47, 48] One of the major challenges within this topic, is how to accurately determine the position of an imaged object within the ultrasound images. Although a wide range of imaging phantom designs have been proposed, the quality of the images that are obtained vary greatly with both operator, ultrasound probe and scanner settings, and this affects the accuracy of the subsequent image segmentation. This was also the problem we experienced when we tried to calibrate the new, small ultrasound probe that we had developed. The calibration method that we were using at the time, which was developed by Chen et al. [49], uses a phantom consisting of multiple thin wires that should be imaged, and it was very hard both to position the probe so that all the wires were visible, and to hold the probe still once the image was satisfactory. In this study, we therefore aimed at developing a calibration method that was sufficiently versatile to be used with any kind of probe with very little user interaction.

In order to achieve this, we employed a simple plastic sphere as an imaging target and mounted it on a small robot arm. The symmetry

of the sphere made it easy to image regardless of the orientation of the ultrasound probe, and the very accurate movements of the robot facilitated both the initial positioning of the sphere and the movement of the sphere through various parts of the image plane in a controlled manner. To evaluate the method, we applied it to three very different ultrasound probes and tested it on a separate imaging phantom, which was specially developed for accuracy measurements.

The calibration method was easily applied to all of the ultrasound probes that we tested it on, and the accuracy of the calibrations was similar or better to that reported by other groups. While being somewhat less automatic than some other methods, we concluded that its versatility made it well suited for a research setting, such as ours, where new and specialised probes are tested on a regular basis.

This paper was published in *Physics in Medicine and Biology*, vol. 60, no. 9, pp. 3499–3513, 2015.

#### **4.4 Paper D: Registration of MR to percutaneous ultrasound of the spine for image-guided surgery**

In this paper, we presented a method for registration of preoperative MR images to the patient by means of intraoperative ultrasound imaging using a tracked ultrasound probe. The main motivation for this approach was to enable the registration to be performed before the surgery starts and without the use of X-ray imaging.

To enable a proper representation of the vertebral column in the MR images, we first developed a customised MR imaging protocol. The method then segmented the posterior bone surface from both the ultrasound images and the MR images. While the ultrasound segmentation was done automatically, the MR segmentation was only semi-automatic. Finally, the two surfaces were registered to each other using a modified version of the Iterative Closest Point algorithm, which is a surface-based, rigid registration method. For this paper, the method was only tested on one subject, but the accuracy of the registration on

this subject was clinically relevant, and we concluded that the method was promising.

This paper was published in J. Yao, T. Klinder and S. Li (eds.) 2014, *Computational Methods and Clinical Applications for Spine Imaging*, Springer, pp. 209–218.



## Chapter 5

# Discussion and future work

This PhD project was established in conjunction with the project Vertebral Intraoperative Repair by Targeted Ultrasound imaging (VIRTUS), which was led by the company SonoWand AS (Trondheim, Norway) and partly funded by the Research Council of Norway. When we started our work towards using 3D ultrasound for guiding spine surgery, there were many technical challenges that needed a solution. Some of them were clear from the start, while others appeared as we went along. This took a lot of focus in the first part of the project. Unfortunately, due to financial problems, SonoWand was forced to close its operations halfway through the project period. The VIRTUS project was thus terminated before it reached its conclusion, which meant that we were not able to realise a complete system ready for clinical testing the way we had planned. As a consequence, the main emphasis of this thesis shifted somewhat away from clinical spine application and towards enabling technologies such as tracking, reconstruction and registration, which are common for many ultrasound-guided procedures. In this section, both the work presented in the included publications and some work that has not been published will be discussed.

## 5.1 Intraoperative ultrasound imaging

A lot of effort was invested in optimising tracking, calibration and reconstruction methods for the new spine probe in order to generate high-quality 3D image volumes for navigation. However, this was not tested on any spine patients in this project. The main reason for this was that the initial tests of the probe, which were conducted before the tracking was in place on a few patients with lumbar disk herniations, were not satisfactory. In some patients, the 2D images provided a good view of the dural sac, the spinal nerves and the hernia pressing on these. An example of this can be seen in Figure 5.1. However, due to the right angle between the ultrasound probe shaft and the image sector, it was very difficult to image structures located deeper than the bottom of the narrow working channel. This includes the greater portion of the intervertebral disc and nerve roots exiting from the spinal canal. In many cases, depending on the position of the herniation, this prohibited a good representation of the relevant anatomy. Also, some of the cases where intraoperative imaging would potentially be most helpful are reoperations, where scar tissue makes it difficult to navigate visually. Unfortunately, the image quality of the new probe was not deemed sufficiently high for these cases.

While the probe was never used for 3D imaging of the spine, it was tested in transsphenoidal surgery of pituitary tumours, as reported by Selbekk et al. [43] and Solheim et al. [44]. In this application, the images' quality was sufficient, with a resolution superior to that of clinical MR imaging. Important structures, such as the optic nerves, the ophthalmic arteries and the cerebral arteries were well visualised, although also here, the right angle of the image sector was a limitation. Later, we were also able to produce some 3D volumes that were of good quality, as can be seen in Figure 5.2. However, the EM tracking proved unreliable in this clinical setting; we regularly experienced large disturbances, often with complete loss of tracking. While we were able to eliminate all major metal objects from the tracking volume during navigation, e.g. by replacing the metal head rest with one made from PEEK plastic, the nasal speculum that was used was still

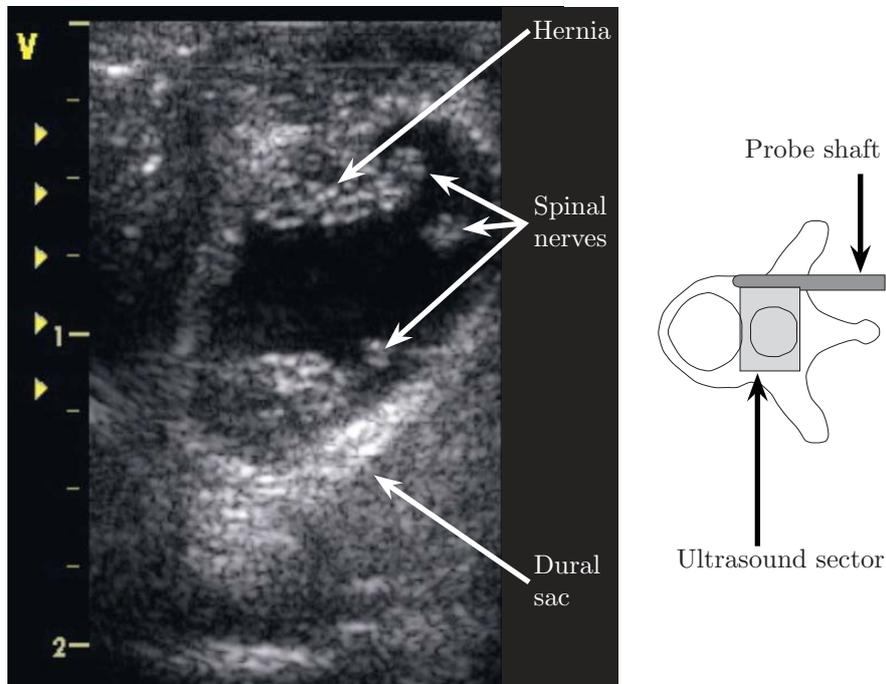


Figure 5.1: Ultrasound image of a lumbar disc herniation. It shows the spinal nerves of the cauda equina (bright spots) floating in the cerebrospinal fluid (black area in the centre) contained by the dural sac. The sketch on the right indicates the approximate position of the ultrasound image sector.

made from steel, and we suspect that this may have been the source of the disturbances.

In order to address the shortcomings of the spine probe presented here, we have, in cooperation with SonoWand AS, made a new prototype ultrasound probe seen in Figure 5.3. The most important change is that it has a curved array, which allows for imaging to the side and forward at the same time. The trade-off is that the tip of the probe is quite a lot bigger than that of the original probe, which means that it does not fit in the narrowest working channels. The probe has already been tested in pituitary surgery with good results, and we hope to be

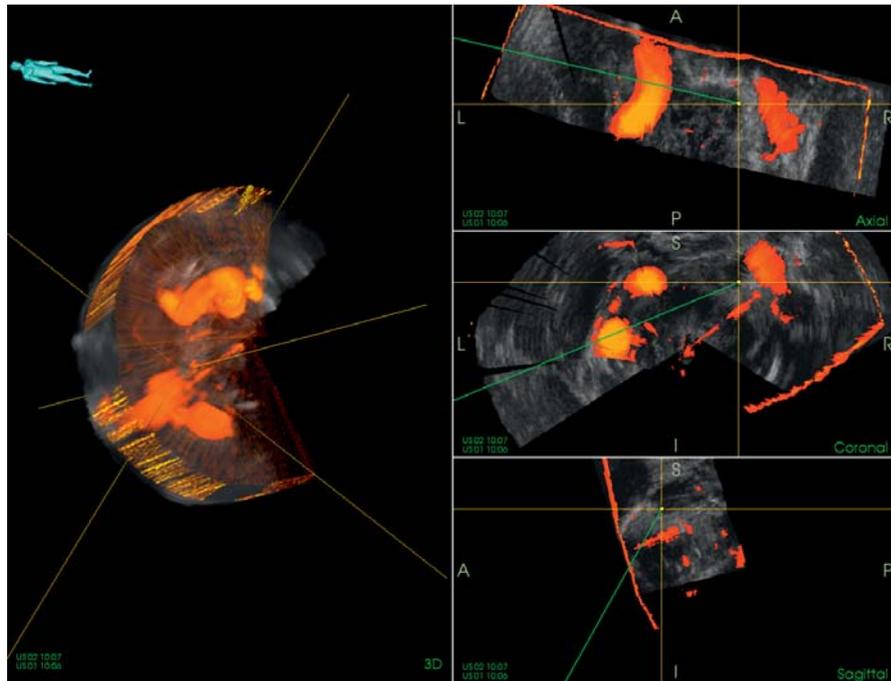


Figure 5.2: Scene from the navigation system CustusX (SINTEF, Trondheim, Norway) [7] during transsphenoidal surgery of a pituitary tumour. The left-hand frame shows a 3D rendering of two reconstructed ultrasound volumes, one based on ordinary B-mode images (grey scale) and one based on power Doppler images, which highlight the blood flow (red and yellow). The yellow lines in the scene indicate the position of the navigated pointer. The right-hand frames show orthogonal 2D slices through the volumes. The ultrasound images were recorded with the spine probe shown in Figure 2.1.



Figure 5.3: Prototype ultrasound probe for pituitary and spine surgery with a curved array, which allows for imaging to the side and forward at the same time.

able to test it in spine surgery as well.

## 5.2 Reconstruction

There are a number of different approaches to reconstruction of 3D image volumes from 2D ultrasound images, and good overviews can be found in the papers by Solberg et al. [50] and Mozaffari et al. [51]. Each approach has its strengths and weaknesses, and it has become evident to us that a reconstruction method that is suited for one application might not be suited for another.

In Paper A, we focused among other things on how the methods handle areas where information is missing. However, when we tried to image the bone surface of the spine, the biggest challenge was how to handle areas with conflicting information where multiple ultrasound images overlapped. The reason for this was that the angle between the ultrasound beam and the imaged surface must be close to  $90^\circ$  in order to get a good reflection from the bone surface. For the central parts of the spine, this means that the ultrasound probe must be angled towards the patient's midline on both sides, resulting in a lot of overlapping images. The bone surfaces also produce a lot of artefacts, such as shadows and reverberations, in the ultrasound images. Some

of these images may therefore contain only noise in areas where other images contain actual information, and it is often impossible to distinguish the two by just comparing the small area where the images overlap. In these cases, the most common solution is therefore to enhance or completely segment the bone surface in the 2D ultrasound images before reconstruction. [52–55] This was also the approach that we chose in Paper D.

In order to address the varying requirements of different applications, our group has more recently developed a new reconstruction method. [56] This uses a more sophisticated and adjustable scheme for determining the contribution from each pixel in the 2D ultrasound images to each voxel in the reconstructed volume; by automatically adjusting the weighting according to the variance of the input data, the method tries to smooth out noise while at the same time retaining detail in high-frequency regions. In cases with conflicting information, it can add extra weight to input data with high signal or to input data that is more recent, for example data from the last of two sweeps with the probe. Which weighting gives the best result, depends highly on the imaged anatomy, imaging technique and equipment, and the parameters of the method must therefore be tuned to the specific application. While we experienced that proper tuning of the parameters improved the quality also of the spine images, it was not sufficient to avoid preprocessing of the 2D ultrasound images.

### 5.3 Tracking accuracy

#### Optical tracking

Pose tracking is an important part of any surgical navigation system, and it is a field that requires constant attention. Optical tracking systems, in particular used with passive, reflective spheres, are much used due to their simplicity and high accuracy under good conditions. Their big weakness is, of course, that they require a free line of sight between the camera and the tracked tool. In practice, both in the OR and in the lab, this can be surprisingly difficult to achieve; when the

tracked tools are combined with other equipment and move around, it takes very little before one of the many reflective sphere disappears out of sight. If it is completely out of sight, tracking typically stops, and it may take considerable effort to reposition the camera so that navigation can continue. While this can be frustrating, it does not affect the tracking accuracy. However, if the sphere is only partially occluded, it can severely affect the accuracy in a way that is difficult to detect during operation.

This was something we experienced during our work with the new probe calibration method presented in Paper C. In the initial lab setup, we used a linear robot that was mounted above the water tank. The distance between the aluminium posts supporting the robot was 50 cm, which for the first tests provided a good field of view for the tracking camera. However, as we tested different probes with different geometries, it was often a lot of work to adjust the positions of the robot, the calibration arm and the probe so that tracking was not obstructed by the aluminium posts. Also, by making tiny adjustments until all tools were just within sight of the tracking camera, the risk that some of the tools were still partially occluded was quite high. As a result, we several times experienced large errors in the calibration results that we were unable to explain. After thorough investigation, it turned out that it was caused by inaccurate tracking, which prompted us to change the entire setup.

The line-of-sight requirement also has major implications for the design of tracked tools. The reflective spheres must be mounted so that the tool can be used in a intuitive way without the user's hands or other objects in the surroundings coming in the way. This often means that the distance from the spheres to the tip of the tool must be relatively long, resulting in a reduced tracking accuracy at the tip and ergonomic challenges for the surgeon. Also, while increasing the distance between the spheres increases the accuracy, it also increases the area that must be kept clear to avoid occlusion, and it may hamper direct visualisation in the surgical field, for example when using surgical microscopy in small surgical corridors. This is an import trade-off.

In surgical navigation, the Polaris system from NDI has com-

pletely dominated the market for many years. Recently, the company Metronor AS (Oslo, Norway) also launched an optical tracking system for medical applications. [57] This system employs a single, small camera instead of two cameras, like most of its competitors, and it uses wireless tools with infrared light-emitting diodes instead of reflecting spheres as light sources. This way the line-of-sight issues are reduced to a minimum, and accuracy is increased significantly. Also, due to the small size of the diodes, partial occlusion is not an issue. We are currently working to integrate the Metronor system with our navigation system CustusX in order to test its capabilities.

### **EM tracking**

For EM tracking, disturbances from metal objects in the surroundings is the main challenge. In Paper B, we studied the accuracy in one of the ORs of the Operating Room of the Future. Here, our group has since applied EM-based navigation to numerous procedures, mainly within interventional radiology and bronchoscopy, without problems. Before initiating the experiments with 3D imaging in pituitary surgery, we did similar investigations in the neurosurgical suite of the Operating Room of the Future. Here, the setup included a different operating table and a different C-arm in addition to a surgical microscope, but the results were still very similar to those presented in the paper.

Based on previous experience, we assumed that smaller metal equipment, such as the surgical tools, would not interfere significantly with the tracking system and could thus be omitted in the accuracy analysis. This seems to be correct in many cases, but our experience from pituitary surgery, and also initial tests of the setup for spine surgery, indicate that it does not always hold true. In particular, we have seen that the retractors or specula used in microsurgery to hold the working channel open pose a problem. While these are small instruments, the other instruments have to pass through them. A tracking sensor integrated in an instrument may thus end up being completely surrounded by metal, which would introduce severe disturbances. This could be solved by integrating the sensor in a part of the instrument that does

not enter the working channel, such as the end of the handle. However, increasing the distance from the sensor to the tip of the instrument in this way has a negative effect on the tracking accuracy. An alternative is to replace the retractor or speculum with a radiolucent or MR-compatible version, which is not ferromagnetic. If this would be sufficient to eliminate the disturbances remains to be seen.

## 5.4 Probe calibration

The probe calibration method that we presented in Paper C, was designed to be easily applied to any ultrasound probe. Since it was finished, we have used it on a number of probes, including both linear array, curved array and phased array probes with both optical and EM tracking, without any need for modifications and with good results. It has, in other words, worked as intended. The method has also been adopted by a cooperating group at TU Delft and the Academic Medical Center in the Netherlands.

One drawback with this method today is that the position of the imaging target is measured with the same optical tracking system that we use for navigation. The accuracy of this measurement is thus of the same order as the tracking of the instruments during navigation, and in practice it therefore adds to the total error of the calibration. An easy way to improve the accuracy of the method would therefore be to manufacture the calibration arm with higher accuracy than that of the tracking system. In particular, this could be facilitated by the increasing availability and accuracy of 3D printers, and this is something we are currently looking into.

In the paper, we mention that the method has a great potential for further automation. However, in our experience, the steps that are most time consuming are really the setup of the ultrasound probe, the ultrasound scanner and the navigation system. There are not many minutes to be saved on the rest of the process. Also, for any automatic step there is risk that it fails when applied to a new probe, which would make the method less versatile. This is thus not something that we

have prioritised.

## 5.5 Registration

While we believe intraoperative ultrasound imaging holds a great potential in spine surgery, it can not replace X-ray or MR imaging completely. The main reason for this is its inability to image through bone. In order to replace intraoperative X-ray imaging, the registration of preoperative CT or MR images to the patient is therefore crucial. As mentioned, the lack of natural landmarks on the back makes registration challenging, and the obvious solution is to register the images to the intraoperative images rather than directly to the patient.

In Paper D, we presented a feature-based registration method for registration of MR images to ultrasound images of the spine. The feature that we used was the bone surface of the spine, and this was segmented from both modalities before registration. The biggest challenge was the segmentation of the ultrasound images. Here, we based our method on work by Jain et al. [58] and Foroughi et al. [59], who defined certain characteristics of bone reflections in ultrasound images and used this to calculate the probability that a given pixel is part of a bone surface. By doing this for each pixel in the image, we created a bone probability map, and we then extracted the bone surface from this map using backwards scan line tracing as described by Yan et al. [60]. While the initial attempts were successful, as reported in the paper, we were not able to find a single set of parameters that would work well for all cases. Also, it was very hard to acquire images of the same high quality from all subjects.

Due to these difficulties, we decided to seek another approach. This time, we wanted to make the ultrasound segmentation more robust by including known information about the general geometry of the spine. We therefore adopted a method developed and implemented by the group of professor Purang Abolmaesumi at the Robotics and Control Laboratory (RCL) at the University of British Columbia (Vancouver, Canada). The method made use of a statistical shape model of the

lumbar spine created from 32 CT images, which had been manually segmented. [61] This was automatically fitted to the ultrasound volume with a technique based on a Gaussian mixture model. [62] An example of this can be seen in Figure 5.4.

In spite of a close collaboration with the group at the RCL, who generously provided us with their own implementation of the method and gave us a careful introduction to its use, we were not able to make it work consistently. Even when the ultrasound volumes seemed to be of good quality and in accordance with the instructions, the registration would regularly fail, at least for a couple of vertebrae. However, as in our previous registration efforts, the main challenge was to produce ultrasound volumes of consistent quality between subjects. In some cases, such as the one shown in Figure 5.4, the bone surfaces would be clearly visible with bright reflections and well-defined shadows, resulting in a good registration. In other cases, these features were much less pronounced in spite of the fact that all of the subjects were below 50 years of age and normal weight. While it was mostly feasible to produce single ultrasound images of good quality, it was often very challenging to move the probe so that this quality was maintained throughout the acquisition.

We can not be certain, but we now believe that the image quality was the main reason we were not able to reproduce the results of the RCL group. This is corroborated by the fact that most of their ultrasound volumes were acquired by a sonographer with vast experience in spine imaging, while we acquired our data ourselves. Ultrasound imaging is notoriously user dependent, and while we as technologists have extensive knowledge of, and also a certain amount of experience with, how to perform ultrasound examinations, we suspect that this application requires more thorough training. For our use, in a registration method for spine surgery, this is a problem: as the registration step is subsidiary to the surgery itself, it should be easy both to learn and to use. A registration method that requires extensive training and is highly user dependent is not likely to be widely accepted and adopted in the clinic.

A recent trend that can assist in making ultrasound segmentation,

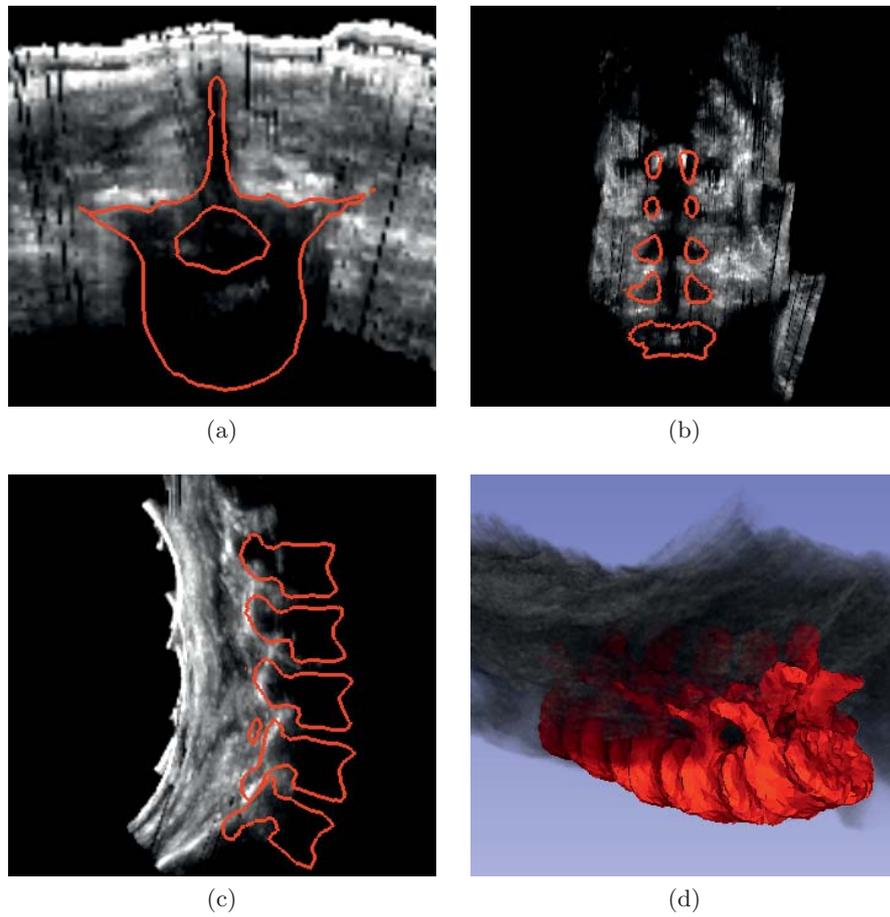


Figure 5.4: A statistical shape model registered to a 3D ultrasound volume of the lumbar back using the method developed and implemented by Rasoulian et al. [63]. The first three figures show (a) axial, (b) coronal and (c) sagittal slices through the ultrasound volume and the registered model, while the last (d) shows a 3D rendering of the volume and the model.

and thus ultrasound registration, less dependent on image quality and user skills is machine learning. Rather than defining the characteristics of bone reflections explicitly, as most of the previous methods do, the machine learning methods can learn the characteristics on their own based on training data consisting of correctly segmented example images. Given that the set of training data is both large enough and, just as important, representative for the kind of images that typical users will produce, such methods tend to be much more robust than methods based on explicitly defined criteria. Our group has previously applied a machine learning technique called deep convolutional neural networks to the segmentation of blood vessels from ultrasound images with good results [64]. More recently, both this method and another machine learning technique called random forests have been used to segment bone surfaces from ultrasound images, showing great potential. [55, 65, 66] We hope to be able to test this in the near future and apply it to the spine surgery application.

Another possibility for improving the detection of bone surfaces using ultrasound is to optimise the imaging method itself. This was also studied as a part of the VIRTUS project, and based on this work, Rodriguez-Molares et al. [67] proposed a method for imaging specular reflectors using an ordinary ultrasound transducer. The method is based on the mirror formula for spherical reflectors, and the results so far suggest that it is capable of reconstructing a specular reflector of arbitrary shape in a robust manner. Recently, the same group has presented a beamforming technique tailored to the physics of specular reflection. [68] The main challenge with both of these methods is that the specular reflector needs to be angled towards some part of the ultrasound transducer so that some of the acoustic energy can be recorded. The small size of ordinary ultrasound transducers, particularly in the elevation direction, thus greatly limits the shapes that can be recorded. A possible solution to this could be to specialised probes with large 2D transducers. [68]



## Chapter 6

# Conclusion

Surgical navigation can in many cases reduce the invasiveness and risk and improve the outcome of medical interventions. In spine surgery, this potential has only been realised to a very limited degree, and two thirds of the surgeons never use it. The main reasons for this are high costs, increased OR time and lack of equipment and training. There is thus a need for solutions that are both cheaper and easier to use. They should also be applicable to a wider range of procedure, which would both help substantiate the cost of the equipment and make the personnel more proficient at its use.

Ultrasound imaging could provide such a solution. Ultrasound scanners are both inexpensive and compact, a development which could be further accelerated with the introduction of CMUT-based probes. It also has the advantage over X-ray imaging that it does not produce any harmful radiation. While the compact equipment may provide a more streamlined workflow, both the acquisition and the interpretation of ultrasound images requires training, and navigation and image processing methods that can assist in this, is therefore important to make the technology accessible.

For intraoperative ultrasound imaging, we have developed and tested technology that enables tracking of small, intraoperative ultrasound probes and allows the generation of 3D volumes suitable for

navigation from such images. The main challenge here is to design a probe that both can reach the target area through a narrow working channel and, at the same time, produces images of a quality that enables imaging of the anatomy in question under varying conditions. We have come close to this with our first spine probe, and hope that the second generation may perform even better. A CMUT probe could provide a great improvement, both when it comes to design and image quality.

We have also investigated the use of intraoperative ultrasound imaging for registration of preoperative CT and MR images to the spine. The methods that we have tested, were not sufficiently robust for clinical use. However, several advances in the field of ultrasound segmentation and interpretation have been reported more recently, particularly based on machine learning techniques, and we therefore believe that a fully automatic registration method suitable for the clinic could be realisable in the near future.

# References

1. Spiegel, PK. The First Clinical X-Ray Made in America—100 Years. *American Journal of Roentgenology* 1995;164:241–243.
2. Birkfellner, W, Hummel, J, Wilson, E, and Cleary, K. Tracking Devices. In: *Image-Guided Interventions: Technology and Applications*. Ed. by Peters, T and Cleary, K. New York, NY, USA: Springer, 2008. Chap. 2:23–44.
3. Wiles, AD, Thompson, DG, and Frantz, DD. Accuracy assessment and interpretation for optical tracking systems. In: *Medical Imaging 2004: Visualization, Image-Guided Procedures, and Display*. Ed. by Robert L. Galloway, J. Vol. 5367. Proceedings of SPIE. SPIE, 2004.
4. Bø, LE, Leira, HO, Tangen, GA, Hofstad, EF, Amundsen, T, and Langø, T. Accuracy of electromagnetic tracking with a prototype field generator in an interventional OR setting. *Medical Physics* 2012;39:399–406.
5. Yaniv, Z, Wilson, E, Lindisch, D, and Cleary, K. Electromagnetic Tracking in the Clinical Environment. *Medical Physics* 2009;36:876–892.
6. Jolesz, FA. Introduction. In: *Intraoperative Imaging and Image-Guided Therapy*. Ed. by Jolesz, FA. New York, NY, USA: Springer, 2014. Chap. 1:1–23.

7. Askeland, C, Solberg, OV, Bakeng, JBL, Reinertsen, I, Tangen, GA, Hofstad, EF, Iversen, DH, Våpenstad, C, Selbekk, T, Langø, T, Hernes, TAN, Leira, HO, Unsgård, G, and Lindseth, F. CustusX: an open-source research platform for image-guided therapy. *International Journal of Computer Assisted Radiology and Surgery* 2016;11:505–519.
8. Pieper, SD. Visualization and Display for Image-Guided Therapy. In: *Intraoperative Imaging and Image-Guided Therapy*. Ed. by Jolesz, FA. New York, NY, USA: Springer, 2014. Chap. 7:107–115.
9. Lærum, E, Brox, JI, and Storheim, K, eds. Nasjonale retningslinjer for korsryggsmerter. Oslo, Norway: Norwegian Directorate of Health and Social Services, 2007.
10. Rampersaud, YR and Macwan, K. Economics of Minimally Invasive Spine Surgery. In: *Minimally Invasive Spine Surgery*. Ed. by Phillips, FM, Lieberman, IH, and Jr., DWP. New York: Springer, 2014. Chap. 3:23–33.
11. Spiker, WR and Lawrence, BD. Lumbar Disc Herniation. In: *Spine Surgery Basics*. Ed. by Patel, VV, Patel, A, Harrop, JS, and Burger, E. Berlin, Germany: Springer, 2014. Chap. 15:203–214.
12. Solberg, TK and Olsen, LR. Nasjonalt kvalitetsregister for ryggkirurgi (NKR): Årsrapport for 2015 med plan for forbedringstiltak 2016. Tech. rep. The University Hospital of North Norway, 2016.
13. Tosteson, ANA, Tosteson, TD, Lurie, JD, Abdu, W, Herkowitz, H, Andersson, G, Albert, T, Bridwell, K, Zhao, W, Grove, MR, Weinstein, MC, and Weinstein, JN. Comparative effectiveness evidence from the spine patient outcomes research trial: surgical versus nonoperative care for spinal stenosis, degenerative spondylolisthesis, and intervertebral disc herniation. *Spine* 2011;36:2061–2068.

14. Corder, HJ. Patient Evaluation and Criteria for Procedure Selection. In: *Image-Guided Spine Interventions*. Ed. by Mathis, JM and Golovac, S. New York, NY, USA: Springer, 2010. Chap. 3:39–55.
15. Alshamari, M, Geijer, M, Norrman, E, Lidén, M, Krauss, W, Wilamowski, F, and Geijer, H. Low dose CT of the lumbar spine compared with radiography: a study on image quality with implications for clinical practice. *Acta Radiologica* 2016;57:602–611.
16. Stone, J. Imaging Techniques in the Adult Spine. In: *Imaging of the Spine*. Ed. by Naidich, T, Castillo, M, Cha, S, Raybaud, C, Smirniotopoulos, J, and Kollias, S. Philadelphia, PA, USA: Elsevier Health Sciences, 2011. Chap. 1.
17. Mathis, JM. Materials Used in Image-Guided Spine Interventions. In: *Image-Guided Spine Interventions*. Ed. by Mathis, JM and Golovac, S. New York, NY, USA: Springer, 2010. Chap. 2:29–38.
18. Moses, ZB, Mayer, RR, Strickland, BA, Kretzer, RM, Wolinsky, JP, Gokaslan, ZL, and Baa, AA. Neuronavigation in minimally invasive spine surgery. *Neurosurgical Focus* 2013;35:E12.
19. Angelsen, BAJ. Principles of medical ultrasound imaging and measurements. In: *Ultrasound Imaging. Waves, signals and signal processing*. Vol. 1. Trondheim, Norway: Emantec, 2000. Chap. 1.
20. Tildsley, P, Lim, MJ, and Sng, BL. Ultrasound of the lumbar spine: Applications and advances. *Trends in Anaesthesia and Critical Care* 2016;7–8:36–40.
21. Loizides, A, Peer, S, Plaikner, M, Spiss, V, Galiano, K, Obernauer, J, and Gruber, H. Ultrasound-guided injections in the lumbar spine. *Medical Ultrasonography* 2011;13:54–58.
22. Dohrmann, GJ and Rubin, JM. Intraoperative Ultrasound Imaging of the Spinal Cord: Syringomyelia, Cysts, and Tumors—A Preliminary Report. *Surgical Neurology* 1982;18:395–399.

23. Kolstad, F, Rygh, OM, Selbekk, T, Unsgaard, G, and Nygaard, ØP. Three-dimensional ultrasonography navigation in spinal cord tumor surgery. *Journal of Neurosurgery: Spine* 2006;5:264–270.
24. Selbekk, T, Solheim, O, Tangen, GA, Solberg, OV, Bø, LE, Unsgård, G, and Nygård, Ø. Navigational high-resolution 3D ultrasound in spine tumour surgery — technical assessment. *International Journal of Computer Assisted Radiology and Surgery* 2009;4. Notes:S56–S57.
25. Zhou, H, Miller, D, Schulte, DM, Benes, L, Bozinov, O, Sure, U, and Bertalanffy, H. Intraoperative ultrasound assistance in treatment of intradural spinal tumours. *Clinical Neurology and Neurosurgery* 2011;113:531–537.
26. Shamov, T, Eftimov, T, Kaprelyan, A, and Enchev, Y. Ultrasound-based neuronavigation and spinal cord tumour surgery – marriage of convenience or notified incompatibility? *Turkish Neurosurgery* 2013;23:329–335.
27. Prada, F, Vetrano, IG, Filippini, A, Bene, MD, Perin, A, Casali, C, Legnani, F, Saini, M, and DiMeco, F. Intraoperative ultrasound in spinal tumor surgery. *Journal of Ultrasound* 2014;17:195–202.
28. Ivanov, M, Budu, A, Sims-Williams, H, and Poeta, I. Using Intraoperative Ultrasonography for Spinal Cord Tumor Surgery. *World Neurosurgery* 2017;97:104–111.
29. StealthStation Surgical Navigation System. Medtronic. URL: <http://www.medtronic.com/us-en/healthcare-professionals/products/neurological/surgical-navigation-systems/stealthstation.html> (visited on 2017).
30. Spinal Navigation Application. Brainlab AG. URL: <https://www.brainlab.com/en/surgery-products/overview-spinal-trauma-products/spinal-navigation/> (visited on 2017).
31. SpineMap 3D Software. Stryker Corporation. URL: <https://nse.stryker.com/products/spinemap-3d-software/> (visited on 2017).

32. Airo. Brainlab AG. URL: <https://www.brainlab.com/en/surgery-products/overview-platform-products/intraoperative-ct/> (visited on 2017).
33. Härtl, R, Lam, KS, Wang, J, Korge, A, Kandziora, F, and Audigé, L. Worldwide survey on the use of navigation in spine surgery. *World Neurosurgery* 2013;79:162–172.
34. Overley, SC, Cho, SK, Mehta, AI, and Arnold, PM. Navigation and Robotics in Spinal Surgery: Where Are We Now? *Neurosurgery* 2017;80:S86–S99.
35. D’Andrea, K, Dreyer, J, and Fahim, DK. Utility of Preoperative Magnetic Resonance Imaging Coregistered with Intraoperative Computed Tomographic Scan for the Resection of Complex Tumors of the Spine. *World Neurosurgery* 2015;84:1804–1815.
36. Lindseth, F, Langø, T, Selbekk, T, Hansen, R, Reinertsen, I, Askeland, C, Solheim, O, Unsgård, G, Mårvik, R, and Hernes, TAN. Ultrasound-based Guidance and Therapy. In: *Advancements and Breakthroughs in Ultrasound Imaging*. Ed. by Gunarathne, G. Rijeka, Croatia: InTech, 2013.
37. Cranial Neurosurgery Navigation. Medtronic. URL: <http://www.medtronic.com/us-en/healthcare-professionals/products/neurological/surgical-navigation-systems/stealthstation/cranial-neurosurgery-navigation.html> (visited on 2017).
38. Cranial EM Navigation Application. Brainlab AG. URL: <https://www.brainlab.com/en/surgery-products/overview-neurosurgery-products/cranial-em-navigation-application/>.
39. Empowering Spine Surgeons. 7D Surgical. URL: <http://7dsurgical.com/7d-surgical-system/spine-navigation/> (visited on 2017).
40. Philips announces new augmented-reality surgical navigation technology designed for image-guided spine, cranial and trauma surgery. Philips. URL: <http://www.philips.com/a-w/about/news/archive/standard/news/press/2017/20170112-philips-announces-new-augmented-reality-surgical-navigation-technology-designed-for->

- image-guided-spine-cranial-and-trauma-surgery.html (visited on 2017).
41. Elmi-Terander, A, Skulason, H, Söderman, M, Racadio, J, Homan, R, Babic, D, Vaart, Nvd, and Nachabe, R. Surgical Navigation Technology Based on Augmented Reality and Integrated 3D Intraoperative Imaging: A Spine Cadaveric Feasibility and Accuracy Study. *Spine* 2016;41:E1303–E1311.
  42. Bonsanto, MM, Metzner, R, Aschoff, A, Tronnier, VM, Kunze, S, and Wirtz, CR. 3D ultrasound navigation in syrinx surgery – a feasibility study. *Acta Neurochirurgica* 2005;147:533–540.
  43. Selbekk, T, Solheim, O, Bø, LE, and Hernes, TAN. Navigation and ultrasound imaging in surgery of pituitary tumours – technical considerations. *Minimally Invasive Therapy & Allied Technologies* 2010;19. Notes:53.
  44. Solheim, O, Selbekk, T, Løvstakken, L, Tangen, GA, Solberg, OV, Johansen, TF, Cappelen, J, and Unsgård, G. Intracellular ultrasound in transsphenoidal surgery: a novel technique. *Neurosurgery* 2010;66:173–185.
  45. FOR. St. Olavs Hospital. URL: <https://stolav.no/en/for> (visited on 2017).
  46. Wilson, E, Yaniv, Z, Zhang, H, Nafis, C, Shen, E, Shechter, G, Wiles, AD, Peters, T, Lindisch, D, and Cleary, K. A Hardware and Software Protocol for the Evaluation of Electromagnetic Tracker Accuracy in the Clinical Environment: a Multi-Center Study. In: *Medical Imaging 2007: Visualization and Image-Guided Procedures*. Ed. by Kevin, RC and Michael, IM. Vol. 6509. Proceedings of SPIE 1. SPIE, 2007:65092T.
  47. Lindseth, F, Tangen, GA, Langø, T, and Bang, J. Probe calibration for freehand 3-D ultrasound. *Ultrasound in Medicine & Biology* 2003;29:1607–1623.

48. Mercier, L, Langø, T, Lindseth, F, and Collins, DL. A review of calibration techniques for freehand 3-D ultrasound systems. *Ultrasound in Medicine & Biology* 2005;31:449–471.
49. Chen, TK, Thurston, AD, Ellis, RE, and Abolmaesumi, P. A Real-Time Freehand Ultrasound Calibration System with Automatic Accuracy Feedback and Control. *Ultrasound in Medicine & Biology* 2009;35:79–93.
50. Solberg, OV, Lindseth, F, Torp, H, Blake, RE, and Hernes, TAN. Freehand 3D ultrasound reconstruction algorithms – a review. *Ultrasound in Medicine & Biology* 2007;33:991–1009.
51. Mozaffari, MH and Lee, WS. Freehand 3-D Ultrasound Imaging: A Systematic Review. *Ultrasound in Medicine & Biology* 2017;43:2099–2124.
52. Amin, DV, Kanade, T, DiGioia, AM, and BranislavJaramaz. Ultrasound Registration of the Bone Surface for Surgical Navigation. *Computer aided surgery* 2003;8:1–16.
53. Rasoulian, A, Abolmaesumia, P, and Mousavi, P. Feature-based multibody rigid registration of CT and ultrasound images of lumbar spine. *Medical Physics* 2012;39:3154–3166.
54. Hacıhaliloglu, I, Rasoulian, A, Rohling, RN, and Abolmaesumi, P. Statistical Shape Model to 3D Ultrasound Registration for Spine Interventions Using Enhanced Local Phase Features. In: *Medical Image Computing and Computer-Assisted Intervention — MICCAI 2013*. Ed. by Mori, K, Sakuma, I, Sato, Y, Barillot, C, and Navab, N. Vol. 8151. Springer, 2013.
55. Baka, N, Leenstra, S, and Walsum, Tv. Ultrasound aided vertebral level localization for lumbar surgery. *IEEE Transactions on Biomedical Engineering* 2017.
56. Øygaard, TK. Improved Distance Weighted GPU-based 3D Ultrasound Reconstruction Methods. PhD thesis. 2014.
57. Orthopedic Navigation. Metronor AS. URL: <http://www.metronor.com/medical/orthopedic-navigation/> (visited on 2017).

58. Jain, AK and Taylor, RH. Understanding Bone responses in B-mode Ultrasound Images and Automatic Bone Surface extraction using a Bayesian Probabilistic Framework. In: *Medical Imaging 2004: Ultrasonic Imaging and Signal Processing*. Ed. by Walker, WF and Emelianov, SY. Vol. 5373. Proceedings of SPIE. SPIE, 2004:131–142.
59. Foroughi, P, Boctor, E, Swartz, MJ, Taylor, RH, and Fichtinger, G. Ultrasound Bone Segmentation Using Dynamic Programming. In: *2007 IEEE Ultrasonics Symposium*. Ed. by Yuhas, MP. IEEE, 2007:2523–2526.
60. Yan, CXB, Goulet, B, Tampieri, D, and Collins, DL. Ultrasound–CT registration of vertebrae without reconstruction. *International Journal of Computer Assisted Radiology and Surgery* 2012;7:901–909.
61. Rasoulia, A, Rohling, RN, and Abolmaesumi, P. Augmentation of Paramedian 3D Ultrasound Images of the Spine. In: *Information Processing in Computer-Assisted Interventions — IPCAI 2013*. Ed. by Barratt, D, Cotin, S, Fichtinger, G, Jannin, P, and Navab, N. Vol. 7915. Springer, 2013.
62. Rasoulia, A, Rohling, R, and Abolmaesumi, P. Group-wise registration of point sets for statistical shape models. *IEEE Transactions on Medical Imaging* 2012;31:2025–2034.
63. Rasoulia, A, Rohling, RN, and Abolmaesumi, P. A statistical multi-vertebrae shape+pose model for segmentation of CT images. In: *Medical Imaging 2013: Image-Guided Procedures, Robotic Interventions, and Modeling*. Ed. by Holmes III, DR and Yaniv, Z. Vol. 8671. SPIE, 2013.
64. Smistad, E and Løvstakken, L. Vessel Detection in Ultrasound Images Using Deep Convolutional Neural Networks. In: *Deep Learning and Data Labeling for Medical Applications*. Ed. by Carneiro, G, Mateus, D, Peter, L, Bradley, A, Tavares, JMRS, Belagiannis, V, Papa, JP, Nascimento, JC, Loog, M, Lu, Z, Car-

- doso, JS, and Cornebise, J. Vol. 10008. Lecture Notes in Computer Science. Springer, 2016:30–38.
65. Baka, N, Leenstra, S, and Walsum, Tv. Random Forest-Based Bone Segmentation in Ultrasound. *Ultrasound in Medicine & Biology* 2017;43:2426–2437.
  66. Salehi, M, Prevost, R, Moctezuma, JL, Navab, N, and Wein, W. Medical Image Computing and Computer-Assisted Intervention — MICCAI 2017. In: *Lecture Notes in Computer Science*. Ed. by Descoteaux, M, Maier-Hein, L, Franz, A, Jannin, P, Collins, DL, and Duchesne, S. Vol. 10434. Springer, 2017.
  67. Rodriguez-Molares, A, Løvstakken, L, Ekroll, IK, and Torp, H. Reconstruction of specular reflectors by iterative image source localization. *Journal of the Acoustical Society of America* 2015;138:1365–1378.
  68. Rodriguez-Molares, A, Fatemi, A, Løvstakken, L, and Torp, H. Specular Beamforming. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 2017;64:1285–1297.



Paper A

# 3D ultrasound reconstruction algorithms from analog and digital data

Ole Vegard Solberg<sup>1,2,3</sup>, Frank Lindseth<sup>1,2,3</sup>, Lars Eirik Bø<sup>1,2</sup>, Sébastien Muller<sup>1,2</sup>, Janne Beate Lervik Bakeng<sup>1,2</sup>, Geir Arne Tangen<sup>1,2</sup> and Toril A. N. Hernes<sup>1,2,3</sup>

<sup>1</sup>SINTEF Technology and Society, Trondheim, Norway

<sup>2</sup>National Centre for 3D Ultrasound in Surgery, Trondheim, Norway

<sup>3</sup>Norwegian University of Science and Technology, Trondheim, Norway

*Ultrasonics*, vol. 51, no. 4, pp. 405–419, 2011

**Abstract**

Freehand 3D ultrasound is increasingly being introduced in the clinic for diagnostics and image-assisted interventions. Various algorithms exist for combining 2D images of regular ultrasound probes to 3D volumes, being either voxel-, pixel- or function-based. Previously, the most commonly used input to 3D ultrasound reconstruction has been digitized analog video. However, recent scanners that offer access to digital image frames exist, either as processed or unprocessed data. To our knowledge, no comparison has been performed to determine which data source gives the best reconstruction quality. In the present study we compared both reconstruction algorithms and data sources using novel comparison methods for detecting potential differences in image quality of the reconstructed volumes. The ultrasound scanner used in this study was the Sonix RP from Ultrasonix Medical Corp (Richmond, Canada), a scanner that allow third party access to unprocessed and processed digital data. The ultrasound probe used was the L14-5/38 linear probe. The assessment is based on a number of image criteria: detectability of wire targets, spatial resolution, detectability of small barely visible structures, subjective tissue image quality, and volume geometry. In addition we have also performed the more “traditional” comparison of reconstructed volumes by removing a percentage of the input data. By using these evaluation methods and data from the specific scanner, the results showed that the processed video performed better than the digital scan-line data, digital video being better than analog video. Furthermore, the results showed that the choice of video source was more important than the choice of tested reconstruction algorithms.

**A.1 Introduction**

The use of 2D ultrasound for a variety of clinical applications is becoming more common. Compared to other imaging modalities like magnetic resonance imaging (MRI) and computed tomography (CT), ultrasound has the advantages of being cheaper, smaller and more flex-

ible, it has no radiation and it is easier to introduce during surgery. 3D freehand ultrasound offers even more flexibility, and combined with position tracking it is found useful in minimally invasive image-guided surgery (IGS). [1]

The most commonly used input to freehand 3D reconstructions is digitized analog video, either first stored on a video tape and digitalized with an analog video frame grabber [2–4] or obtained instantly by connecting the video frame grabber directly to the video output of the ultrasound scanner [4–18]. Some researchers state that they use digital data from the scanner. In most cases, such data is unavailable to third party users, but exceptions exist like the Sonix RP scanner (Ultrasonix Medical Corp., Richmond, Canada). Some groups also gain access to the digital data by collaborating with the ultrasound scanner manufacturers. [19–23]

Different algorithms for reconstructing 3D volumes from freehand ultrasound exist. In summary, these are [24]:

- **Voxel-Based Methods (VBM).** Include the Voxel Nearest Neighbor (VNN) where each voxel is assigned the nearest pixel [8] and algorithms where each voxel is assigned a value based on several of the nearest pixels [19, 21, 25–27]. In this group are also algorithms that skip the creation of a voxel volume and reconstruct a 2D slice or surface directly. [28, 29]
- **Pixel-Based Methods (PBM).** Range from algorithms like the one-step Pixel Nearest Neighbor (PNN) where 2D input images are inserted directly into a target volume [25] to two-step PNN where a second step fills empty voxels afterwards [2, 3, 30–32] to algorithms where input pixels are added with a 3D kernel [6, 17, 33, 34].
- **Function-Based Methods (FBM).** Algorithms where functions are made based on the input pixels and the target volume is created by evaluating these functions at regular intervals. [35, 36]

For 3D probes [37, 38], the process is digital from acquisition to reconstruction. The DICOM 2008 standard (DICOM 2008, Suppl. 43) also has defined structures for 3D ultrasound volumes. 3D probes still have the disadvantage of poorer resolution compared to 2D probes. 3D probes on the other hand have the advantage of allowing real-time 3D volumes, while 2D probes may only provide real-time 2D images. In IGS, the positions of the data are also necessary in order to navigate in the 3D volume, so a position sensor must be attached to the ultrasound probe and the relation between the sensors and image data must be defined (probe calibration). In the present study we have evaluated the differences between image volumes originating from different ultrasound data sources and various 3D reconstruction algorithms using a broad range of both quantitative and qualitative comparisons methods.

## A.2 Materials and methods

### Ultrasound data import

The ultrasound scanner used (Figure A.1, Sonix RP, Ultrasonix Medical Corp., Richmond, Canada) has a research interface allowing real-time access to digital data from the scanner. A 7.5 MHz linear probe (Figure A.1, L14-5/38, Prosonic, Gyeongbuk, South Korea) [39] operating at 10 MHz scanning frequency was used to acquire all images analyzed in this study.

Three different video streams were imported simultaneously:

- Analog video, converted from PAL S-Video with a video-to-FireWire converter (DFG/1394-1e, The Imaging Source, Germany).
- Digital scan converted video, processed by the ultrasound scanner for viewing on a screen.
- Unprocessed digital data delivered as 1D scan lines, only envelope detected and log-compressed.

The digital image sources were imported directly from the ultrasound scanner over a crossed LAN (Local Area Network) cable. The three video streams were imported simultaneously in different threads on a PC with four CPU kernels (Figure A.1, Intel Core 2 Quad Processor Q6700 2.66 GHz). Each time an image was received on the computer, a time stamp was created and assigned to the image. A depth setting of 4 cm was used on the scanner for all image acquisitions. This resulted in a pixel size (width  $\times$  depth) of 0.147 mm  $\times$  0.147 mm for the analog video, 0.097 mm  $\times$  0.097 mm for the digital video, and 0.150 mm  $\times$  0.077 mm for the unprocessed “video”. The images from the analog and digital video were cropped to only contain the ultrasound data. The difference in pixel sizes between the unprocessed video and digital video is due to the ultrasound scanners internal processing of the video. The pixel size difference between the digital and analog video comes from lesser resolution in the video grabbing hardware. All data sources supplied 8-bit pixels.

### **Position tracking and probe calibration**

For freehand 3D ultrasound reconstruction, positions and orientations of the 2D images are needed. Several methods for obtaining these positions exist. [24, 40] In our study we used an optical positioning system (Figure A.1, Polaris Spectra, Northern Digital Inc., Canada), consisting of a tracking frame attached to the ultrasound probe and a camera unit that were used for calculating the position and orientation of this frame. The positions were obtained using the Image-Guided Surgery Toolkit (IGSTK) [41–43]. As with the video streams, time stamps for the positions were created and assigned by the software at the time the positions were received. The positions were imported with the same application importing the three video streams, but run as a separate thread.

For the 2D ultrasound image to be correctly aligned with the output from the positioning system, a calibration is necessary. We used the spatial calibration method and phantom developed by Chen et al. [18], with threads stretched between the sidewalls in two parallel, ‘N’-

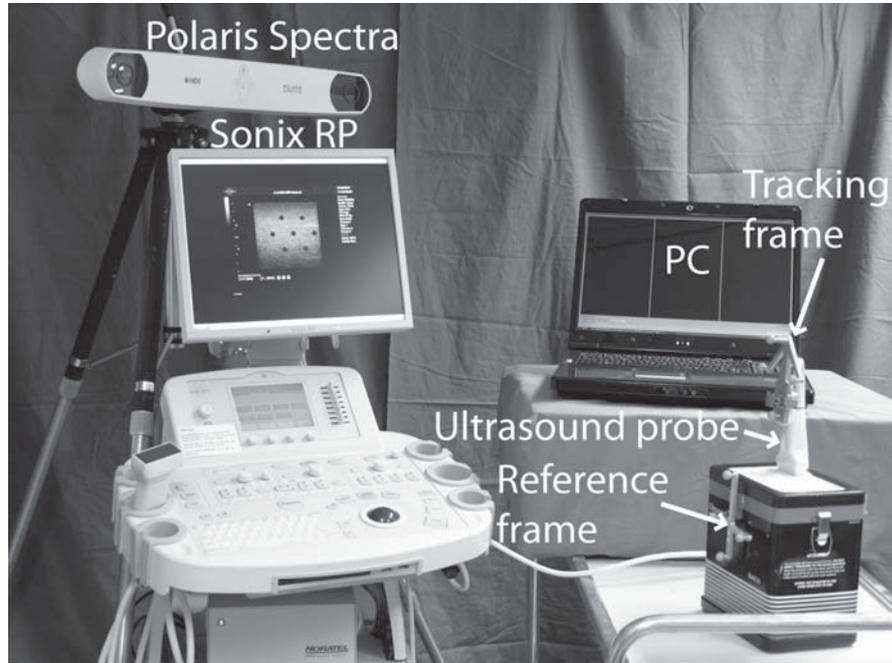


Figure A.1: System setup. Polaris Spectra optical positioning system, Sonix RP ultrasound scanner, ultrasound probe with tracking frame, ultrasound phantom with tracking reference frame, and PC for data import.

shaped configurations. A tracking frame was mounted on the phantom and the positions of the threads relative to this frame were measured. The corresponding structures were identified in the ultrasound images using an automatic segmentation algorithm. The relationship between the image plane and the positioning system was then found using a least-squares minimization method. We also implemented a temporal calibration based on the work by Treece et al. [44] except that we detected a point instead of a line. This temporal calibration was used to synchronize the imported positions with the images. The temporal calibration method matched the vertical movement of one of the segmented points in the 2D images to the vertical movement of the

ultrasound probe as reported by the positioning system. Using the temporal calibration, 23.86 ms was subtracted from the analog video time-tags to match the time-tags of the positions, 2.56 ms was added to the digital video time-tags and 0.32 ms to the time-tags of the unprocessed images. The total mean error from the spatial calibration was 1.05 mm, with a root mean square (RMS) value of 1.13 mm and a standard deviation of 0.42 mm. These results were used in the reconstructions and display of data in the following tests.

### **Reconstruction algorithms**

Two pixel-based 3D reconstruction algorithms were compared in this study. The first is called Pixel Nearest Neighbor [2, 3, 24, 30] and is a relatively fast two-step method. The first step inserts each image pixel in the input 2D images into the target 3D volume based on the position and orientation of the images. The chosen implementation overwrites any existing data in the 3D volume with the most current 2D image. The second step is an interpolation step that traverses the voxels of the target volume and attempt to fill empty voxels with the average value of the nearby voxels. The interpolation first tries to interpolate a voxel with the voxel values from the  $3 \times 3 \times 3$  grid around it. If all these voxels are empty a  $5 \times 5 \times 5$  grid is used and after that a  $7 \times 7 \times 7$  grid, and if there is still no voxel values within this range the voxel is left empty. The second reconstruction algorithm uses a 3D kernel around the input pixels. Several variations of the input kernel are described [6, 17, 24, 33, 34, 45], and we used an ellipsoid truncated Gaussian-weighted kernel around the input pixels [17, 24]. The size of the kernel is usually set to fill holes in the volume, but we used a novel method of approximately matching the theoretical ultrasound resolution in all three dimensions. For the comparisons we used two slightly different sizes of this kernel, resulting in three different reconstructed volumes:

- Pixel Nearest Neighbor.
- Small 3D ellipsoid kernel around input pixels.

- Large 3D ellipsoid kernel around input pixels.

The reconstruction algorithms were set to create volumes that used the full range of 8 bits to produce volumes that were similar in intensity. The voxel size of all reconstructed volumes were set to 0.2 mm as a compromise between resolution and processing time, the voxel sizes being larger than the input pixel sizes.

*Determining kernels for the 3D kernel based reconstruction algorithm*

As an approximation of the two-way pulse-echo response of the ultrasound imaging system we have decided to use a 3D Gaussian function. This function is best known as the probability density function for a normal distribution, and in 1D it is given by the formula

$$f(x | \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}, \quad (\text{A.1})$$

where  $\mu$  is the mean and  $\sigma$  is the standard deviation. In 3D this becomes

$$f(\mathbf{x} | \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{(\sqrt{2\pi})^3 \sqrt{|\boldsymbol{\Sigma}|}} e^{-\frac{(\mathbf{x}-\boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x}-\boldsymbol{\mu})}{2}}, \quad (\text{A.2})$$

where  $\boldsymbol{\mu} = (\mu_x, \mu_y, \mu_z)$  is the mean vector, and  $\boldsymbol{\Sigma}$  is the covariance matrix.

The actual pulse-echo response may be approximated as a product of two sinc functions. The formula for a rectangular aperture in the lateral and elevation direction (Figure A.2) may be derived from Angelsen [46, p. 5.54]:

$$H(x, \lambda) = \text{sinc}\left(\frac{x}{\lambda_t f_{\#t}}\right) \text{sinc}\left(\frac{x}{\lambda_r f_{\#r}}\right), \quad (\text{A.3})$$

where  $H$  is the two-way pulse-echo response,  $x$  is distance in either the lateral or elevation direction (Figure A.2).  $\lambda = c/f$  is the wavelength,  $f_{\#} = F/D$  is the  $f$ -number,  $c$  is the speed of sound,  $f$  is the ultrasound center frequency,  $F$  is the focal position (depth for the calculation)

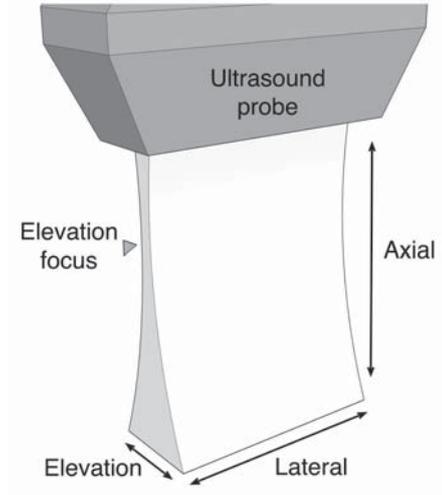


Figure A.2: Illustration of the imaging sector of an ultrasound probe. The lateral, axial and elevation direction are marked in addition to the fixed elevation focus.

and  $D$  is the effective aperture size. The transmit ( $t$ ) and receive ( $r$ ) directions may have different  $\lambda$  and  $f_{\#}$ . Acoustic absorption reduces the pulse center frequency, resulting in a lower center frequency for the received pulse compared to the transmitted pulse. For a Gaussian pulse envelope the frequency is reduced by the following formula [47]:

$$f_r(z) = f_t - \frac{\alpha B^2 z}{4 \ln 2}, \quad (\text{A.4})$$

where  $z$  is the propagation distance,  $\alpha = (a \ln 10)/20$  is the average constant of the absorption coefficient. A commonly used value for  $a$  is 0.5 dB/(cm MHz).  $B$  is the  $-6$  dB bandwidth of the imaging pulse. In the axial direction (see Figure A.2), the pulse-echo response is dependent of the form of the transmitted pulse, but is often approximated as a Gaussian function.

To calculate the theoretical size of the focus of an ultrasound probe in the lateral or elevation direction the function (A.3) can be evaluated

at  $-6$  dB. However, for the Gaussian function we use the width of the main lobe, equal the first zero in the narrower of the two sinc functions in (A.3). The resulting value was used as the limit for a 99.7% confidence level for a 1D Gaussian distribution (total pulse width obtained from (A.4) corresponds to 6 standard deviations, so the standard deviation ( $\sigma$ ) is estimated simply by dividing the width of the main lobe by 6). The ultrasound scanner uses a Gaussian-like apodization on the elements to dampen the side lobes of the transmitted pulse. This apodization will also create a wider main lobe than indicated by the formula (A.3), which are for a situation without apodization [46, pp. 5.56 – 5.58].

The ultrasound probe had a fixed elevation focus of 16 mm (Ultrasonix Medical Corp., Richmond, Canada) and the ultrasound scanner allows the operator to set several variable focus points in the axial direction. Three focus points were used and at least one was set near the fixed elevation focus. The probe size in the lateral direction is 38 mm, but the effective aperture size varies according to depth (dynamic aperture). The aperture sizes for both transmit and receive were calculated with parameters read from the scanner and code obtained from Ultrasonix. According to information obtained from Ultrasonix the fractional Bandwidth of the probe is minimum 70% of the center frequency at  $-6$  dB, the center frequency being 7.2 MHz, resulting in a bandwidth of 5.04 MHz. The theoretical lateral and elevation resolutions in focus for the used ultrasound probe are illustrated in Figure A.3A by evaluating formula (A.3) at  $-6$  dB, while the calculations of the total width of the main lobe are shown in Figure A.3B. The imaging frequency used is 10 MHz, and the speed of sound in tissue is set to 1540 m/s.

The size in the elevation direction (element height) is 4 mm (Ultrasonix Medical Corp., Richmond, Canada). The theoretical resolution in the axial direction equals half the pulse length [46, p. 1.22], i.e.

$$\Delta a = \frac{cT_p}{2} = \frac{c}{2B}, \quad (\text{A.5})$$

where  $T_p$  is the pulse length in time. With a bandwidth of 5.04 MHz

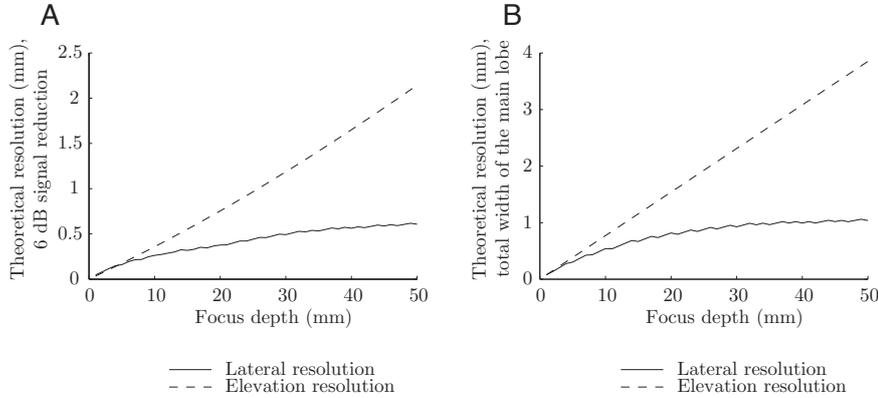


Figure A.3: Plot of the theoretical lateral and elevation resolution in focus for a few selected depths for the L14-5/38 probe operating at 10 MHz. (A) The resolutions shown as a 6 dB signal reduction, calculated from (A.3) and (A.4). The corresponding axial resolution is 0.154 mm for all depths, calculated from (A.5). (B) The resolutions illustrated as the total width of the main lobe from (A.3).

the best theoretical axial resolution is 0.153 mm (A.5). However, as we wanted to find a value to use for the limit of the 99.7% confidence level of a 1D Gaussian function, the full pulse length should be used: 0.306 mm.

The resolution values in all three dimensions were used to create the covariance matrix  $\Sigma$  for the 3D Gaussian function (A.2) and thus define the extent of the truncated kernel. For the small kernel we used an imaging depth of 19 mm to calculate the lateral resolution (= 0.78 mm), and for the large kernel we used a depth of 32 mm (= 0.99 mm). To speed up reconstruction computation time we reduced the total kernel size by calculating the elevation resolutions at somewhat shallower depths. The elevation resolutions were calculated for the depth of 18 mm (= 1.39 mm) for the small kernel and for the depth of 24 mm (= 1.85 mm) for the large kernel.

In the 3D reconstruction, the 3D Gaussian function was oriented according to the 2D ultrasound images, and discrete values of the

function were used to match the target 3D voxel grid. For the axial resolution we used a resolution of 0.171 mm instead of 0.306 mm, due to initially calculating the axial resolution for 9 MHz instead of 5 MHz. We truncated the 3D Gaussian kernel at 95% confidence level and then increased the kernel size to include the same number of voxels on both sides of the input pixel. Since the axial resolution was so high this resulted in a kernel size of only one voxel in the axial direction. However, the current implementation of the reconstruction requires a larger kernel in this direction, so we increased the kernel size in the axial direction to three voxels.

### **Data collection**

All acquisitions were performed as freehand translation sweeps with the three different data streams imported simultaneously. To minimize errors from the position-to-image synchronization, and to make sure that the reconstructions got enough data to fill holes, all sweeps were performed with a slow, smooth motion. One scan thus resulted in input data from three different data sources, all showing images from the same structures. The digital video after scan conversion was the same images as presented on the screen of the ultrasound scanner. These images were processed by the scanner for viewing and all user-controlled functions, such as the gain function on the scanner's control panel were applied by the scanner. The analog data is a video signal output from the ultrasound scanner. The signal was converted from digital to analog by the scanner, and again to a digital signal by the video grabber.

Throughout this paper the images used for the 3D reconstructions are called "input images" while images used directly in the comparisons are called "original images". In the comparisons the same images were never used as both input and original images even if the images may be similar and acquired in approximately the same positions. Both input images and original images came from all three data sources, and while input images were acquired as a stack of several images original images were only acquired as single images.

### Tests for comparing reconstruction algorithms and data sources

1. *Compare the reconstruction algorithms' ability to correctly recreate removed data (test 1):* To compare the quality difference of different reconstructions based on the same input data the method of removing a percentage of input data [31, 32, 35, 48, 49] was used. We scanned a section of the underside of the forearm on two healthy volunteers with two translation scans along the arm of one volunteer and two translation scans across the arm of the other volunteer. The scans along the forearm gave images that only changed slightly from one image to the next, while the images from the scans across the forearm changed more rapidly. Before reconstruction, 0 %, 25 %, 50 %, 75 %, 100 %, 300 %, 500 % and 700 % of the data of one of the input 2D ultrasound images was randomly removed. Random pixels were removed from one slice for the percentages below 100 %, and whole slices were removed for the percentages 100 % to 700 %, e.g.: for 300 % three slices were removed. After reconstruction, all pixel values of this input image were compared with the voxel values from the corresponding positions in the reconstructed volumes and the RMS value of the differences was calculated. The orientation of the reconstructed volume was based on the orientation of the input image from which the data was removed. This procedure was performed on four different positions without holes in each volume, and statistical analyses were used to compare the performance of the reconstruction algorithms.
2. *Compare the reconstructed volumes' ability to display existing structures (test 2):* Measurements of the reconstructed volumes' ability to retain the resolution in the 2D images and of the resolution they manage to obtain in the elevation direction were performed. A comparison was performed by human observers on how well tissue was visualized. The human observers were technical researchers with knowledge of ultrasound imaging ranging

from medium to expert. The reconstructed volumes from one scan resulted in a set of volumes with various combinations of the three data sources and three reconstruction algorithms. Original 2D ultrasound images from the ultrasound scanner were acquired simultaneously as a selection of 2D anyplane slices from the reconstructed volumes in the same position and orientation (Figure A.4C and D). Statistical comparisons were performed between reconstruction algorithms/original 2D ultrasound images and video sources. All volunteers were presented with the same set of images on the same computer with brightness and contrast levels unchanged.

- a) *Visual comparisons of structures placed increasingly closer (test 2a)*: Measuring the volume resolutions was accomplished by scanning an ultrasound resolution phantom, with six small structures (threads) that are placed with distances of 1 mm, 2 mm, 3 mm, 4 mm and 5 mm (Model 040, CIRS Inc., VA, USA). The threads of the CIRS 040 phantom are 0.1 mm in diameter and made of Nylon Monofilament. The threads were scanned with three translation sweeps both along (Figure A.4A) and across (Figure A.4B) the threads to measure resolutions both in the lateral and elevation direction (Figure A.2) of the ultrasound input planes. 2D images were obtained by collecting original ultrasound images with position and orientation (Figure A.4C). This position and orientation were used to create 2D anyplanes through the reconstructed 3D volumes (Figure A.4D). These images were presented in random order and evaluated by eight volunteers that were given two questions to answer for each image: “Count how many separate bright structures you can see and rate how easy it is to identify those you can see (1–5, where 1 is easy and 5 hard)”.
- b) *Comparisons based on image measurements of spatial resolution (test 2b)*: The images in test 2a of the CIRS 040 phantom also contained a separate thread, which was used

for measurements of resolution: Both the original and a re-sampled (downsampled) version of the original image were used, and resolution were measured in positions corresponding with the anyplanes through the 3D volumes. Axial and lateral spatial resolutions were measured directly from the 2D ultrasound images by plotting gray level profiles through the center of the scanned wires [50]. The maximum pixel value was used as the center value for each thread. Parabolic curves were matched to these plots and evaluated at  $-6$  dB and  $-20$  dB. Only the values above  $-6$  dB were used for matching the parabolic curves both in the original 2D images and in 2D anyplanes obtained from the 3D volumes. Measurements in the elevation direction were possible to perform on anyplanes through the 3D volumes scanned across the threads (Figure A.4B). Each measurement was performed on the same wire at a depth of approximately 27 mm in three different scans.

- c) *Visual comparisons of small, barely visible structures (test 2c)*: The CIRS 044 (Model 044, CIRS Inc., VA, USA), an ultrasound phantom with small cylinders of varying size, was used to compare the visibility of small structures after a reconstruction. Three ultrasound translation sweeps were performed on the smallest cylinders on the CIRS 044 phantom, both lengthwise and crosswise. These small cylinders have measurements of 1.5 mm (diameter) and 2.4 mm (length) specified in the fabrication sheet. Original single 2D images were acquired showing as many cylinders as possible, all cylinders appearing as circles in the ultrasound image. Anyplane images (Figure A.4D) through the 3D volumes were created from the same position and orientation as the collected original images (Figure A.4C). Eight people were presented 2D images, in random order, showing small objects and presented with two questions for each image: “Count how many separate dark structures you can see,

and rate how easy it is to identify those you can see (1–5, where 1 is easy and 5 hard).” When presented with several images of varying quality, the test person may “learn” where the structures should be, and this may enable them to identify more structures in a poor image that they would do otherwise. To allow for this the test subjects were also asked to rate (1–5) how easy it was to identify the structures they could identify, and they were also shown an illustration of the corresponding section of phantom beforehand. The 2D images were either an original 2D ultrasound image showing the structures or a 2D anyplane obtained from a 3D volume. The anyplanes from the 3D volumes were either approximately orthogonal or parallel to the 2D images from which the 3D volume was created.

- d) *Visual comparisons of tissue data orthogonal to the scanning direction (test 2d)*: A section of the underside of the forearm of two healthy volunteers was scanned with free-hand translation sweeps. One person was scanned with two sweeps along the arm while the other was scanned with two sweeps across the arm. Original single 2D images were acquired approximately orthogonal to the 3D acquisition sweeps for comparison (Figure A.4C) and corresponding anyplane images through the 3D volumes were collected (Figure A.4D). A group of eight people were presented with different 2D images in random order, showing tissue data from the same location. They were presented with sets of 3 or 4 images and asked the following question: “Sort the images according to quality and give each image a quality score (1–5), where 1 is best and 5 is worst.” The images were original 2D ultrasound images of the tissue or anyplanes obtained from reconstructed 3D volumes approximately orthogonal to the input 2D images. The sets of 3 images showed the 3 different input sources, being all original 2D images or all anyplane images from a specific reconstruc-

tion algorithm. The sets of 4 images showed an original 2D ultrasound image and anyplane images from different reconstruction algorithms, all images being from the same input source.

3. *Compare the correctness of the reconstructed volume's geometry (test 3):* A geometry comparison is a measurement on how well a reconstruction manages to recreate a known phantom geometry. Statistical comparisons based on both data sources and reconstruction algorithms were performed.

The CIRS 044 have a set of cylinders with specified measurements of 3 mm (diameter)  $\times$  6 mm (length). Three ultrasound translation sweeps were performed on these structures in both the lengthwise and crosswise direction. Original single 2D ultrasound images with 3D positions and orientations were acquired for each data set both along and across the cylinder. Both height and width of the structure in the image was measured, resulting in measurements of cylinder length in the lateral direction and diameter height in the axial direction for the images along the cylinder, and measurements of diameter width in the lateral direction and diameter height in the axial direction. Three measurements were performed in the original image and three in an image processed with a levels function of an image processing application. The levels function was used to spread out the pixel intensities so that they fill the whole 8 bit range. The reason for this was to try to emulate the processing in the image reconstructions where a similar function was used for the whole volume.

2D anyplane images were created through the reconstructed 3D volumes at positions and orientations matching that of the original 2D images. The same structure measurements were done on these anyplane images and the differences were compared and tested for variations in reconstruction algorithm and data source quality. The anyplane images covered measurements in the ele-

vation direction in addition to measurements in the lateral and axial direction. The elevation measurements were received from cylinder length in the scans along the cylinder and from the cylinder width measurements in the scans across. The same person performed all analyses, repeating each measurement three times. The cylinder measurements performed on the original ultrasound images were used as the gold standard for the comparisons.

### **Statistical comparisons**

All data were compared using statistical methods using the SPSS Statistics software (SPSS 16 for Mac, SPSS Inc., IL, USA). All statistical tests were performed with a 5% confidence level ( $\alpha = 5\%$ ). To check if all groups came from the same distribution, an overall test was applied. In case of rejection, each group was tested against the others.

Each group was tested for normal distribution using the Shapiro-Wilk (SW) test to check for the possibility to use parametric statistics. For unrelated samples, if normality was accepted for all groups the Analysis of Variance (ANOVA) test was used for the overall statistics and the Bonferroni multiple-comparisons procedure was used for the pairwise comparisons. If normality was rejected the non-parametric Kruskal-Wallis (KW) test was used for the overall statistics and the Mann-Whitney-Wilcoxon (MWW) test was used for the pairwise comparisons. For test 1 with related samples the Friedman test was used for the overall statistic and the Wilcoxon Signed Rank (WSR) test was used for the pairwise comparisons. All statements in the text regarding differences (e.g.: performed better/poorer than, better/worse result, harder to identify, more accurate detection) are based on statistically significant results, even if this not mentioned explicitly every time.

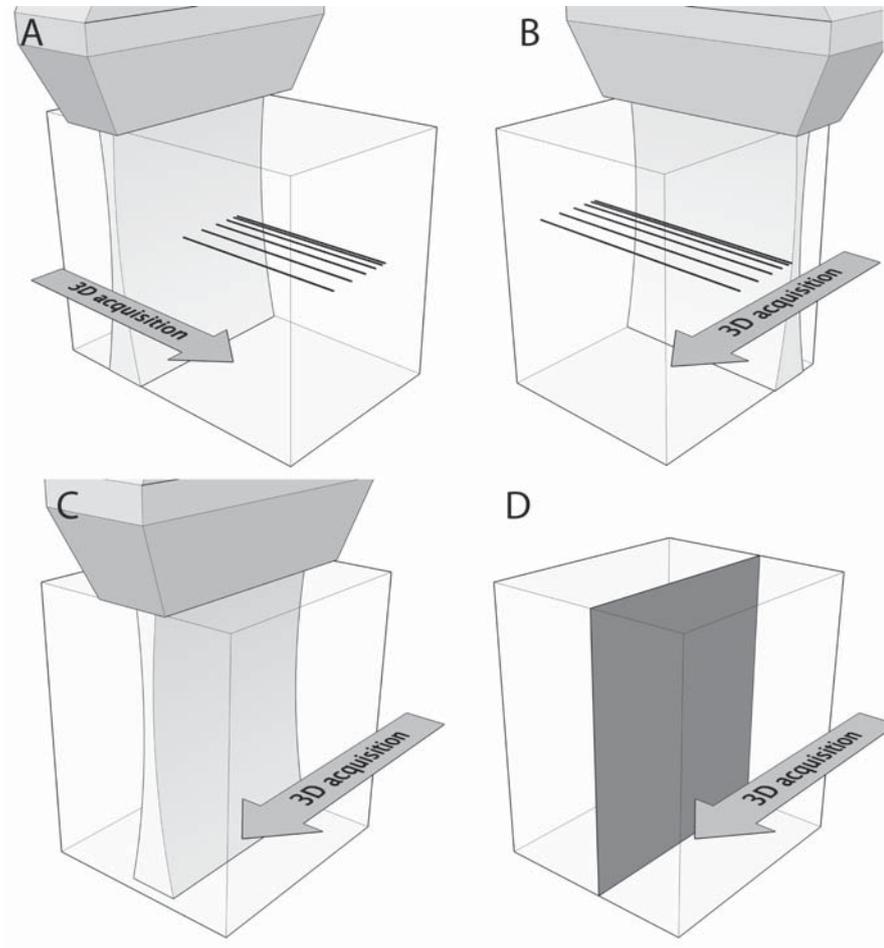


Figure A.4: Illustrations of ultrasound sectors. The 3D acquisition scan direction is indicated, and the reconstruction volume is shown as a transparent box. (A) A 3D scan along the threads of a resolution phantom. The threads are seen as points in the ultrasound sector. (B) A 3D scan across the threads of a resolution phantom. The threads are seen as lines in the ultrasound sector. (C) Illustration of a single original ultrasound image acquired orthogonal to the input images in the 3D reconstruction. (D) Illustration of a 2D anyplane through the 3D volume in the same position as the original image in (C).

### A.3 Results

#### Compare the reconstruction algorithms abilities to correctly recreate removed data (test 1)

The differences between the slice with the removed data and the data values from the anyplanes in the same positions in the reconstructed volumes were plotted as curves with RMS results for each removed percentage. See Figure A.5 as one example showing the RMS values from the different reconstructed volumes using the digital video as input in the scans across the forearm. All RMS values for the different removed percentages were combined for the statistical comparisons (Table A.1). Table A.1 shows results from the scans taken either across or along the forearm.

Statistically significant differences were found between the following reconstruction algorithms: In the scans across the forearm, representing data with high degree of variation, the large kernel performed poorer than the other two reconstruction algorithms for the digital data, and for the unprocessed video the small kernel performed better than the large. Also when combining the data from all three video sources the small kernel performed better than the large. In the scans along the forearm, representing data with little variation, the PNN reconstruction gave a better result than the others for the analog and digital video while it gave a worse result for the unprocessed data. For the unprocessed data in the scan along the forearm the large kernel reconstruction performed better than the small kernel that performed better than the PNN reconstruction. The differences between the scan across and along the forearm were visible as higher mean values in the scans across the arm (Table A.1). When comparing values in Table A.1 it was only possible to compare values from the same data source because of the different characteristics of the data sources.

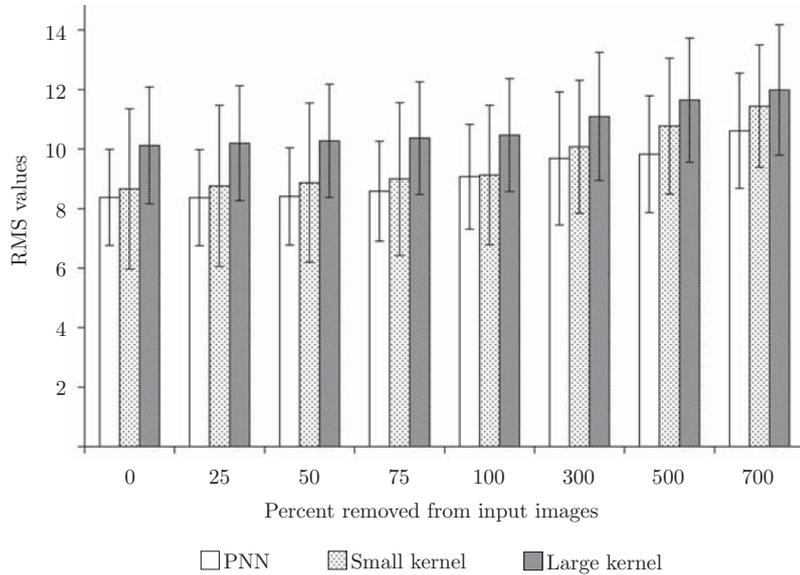


Figure A.5: RMS values for the three different reconstruction algorithms for the digital video acquired across the forearm. Mean and standard deviation values are presented for each removed percentage. All RMS measurements for the large kernel are larger than the other two, and the difference is statistically significant.

### Visual comparisons of structures placed increasingly closer (test 2a)

The data from the comparisons can be illustrated as graphs showing the number of threads and the difficulty to identify the threads (Figure A.6, Table A.2). The compared images came from scans both along (Figure A.4A) and across (Figure A.4B) the threads (See also Figure A.7A and B for examples of the images), the scans along the threads representing a comparison in the lateral direction of the ultrasound images and the scans across the threads representing a comparison in the elevation direction as related to the ultrasound image (Figure A.2).

Table A.1: RMS values obtained by removing a percentage of data (0% to 700% removed) from selected input images. The scans across the forearm represent data with large variations while the scans along the forearm represent data with little variation from one image to the next. The values in the table between data sources are not comparable.

|   | Analog    | Digital    | Unprocessed |
|---|-----------|------------|-------------|
| <i>Scan across arm (large variation between images)</i> |           |            |             |
| PNN   | 8.38±2.42 | 9.12±2.45  | 16.93±3.43  |
| Small kernel  | 7.57±1.46 | 9.59±1.98  | 16.10±2.96  |
| Large kernel  | 7.50±1.37 | 10.77±2.01 | 17.12±3.19  |
| <i>Scan along arm (small variation between images)</i>  |           |            |             |
| PNN   | 5.04±1.01 | 6.83±1.31  | 14.07±0.99  |
| Small kernel  | 5.68±0.96 | 7.41±1.52  | 13.21±1.40  |
| Large kernel  | 5.83±1.04 | 7.49±1.52  | 12.86±1.17  |

The results showed no statistically significant differences between the reconstruction algorithms, each with a median of 5 identified structures (Table A.2). The identification of only 5 structures means that the structures with the distance of 1 mm could not be separated. Comparing the results from the anyplanes through the reconstructed volumes with the original 2D ultrasound images, the original images gave statistically significantly better results with a median of 6 identified structures (Table A.2), meaning that the structures with the smallest distance of 1 mm could be separated in the majority of the observations. The comparison of the data sources showed that for all tests except the structure count from the scan along the threads, the unprocessed data produced volumes where fewer structures could be identified and they were harder to identify.

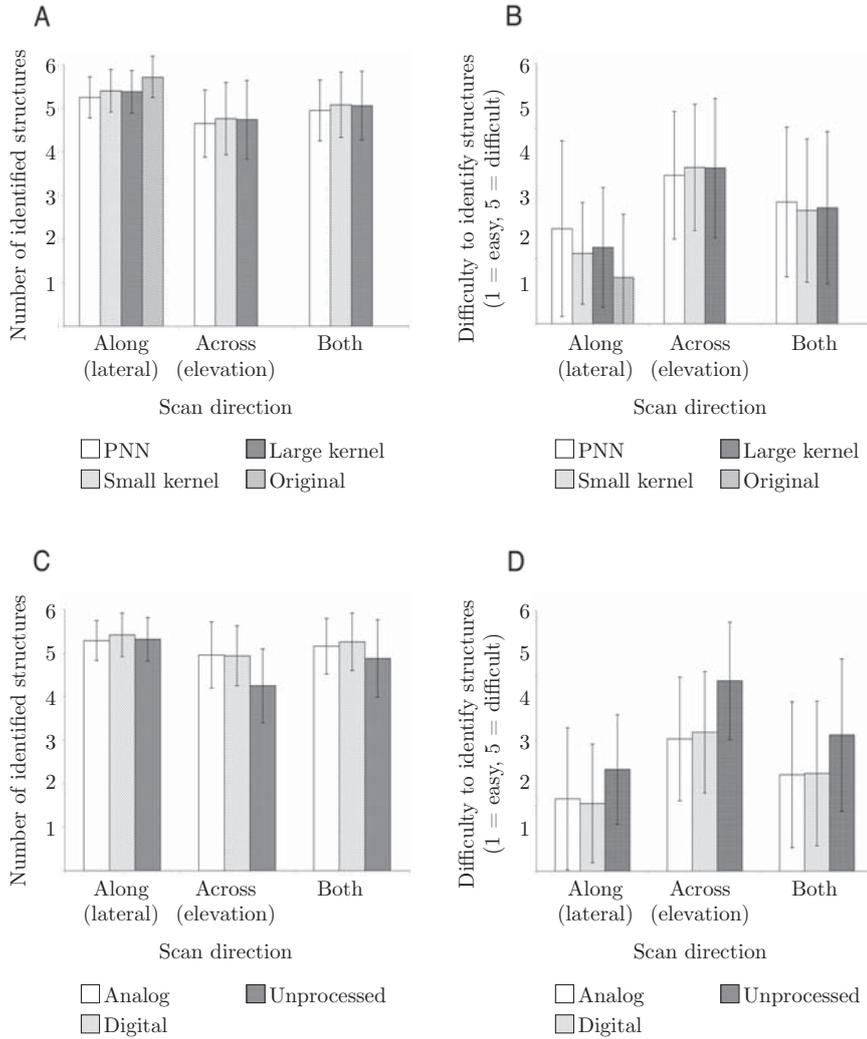


Figure A.6: Identification of 6 structures placed increasingly closer. Results from the scans along, across and both combined. Mean and standard deviation values are shown in the graphs. The six threads were identified by eight volunteers. (A) Number of identified threads for the reconstruction algorithms and original images. (B) The difficulty to identify the threads for the reconstruction algorithms and original images. (C) Number of identified threads for the different data sources. (D) The difficulty to identify the threads for the different data sources.

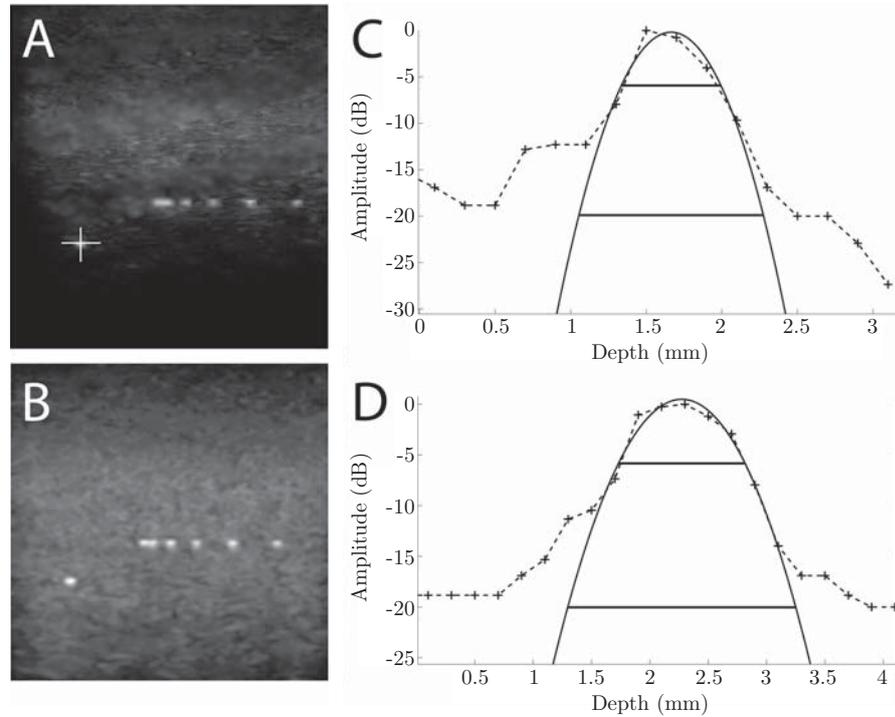


Figure A.7: Example of thread size as a measurement of spatial resolution in the axial and lateral direction. The images are based on an anyplane image from a PNN reconstructed volume with analog video as input. (A) The anyplane slice through the ultrasound volume with an illustration of where (crosshairs) the intensity plots in the axial and lateral direction were collected. The horizontal measurement give the elevation resolution when the anyplanes come from the volumes scanned across the threads. The other structures in this image were used in test 2a. (B) The original ultrasound image from the analog video. Pixel values have been modified with the levels function of an image processing application to provide an image with more contrast in the paper. (C) A dB plot of the intensity values in the axial direction from (A) with a fitted parabolic curve. (D) A dB plot of the intensity values in the lateral direction from (A) with a fitted parabolic curve.

Table A.2: Number of identified small structures. Mean and standard deviation values for the different combinations of reconstruction algorithm and data source.

|   | Analog    | Digital   | Unprocessed |
|---|-----------|-----------|-------------|
| <i>Scan along structures (lateral direction)</i>    |           |           |             |
| PNN   | 5.17±0.38 | 5.42±0.50 | 5.17±0.48   |
| Small kernel  | 5.33±0.48 | 5.46±0.51 | 5.42±0.50   |
| Large kernel  | 5.38±0.49 | 5.38±0.49 | 5.37±0.49   |
| Original  | 5.50±0.52 | 5.94±0.25 | 5.69±0.48   |
| <i>Scan across structures (elevation direction)</i> |           |           |             |
| PNN   | 4.83±0.82 | 4.88±0.74 | 4.25±0.61   |
| Small kernel  | 5.00±0.72 | 4.92±0.65 | 4.38±0.97   |
| Large kernel  | 5.04±0.76 | 5.04±0.69 | 4.12±0.95   |
| <i>Both scan directions combined</i>                |           |           |             |
| PNN   | 5.00±0.65 | 5.15±0.68 | 4.71±0.71   |
| Small kernel  | 5.17±0.63 | 5.19±0.64 | 4.90±0.93   |
| Large kernel  | 5.21±0.65 | 5.21±0.62 | 4.75±0.98   |

### Comparisons based on image measurements of spatial resolution (test 2b)

Parabolic curves were matched to plots through the centers of the scanned wires (Figure A.7) and these curves were evaluated at 6 dB and 20 dB levels below the maximum value. The 6 dB values are presented in Table A.3.

There were no statistically significant differences between the thread measurements from the different reconstruction algorithms, but they performed better than the original images in some of the tests: In the lateral direction all reconstruction algorithms gave better results than both the original images and the resampled version of these images for both 6 dB and 20 dB. For 6 dB, the measurements in the lateral direction based on the anyplanes through the reconstructed volumes

Table A.3: Image measurements of spatial resolutions. The results are presented as mean and standard deviation values in mm, and the measurements are from a fitted parabolic curve evaluated as 6 dB (see Figure A.7 for example).

|                    | Axial     | Lateral   | Elevation |
|--------------------|-----------|-----------|-----------|
| <i>Analog</i>      |           |           |           |
| PNN                | 0.65±0.09 | 1.15±0.04 | 1.86±0.38 |
| Small kernel       | 0.69±0.05 | 1.18±0.05 | 2.10±0.31 |
| Large kernel       | 0.69±0.06 | 1.16±0.07 | 2.10±0.32 |
| Original           | 0.74±0.01 | 1.28±0.01 | –         |
| Resampled          | 0.76±0.01 | 1.27±0.02 | –         |
| <i>Digital</i>     |           |           |           |
| PNN                | 0.66±0.15 | 1.08±0.15 | 1.58±0.50 |
| Small kernel       | 0.62±0.07 | 1.11±0.06 | 2.31±0.62 |
| Large kernel       | 0.65±0.08 | 1.16±0.08 | 2.24±0.50 |
| Original           | 0.69±0.03 | 1.28±0.03 | –         |
| Resampled          | 0.73±0.02 | 1.25±0.01 | –         |
| <i>Unprocessed</i> |           |           |           |
| PNN                | 0.76±0.20 | 1.22±0.19 | 1.90±0.43 |
| Small kernel       | 0.71±0.05 | 1.23±0.08 | 2.37±0.51 |
| Large kernel       | 0.74±0.07 | 1.28±0.06 | 2.34±0.47 |
| Original           | 0.98±0.03 | 1.80±0.03 | –         |
| Resampled          | 0.78±0.02 | 1.45±0.02 | –         |

had a mean thread width of 1.17 mm while the original images had a mean width of 1.32 mm. In the axial direction for the 6 dB evaluations, only the small kernel was better than the original images with a mean thread height of 0.67 mm compared to the mean height of the original images of 0.80 mm. Both the small and the large kernel performed better than the resampled original images. When fewer groups were compared compared without regard to the other groups more differences were detected also for the axial resolution: For 6 dB the PNN algorithm performed better than both the resampled and original images, and also the large kernel performed better than the original images. For 20 dB both the PNN and small kernel reconstruction performed better than the original images. For the 20 dB evaluations in the axial direction, all reconstruction algorithms gave better results than the resampled original images while none showed statistically significant difference from the original images. The results from the data sources showed that the unprocessed ultrasound performed poorer than both analog and digital video in both axial and lateral direction for 6 dB and 20 dB with a measured mean axial thread height for 6 dB of 0.79 mm and a measured mean lateral thread width of 1.49 mm compared to a measured axial thread height of 0.69 mm and measured lateral width of 1.21 mm for the processed video sources. The only difference in the elevation direction was that analog video performed better than unprocessed video for 20 dB.

### **Visual comparisons of small, barely visible structures (test 2c)**

Examples of the images presented to the eight volunteers are shown in Figure A.8. The original images allowed more accurate structure detection than the anyplanes through the reconstructed volumes both for the number of structures identified (Figure A.9, Table A.4) and the difficulty to identify the structures, by deviating from the true number of structures (eleven) by a median of 0. The deviation for the large kernel had a median of 2.5 structures, which was statistically significantly better than the PNN reconstruction with a median deviation

of 4 structures from the true number. The small kernel could not be separated from the other reconstructions with a median deviation of 3 structures. For the scans along the cylinders the anyplanes from the PNN reconstruction did not allow the volunteers to identify the same number of structures as the other reconstructions with a median deviation of 3 structures compared to the others with medians of 1 structure deviation. The structures in anyplanes from the PNN reconstructions were more difficult to identify than the small kernel. In the scans across the cylinders there were no statistically significant differences between the reconstruction algorithms with a combined median deviation of 6 structures. The results from comparing only two groups with each other also showed that the structures in the volumes from the PNN reconstruction were statically more difficult to identify than both the other reconstruction algorithms for the scan direction along the cylinders and for both directions combined. The image sources could not be statistically significantly separated for the scan along the cylinders with a combined median deviation from the true number of 1 structure. In the scan across the cylinders, the unprocessed data performed poorer than the other image sources with a median deviation of 7 structures compared to a median of 6 for the digital video and 5 for the analog video. The digital and analog video could, however, not be separated through statistical significance. When combining both scan directions, the analog video with a median deviation of 3 performed statistically significantly better than the unprocessed data with a median deviation of 4. The digital video with a median deviation of 3 could not be separated from the other algorithms.

### **Visual comparisons of tissue data (test 2d)**

Figure A.10 contains examples of the images used in the comparisons showing both original ultrasound images and anyplanes through reconstructed volumes, all originating from the same position. The differences between the reconstruction algorithms were not statistically significant, but all performed worse than the original images both with respect to ordering and quality. The data sources could not be sepa-

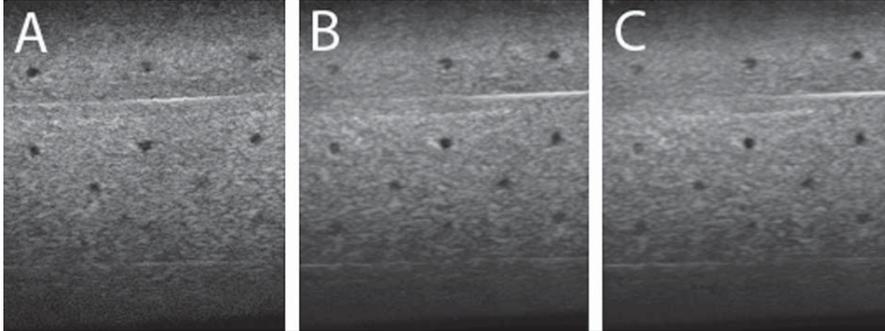


Figure A.8: Example of anyplane images of small structures (cylinders). The anyplanes are from reconstructions based on digital video as input. The scan direction of the input images are along the cylinders. (A) PNN reconstruction. (B) Small kernel reconstruction. (C) Large kernel reconstruction.

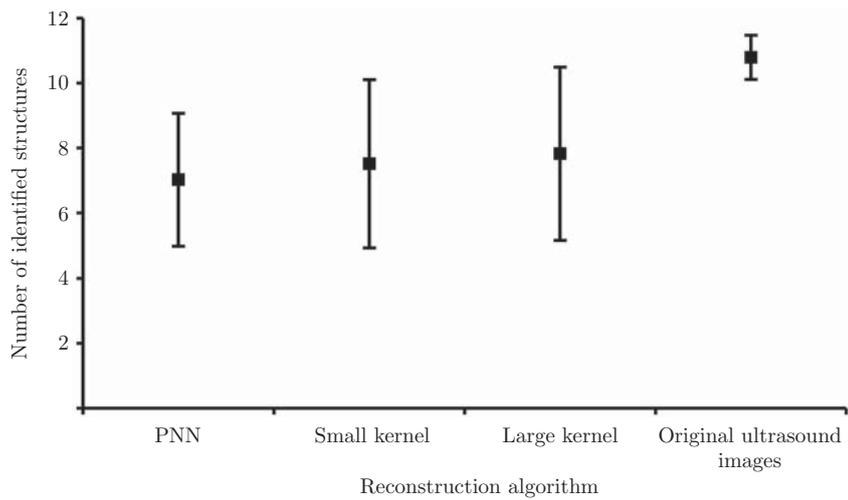


Figure A.9: Number of identified structures. The results are shown as the deviation from the true number of structures (=11) presented as mean and standard deviation values. Results from the three different reconstruction algorithms compared with original ultrasound images. The results from both the scans across and along the structures are combined.

Table A.4: Comparisons based on the number of identified structures. The values shown are the deviation from the true number of structures (=11) presented as mean and standard deviation values.

|  | Analog    | Digital   | Unprocessed |
|--|-----------|-----------|-------------|
| <i>Both scan directions combined</i>               |           |           |             |
| PNN  | 3.63±1.47 | 3.75±2.02 | 4.52±2.46   |
| Small kernel                                       | 3.17±2.30 | 3.23±2.55 | 4.04±2.82   |
| Large kernel                                       | 2.83±2.30 | 2.90±2.78 | 3.77±2.81   |
| Original   | 0.25±0.77 | 0.19±0.54 | 0.19±0.75   |
| <i>Scan across cylinders (elevation direction)</i> |           |           |             |
| PNN  | 4.54±1.28 | 5.04±1.73 | 6.62±0.97   |
| Small kernel                                       | 5.12±1.36 | 5.46±1.56 | 6.58±0.93   |
| Large kernel                                       | 4.88±1.39 | 5.29±1.83 | 6.42±0.93   |
| <i>Scan along cylinders (lateral direction)</i>    |           |           |             |
| PNN  | 2.71±1.00 | 2.46±1.35 | 2.42±1.47   |
| Small kernel                                       | 1.21±0.98 | 1.00±0.66 | 1.50±1.38   |
| Large kernel                                       | 0.79±0.41 | 0.50±0.72 | 1.12±0.80   |

rated in the scans along the forearm, and the quality score could also not be separated in the scans across the forearm. The results from the ordering in the scans across the forearm showed that the digital video was preferred, followed by the unprocessed data, and the analog video was evaluated as the worst quality. When combining the answers from both scan directions the ordering gives the same result as for scans across the forearm, but now also the quality score shows the analog video to perform poorer than the other sources.

### Compare the correctness of the reconstructed volumes geometry (test 3)

None of the reconstruction algorithms could be separated from the original ultrasound images (used as gold standard) while also being

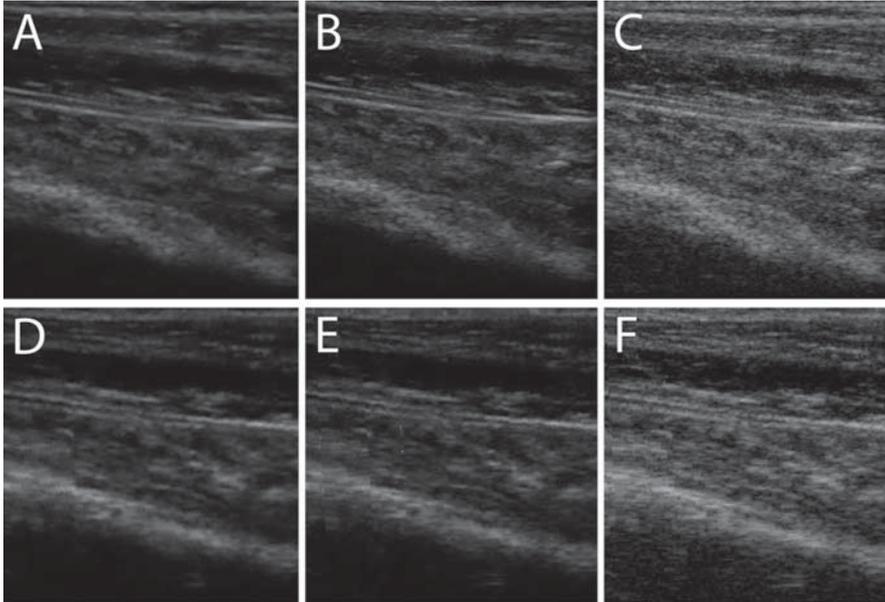


Figure A.10: Original ultrasound images and anyplanes of the underside of the forearm. All images are originating from the same position. All anyplane images are reconstructed with a small kernel, and the anyplanes are obtained orthogonal to the reconstruction input images (See Figure A.4C and D). (A) Original analog image. (B) Original digital image. (C) Original unprocessed image. (D) Anyplane with analog ultrasound video as reconstruction input. (E) Anyplane with digital ultrasound video as reconstruction input. (F) Anyplane with unprocessed ultrasound data as reconstruction input.

Table A.5: Comparisons of reconstruction algorithms based on the measured length (= 6.66 mm), width (= 3.06 mm) and height (= 3.07 mm) of a small cylinder. The absolute of the differences in mm are shown for the results that are statistically significantly different from the original images (“gold standard”). No results were statistically significant when combining both scan directions. NS means not statistically significant.

| US Scan direction      | Elevation (length) | Lateral (width)  | Axial (height) |
|------------------------|--------------------|------------------|----------------|
| <i>Along cylinder</i>  |                    |                  |                |
| PNN                    | 0.27               | NS               | NS             |
| Small kernel           | 0.36               |                  |                |
| Large kernel           | 0.42               |                  |                |
|                        | Elevation (width)  | Lateral (length) | Axial (height) |
| <i>Across cylinder</i> |                    |                  |                |
| PNN                    | NS                 | 0.13             | NS             |
| Small kernel           |                    | 0.22             |                |
| Large kernel           |                    | 0.25             |                |

different from the other reconstruction algorithms (Table A.5). For the scans across the cylinder the analog and digital video differs from the original images in the length measurements, by a difference in measured mean value of 0.33 mm for the analog video and 0.16 mm for the digital video (Table A.6). The unprocessed data differs from the original images in the width measurements, by a difference in measured mean width of 0.28 mm (Table A.6).

## A.4 Discussion

In image-guided surgery it is important to have reliable images available from all directions, and the quality of the reconstructed 3D ultrasound data to be used for guidance and therapy is important. In this paper we have studied if the choice of 3D reconstruction algorithm

Table A.6: Comparisons of input ultrasound sources based on the measured length (= 6.66 mm), width (= 3.06 mm) and height (= 3.07 mm) of a small cylinder. The absolute of the differences in mm are shown for the results that are statistically significantly different from the original images (“gold standard”). No results were statistically significant when combining both scan directions. NS means not statistically significant.

| US Scan direction      | Elevation (length) | Lateral (width)  | Axial (height) |
|------------------------|--------------------|------------------|----------------|
| <i>Along cylinder</i>  |                    |                  |                |
| Analog                 | 0.29               | NS               | NS             |
| Digital                | 0.32               |                  |                |
| Unprocessed            | 0.44               |                  |                |
|                        | Elevation (width)  | Lateral (length) | Axial (height) |
| <i>Across cylinder</i> |                    |                  |                |
| Analog                 | NS                 | 0.33             | NS             |
| Digital                | NS                 | 0.16             |                |
| Unprocessed            | 0.28               | NS               |                |

and input data source may affect image quality and resolution in a 3D volume as evaluated by various methods. A quick summary of the results is presented in Table A.7.

### Comparing the performance of the reconstruction algorithms

Our study showed that the performance of the 3D ultrasound reconstruction algorithms varied from test to test. The small kernel reconstruction algorithm performed slightly better overall than the two other tested algorithm implementations, but a conclusion depends on how the tests are being weighted due to importance. If all tests are to be equally weighted, the choice of reconstruction algorithm doesn't matter, and a good choice may be the fastest algorithm, PNN. All

Table A.7: Short summary of the most important results from the different comparison tests.

|   | Reconstruction algorithms                               | Data sources                                   |
|---|---|--|
| Test 1a: Remove data                        | Large k. worst (for varying images)                     | –  |
| Test 2a: Identification of close structures | Equal performance                                       | Unprocessed worst                              |
| Test 2b: Image resolution measurements      | Equal performance                                       | Unprocessed worst                              |
| Test 2c: Identification of small structures | PNN worst   | Unprocessed worst                              |
| Test 2d: Tissue data comparisons            | Equal performance                                       | Digital best, Unprocessed middle, Analog worst |
| Test 3: Geometry measurements               | PNN barely best, Large k. middle, Small k. barely worst | Indecisive (Digital slightly better)           |

quality measures may, however, not be equally important, and the preferred algorithm may be selected based on intended application.

The method of removing input data and measuring the reconstruction algorithms success at recreating removed data is the “classical” comparison test of reconstruction algorithms [35]. For the scan along the forearm, the PNN algorithm surprisingly has the best performance for the analog and digital video. However, this scan direction creates images that change only slightly from one image to the next. For the unprocessed data, however, the PNN reconstruction has the worst performance, probably because of the increased level of information in

these images. The scans across the forearm were more interesting, since they gave more variation from one image to the next, and the ability to perform well under such situations is more important in a practical situation, so these results were given more weight in Table A.7 and our conclusions. For the scans across the forearm the reconstruction with the largest ellipsoid kernel around input pixels performed worse than the smallest kernel for the digital video and the unprocessed data. The reconstruction with the largest kernel performed also worse than the PNN reconstruction for the digital video. The reason for the poorer performance of the large kernel may be that we used a lateral resolution for the reconstruction based on a depth of 32 mm when the image depth was 40 mm, and thus blurring the data too much in a large part of the reconstructed volume.

In addition to test 1, only test 2c and test 3 showed any statistical differences in reconstruction algorithms. All the comparisons of counted small structures in test 2c showed that the PNN reconstruction performed poorer than the small kernel in two cases and also poorer than the large kernel in two cases (Table A.4). When only comparing two groups at a time in test 3, the PNN algorithm performed better than the small kernel in one comparison and better than the large kernel in three. However, the differences in measured mean values were very small: less than the size of the volume cells ( $< 0.2$  mm), so even if the differences were statistically significant they may not result in much practical difference.

In image-guided surgery, it is most important to identify very small structures in the correct location, so the small or large kernel may be a better choice than the current implementation of the PNN algorithm. However, it should be noted that the small structures in test 2c were so small that a minor position change made the structures disappear. In the test we used the position of a single ultrasound image showing all the structures. The nature of the PNN reconstruction algorithm implementation just replacing existing data in the 3D volume with the latest data might have changed the position of the small structures just enough that they could not be identified. This does not mean that they could not have been identified in a neighboring position, just that the

PNN reconstruction algorithm introduced a small position bias. It is possible to remove this bias by changing the bin-filling [24] of the PNN implementation to an averaging [30, 34] or by keeping the maximum value [30] instead of using only the most recent value.

### **Comparing reconstructions with directly acquired original ultrasound images**

All tests, except test 1, may be used to test for differences between original, directly acquired 2D ultrasound images and anyplanes through the reconstructed volumes. Most tests showed the original ultrasound images to perform better than the anyplanes, but test 2b showed the opposite.

An interesting result from test 2b was that the means acquired for resolution measurements on anyplanes from the reconstructions showed better resolutions than those measured on the original images (Table A.3). The downsampling of the original ultrasound images seemed to lead to slightly better results, but this difference was not enough to explain why the anyplanes in the lateral direction from the reconstructions performed better than original ultrasound images for this test. Also, in the axial direction for the 6 dB reduction, the small kernel obtained better results than both sets of original data, and when fewer groups were compared also the other algorithms performed better than both sets of original data. However, the explanation to the differences may be that the reconstruction algorithms processed the input data to make full use of the 8-bit range, while the original images were unprocessed. This processing increased the distance between the pixel values and the 6 dB and 20 dB measurements were not at the same level for the anyplanes and the original images.

Another interesting observation was obtained when comparing results for 6 dB reductions in Table A.3 with the theoretical results shown in Figure A.3. The differences between measured and theoretical resolutions were quite large: The measured axial resolution was about 4.7 times larger than the theoretical, the measured lateral resolution about 2.7 times larger and the measured elevation resolution about 2.1

times larger. For the axial direction, the assumption about a very short transmitted pulse for the theoretical calculation was clearly too optimistic, and the transmitted pulse was probably several wavelengths long. The thickness of the thread may also lead to a slightly increased measurement, even if the thread diameter was only 0.1 mm. When imaging the thread we made sure that a focus point was at the approximate depth of the thread so that the lateral resolution should be comparable to the theoretical values. However, the formula for the lateral resolution (A.3) did not take into account apodization, and this may explain some of the difference. The available data of the ultrasound probe may also provide additional information: The probe center frequency is 7.2 MHz, and even if the imaging frequency was set to 10 MHz, the real center frequency of the submitted pulse may be closer to 7.2 MHz. In addition the absorption in the imaged area of the ultrasound phantom was 0.7 dB/(cm MHz). Using  $f_t = 7.2$  MHz in (A.3) and  $a = 0.7$  dB/(cm MHz) in (A.4) resulted in the measured lateral resolution being only 1.8 times larger than the theoretical resolution, and measured elevation resolution being only 1.4 times larger than the theoretical. The elevation focus is fixed at depth 16 mm, and compared to the wire depth of 27 mm the difference for the elevation direction was quite understandable. To explore the differences between the theory and practical results further it may be necessary to perform hydrophone measurements of the transmitted pulse, especially to determine the differences in axial resolution.

Test 2c showed that directly acquired ultrasound images performed better than the reconstruction algorithms for the purpose of identifying small structures in phantoms. Test 2d showed that the human observers preferred directly acquired ultrasound images to anyplanes through reconstructions, also for tissue data. Test 3 measuring all axes of a cylinder showed that in most cases the distances measured on anyplanes was not different from directly acquired ultrasound images. Still, Tables A.5 and A.6 showed that images related to the lateral and elevation reconstruction directions gave significant differences from the original ultrasound data for larger distances (lengths measurements) but not for smaller measurements (width measurements),

except for the unprocessed data. Another interesting result was that all the length measurements in anyplanes in the lateral direction (scans across cylinder) were larger than the measurements in the original images while the length measurements in the elevation direction (scans along cylinder) were smaller. The reason for this difference may be inaccuracies in the probe calibration, as this comparison test is dependant of accurate orientation. The original images for the length measurements were obtained orthogonal to the input images in the scan along the cylinder, and a small difference in orientation may have resulted in a relatively large difference in length measurement.

### **Difference between data sources**

Both processed video sources performed generally better than the unprocessed data. The digital video was the data source with the best results throughout all the tests. However, it could not be separated from the analog video in tests 2a, 2b, 2c and 3. On the other hand, it performed better than the analog video in the tissue comparison test 2d. The analog video showed good results in several tests and was often better than the unprocessed data, but not statistically significantly better than the digital video in any of the comparisons. However, in the volunteers evaluation of the tissue data, the analog video performed poorer than both the digital and unprocessed video.

When observing the input video sources, it was obvious that the unprocessed data had more information with the higher resolution, while the processed video sources represented a resampling of the sampled data, which could lead to a loss of information. So far not much work is published comparing the quality of the unprocessed with processed data. Use of the unprocessed data could probably lead to better results in comparisons; however, our results indicate the opposite conclusion. The reason for this may be that the reconstruction algorithms do not process the data as thoroughly as the ultrasound scanner does. This result may be seen as a proof that the image processing done in the ultrasound scanner does improve the image quality as seen from the user perspective even if the resolution may be somewhat reduced. An-

other important aspect of this is that if unprocessed data are used directly there may be a reduction of quality and the images may be harder to interpret even if the data contain more information. To prevent this, a processing similar to that of the ultrasound scanner could be performed. It should however be noted that all comparisons are based on data from a single scanner using one ultrasound probe, and the quality of the unprocessed and processed data may vary between different probes and scanners, and the results of our study can hence not be generalized.

### Comparing input sources and reconstruction algorithms

Several studies try to compare the performance and quality of the reconstruction algorithms [24, 31, 32, 35, 36, 48, 49]. The most commonly used comparison method is the method of removing a percentage of the input data and then determining the algorithms ability to recreate this data. An RMS error value or an equivalent value is usually used to disclose the quality difference. A specific slice is selected, usually in the middle of the ultrasound volume, and a set of different percentages are selected for removal of data [35]. Coupé et al. [48] uses a variation over this and removes input slices (1, 2, 3, 4, 5 and 6) over the whole input data set and calculates a mean and standard deviation for the MSE (Mean Squared Error) values of the different slices. The advantage of this method is that it is an objective method that may be done automatically and is usable with most kinds of ultrasound data.

Drawbacks of this method are that it only compares one aspect of reconstruction quality and the results are dependant of the imaged tissue. In addition, this comparison method is not suitable to compare different ultrasound data sources as we have done in this paper. We have still used this existing method, but in addition we have devised several additional tests for the purpose of comparing both reconstruction algorithms and data sources for the reconstructions.

It is important to note that in our use of the test we performed the statistical comparisons over the full range of removed data (0 %, 25 %, 50 %, 75 %, 100 %, 300 %, 500 % and 700 %) without regard to

the data having more samples in the range from 0% to 100%, leading to a higher weight given to this range. Also, a situation may arise where the RMS values of one reconstruction algorithm are better for one range of percentages but worse for another range when compared to another reconstruction algorithm. In this situation, the statistical tests of the whole range may be indecisive, while a test of a smaller range might give a result. The described situation happened for the combination of all data sources in the scan along the forearm, and as may be seen in Table 1 this test could not differentiate between the data sources. We decided not to perform any additional tests of specific smaller ranges as we wanted to focus on the overall score.

Several of the tests we have used in this paper deal with the practical resolution of the ultrasound images and volumes, especially tests 2a, 2b, 2c and 3. However, test 2c was also highly dependent on accurate positions of the small structures or the ultrasound probe positions related to the structures. Some of the differences between the reconstruction algorithms and original 2D ultrasound images may be due to small errors in the probe calibration having a total mean error of 1.05 mm. The structures being only 1.5 mm in diameter and 2.4 mm long may easily be missed with this calibration accuracy since the comparisons were based on a position and orientation with the structures visible in the original images. A small error in either position or orientation may lead to several missing structures. One of the important features of ultrasound is how good the image quality appear to a human and how easy it is to interpret the image. Test 2d try to test these aspects, but the test is very unspecific. To detect more important differences a set of specific tests could be created asking volunteers (preferably clinical ultrasound users) to evaluate sets of images from different clinical cases.

We have tried to create tests that may represent different uses of the ultrasound volumes, but the range of tests is not exhaustive. Our comparison tests also have a varying degree of manual interaction and may not be practical in all situations. Several other tests could be devised in addition to the tests we have performed in this paper. One idea is to perform an automatic segmentation of a known structure and

compare volume sizes [33]. However, our initial experiments showed that the segmentation algorithms that we had access to created quite different segmented volumes from the same reconstructed ultrasound volume with only small changes in parameters, and it was difficult to get comparable segmentations from volumes originating from different input sources. Another possible test is the comparison of measurements of relatively large distances. An automated method like the one described in Lindseth et al. [51] might be possible, and this test may detect more and larger differences than the test we have performed in this paper, which only compares measurements of small distances.

The general problem with comparison tests for reconstructed ultrasound volumes is that there are several error sources that may affect the results, and these error sources may be larger than the factors we want to compare. Our method of comparing several aspects of the reconstructed volumes is a way to limit the error sources when looking at all tests combined. Still the error sources should be minimized before the comparisons. We have also applied statistical comparisons of all the data collected in the tests since this gives a better way of knowing if one value really is better than another or if it only is natural variation.

### Relevance and further work

3D probes may be used for acquisition of 3D data instead of 2D probes. 3D probes will allow both easy acquisition of near real-time 3D data and the possibilities of 3D data with time as a fourth dimension. However, in IGS, tracked 2D probes have a few advantages over 3D probes: The data stream from the ultrasound probe is limited, so still it is possible to get a higher resolution volume from a 2D probe. In IGS, positioned intraoperative data is easier to correlate to preoperative data, so a tracking sensor is usually needed. 3D probes are shown to interfere more with electromagnetic tracking systems than 2D probes [52]. The volume covered by a 3D probe is much smaller than is practical in most IGS applications, and to overcome this limitation, an application must be created that combines data from the 3D probe.

The data from the 3D probe are not easily available to third party ISG applications, needed to integrate the data with preoperative data, or for combining the data into larger volumes.

Most new ultrasound scanners support the DICOM standard so that it should be possible to get access to digital data for 3D reconstructions. Even if only a stack of 2D images is stored, a 3D reconstruction is still possible with this data if the 2D images are tagged with accurate global time-tags. The positions may then be recorded at the same time with their own time-tags, and a time calibration [44, 53] may be performed to match the positions with the images. Most applications for 3D ultrasound reconstructions require a near real-time implementation, and currently the video grabber approach is the fastest solution in most cases for freehand 3D reconstructions, as DICOM dont support real-time images yet. However, DICOM working group 24 focuses on developing DICOM standards and services for Image Guided Surgery. [54] This work may result in real-time protocols that open up for easily accessible real-time digital ultrasound images suitable for 3D reconstructions. Another option is the method we have used in this paper by using an ultrasound scanner that allows real-time access to digital data.

All comparisons of data in this paper are performed with data from a single scanner using one ultrasound probe. To get more general results it would be interesting to collect data from more scanners and probes. Only a few ultrasound scanners supply unprocessed data (e.g. RF) to third parties without special collaboration agreements. However, a study using data from various scanners and probes could make it easier to arrive to general conclusions regarding choice of input data. Examples of scanner producers with scanners that may supply unprocessed data are: Winprobe (North Palm Beach, FL, US), VeraSonics (Redmond, WA, US), and Terason (Burlington, MA, US), in addition to the Sonix RP scanner from Ultrasonix used in our study.

## A.5 Conclusion

The present study shows that the choice of data source may be more important than the choice of reconstruction algorithm in order to achieve high quality image volumes from tracked freehand 2D ultrasound data. Furthermore, scan converted digital and analog data gave better results than unprocessed ultrasound data in our comparison tests performed with one scanner. Overall, digital video performs slightly better than analog video, but in most cases the two video sources were difficult to separate by the comparison methods, indicating that the quality loss of using flexible video-grabbing solutions over scanners with a digital interface may not be significant. It must be taken into account that the conclusions on data source quality were based on comparisons performed on data from a single scanner using one ultrasound probe, and that quality and processing may differ between probes and scanners. More work including comparison tests on several scanners and probes should therefore be performed in order to obtain more general conclusions. By giving all the comparisons the same weight, no reconstruction algorithm of those tested performs significantly better than the others in terms of image quality indicating that the fastest reconstruction method should be chosen, e.g. PNN. However, in general, the reconstruction algorithm must be selected according to the application and the intended usage of the 3D volume.

## Acknowledgements

This work was supported by the Research Council of Norway, through the FIFOS Programme Project 152831/530 and Medical Imaging Laboratory (MI Lab); the Ministry of Health and Social Affairs of Norway, through the National Centre of 3D Ultrasound in Surgery; and by SINTEF Technology and Society. We want to thank Thomas Kuiran Chen and Purang Abolmaesumi at Queen's University, Kingston, Canada, for supplying us with both a phantom and C++ code for ultrasound probe calibration. We also want to thank Reidar Brekken (SINTEF,

NTNU), Rune Hansen (SINTEF, NTNU) and Hans Torp (NTNU) for discussions regarding ultrasound theory, Sjur Gjerald (NTNU) for aid with laboratory work, Thor Andreas Tangen (NTNU) and Jochen Deibele (NTNU) for help with the Ultrasonix scanner, and we like to thank the rest of the people at SINTEF Technology and Society, Medical Technology for volunteering for some of the comparison tests. The authors are thankful to the reviewers, who provided comprehensive and constructive suggestions and indicated that in addition to Ultrasonix there are two other systems on the market that allow access to the RF data: Winprobe (and possibly Verasonics).

## References

1. Unsgård, G, Rygh, OM, Selbekk, T, Müller, TB, Kolstad, F, Lindseth, F, and Hernes, TAN. Intra-operative 3D ultrasound in neurosurgery. *Acta Neurochirurgica* 2006;148:235–253.
2. McCann, HA, Sharp, JC, Kinter, TM, McEwan, CN, Barillot, C, and Greenleaf, JF. Multidimensional ultrasonic imaging for cardiology. *Proceedings of the IEEE* 1988;76:1063–1073.
3. Hottier, F and Billon, AC, eds. 3D echography: status and perspective. Vol. F 60. *3D Imaging in Medicine*. Springer-Verlag Berlin Heidelberg, 1990.
4. Nadkarni, SK, Boughner, D, and Fenster, A. Image-based cardiac gating for three-dimensional intravascular ultrasound imaging. *Ultrasound in Medicine & Biology* 2005;31:53–63.
5. Prager, RW, Gee, A, Treece, G, and Berman, L. Freehand 3D ultrasound without voxels: volume measurement and visualisation using the Stradx system. *Ultrasonics* 2002;40:109–115.
6. Meairs, S, Beyer, J, and Hennerici, M. Reconstruction and Visualization of Irregularly Sampled Three- and Four-Dimensional Ultrasound Data for Cerebrovascular Applications. *Ultrasound in Medicine & Biology* 2000;26:263–272.

7. Welch, JN, Johnson, JA, Bax, MR, Badr, R, and Shahidi, R. A real-time freehand 3D ultrasound system for image-guided surgery. In: *Ultrasonics Symposium*. Ed. by Schneider, SC, Levy, M, and McAvoy, BR. Vol. 2. IEEE, 2000.
8. Sherebrin, S, Fenster, A, Rankin, RN, and Spence, D. Freehand three-dimensional ultrasound: implementation and applications. In: *Medical Imaging 1996: Physics of Medical Imaging*. Ed. by Van Metter, RL and Beutel, J. Vol. 2708. SPIE, 1996:296–303.
9. Ogawa, S, Itoh, K, Omoto, K, Cheng, X, Ohya, A, and Akiyama, I. Three dimensional ultrasonic imaging for diagnosis of breast tumor. In: *1998 IEEE Ultrasonics Symposium*. Ed. by Schneider, SC, Levy, M, and McAvoy, BR. Vol. 2. IEEE, 1998:1677–1680.
10. Amin, VR, Wang, B, Sonka, M, and Lauer, RM. Three-dimensional ultrasound imaging of vessel wall for evaluating atherosclerosis risk and disease. In: *Medical Imaging 2002: Ultrasonic Imaging and Signal Processing*. Ed. by Insana, MF and Walker, WF. Vol. 4687. SPIE, 2002:255–263.
11. Rohling, RN, Gee, AH, and Berman, L. Automatic registration of 3-D ultrasound images. *Ultrasound in Medicine & Biology* 1998;24:841–854.
12. Gilja, OH, Hausken, T, Olafsson, S, Matre, K, and Odegaard, S. In vitro evaluation of three-dimensional ultrasonography based on magnetic scanhead tracking. *Ultrasound in Medicine & Biology* 1998;24:1161–1167.
13. Detmer, PR, Bashein, G, Hodges, T, Beach, KW, Filer, EP, Burns, DH, and Jr., DS. 3D ultrasonic image feature localization based on magnetic scanhead tracking: In vitro calibration and validation. *Ultrasound in Medicine & Biology* 1994;20:923–936.
14. Moskalik, AP, Carson, PL, Meyer, CR, Fowlkes, JB, Rubin, JM, and Roubidoux, MA. Registration of three-dimensional compound ultrasound scans of the breast for refraction and motion correction. *Ultrasound in Medicine & Biology* 1995;21:769–778.

15. Laporte, C and Arbel, T. Combinatorial and Probabilistic Fusion of Noisy Correlation Measurements for Untracked Freehand 3-D Ultrasound. *IEEE Transactions on Medical Imaging* 2008;27:984–994.
16. Solberg, OV, Tangen, GA, Lindseth, F, Sandnes, T, Enquobahrie, AA, Ibanez, L, Cheng, P, Gobbi, D, and Cleary, K. Integration of a real-time video grabber component with the open source image-guided surgery toolkit IGSTK. In: *Medical Imaging 2008: PACS and Imaging Informatics*. Ed. by Andriole, KP and Siddiqui, KM. Vol. 6919. SPIE, 2008:69190Z–69199.
17. Ohbuchi, R, Chen, D, and Fuchs, H. Incremental volume reconstruction and rendering for 3-D ultrasound imaging. In: *Visualization in Biomedical Computing '92*. Ed. by Robb, RA. Vol. 1808. SPIE, 1992:312–323.
18. Chen, TK, Thurston, AD, Ellis, RE, and Abolmaesumi, P. A Real-Time Freehand Ultrasound Calibration System with Automatic Accuracy Feedback and Control. *Ultrasound in Medicine & Biology* 2009;35:79–93.
19. Berg, S, Torp, H, Martens, D, Steen, E, Samstad, S, Hoivik, I, and Olstad, B. Dynamic three-dimensional freehand echocardiography using raw digital ultrasound data. *Ultrasound in Medicine & Biology* 1999;25:745–753.
20. Martens, D and Gilja, OH. The EchoPAC-3D software for image analysis. In: *Basic and New Aspects of Gastrointestinal Ultrasonography*. Ed. by Ødegård, S, Gilja, OH, and Gregersen, H. Advanced series in Biomechanics. World Scientific Publishing, 2005. Chap. 10:305–329.
21. Thune, N, Gilja, OH, Hausken, T, and Matre, K. A practical method for estimating enclosed volumes using 3D ultrasound. *European journal of ultrasound* 1996;3:83–92.
22. Grønningsæter, Å, Kleven, A, Ommedal, S, Aarseth, TE, Lie, T, Lindseth, F, Langø, T, and Unsgård, G. SonoWand, an ultrasound-based neuronavigation system. *Neurosurgery* 2000;47:1373–1739.

23. Lindseth, F, Bang, J, and Langø, T. A robust and automatic method for evaluating accuracy in 3-D ultrasound-based navigation. *Ultrasound in Medicine & Biology* 2003;29:1439–1452.
24. Solberg, OV, Lindseth, F, Torp, H, Blake, RE, and Hernes, TAN. Freehand 3D ultrasound reconstruction algorithms – a review. *Ultrasound in Medicine & Biology* 2007;33:991–1009.
25. Trobaugh, JW, Trobaugh, DJ, and Richard, WD. Three-dimensional imaging with stereotactic ultrasonography. *Computerized Medical Imaging and Graphics* 1994;18:315–323.
26. Wein, W, Pache, F, Röper, B, and Navab, N. Backward-Warping Ultrasound Reconstruction for Improving Diagnostic Value and Registration. In: *Medical Image Computing and Computer-Assisted Intervention — MICCAI 2006*. Ed. by Larsen, R, Nielsen, M, and Sporring, J. Springer, 2006:750–757.
27. Karamalis, A, Wein, W, Kutter, O, and Navab, N. Fast hybrid freehand ultrasound volume reconstruction. In: *Medical Imaging 2009: Visualization, Image-Guided Procedures, and Modeling*. Ed. by Miga, MI and Wong, KH. Vol. 7261. SPIE, 2009:726114–726118.
28. Treece, GM, Prager, RW, Gee, AH, and Berman, LH. Correction of probe pressure artifacts in freehand 3D ultrasound. *Medical Image Analysis* 2002;6:199–214.
29. Zhang, Y, Rohling, R, and Pai, DK. Direct surface extraction from 3D freehand ultrasound images. In: *IEEE Visualization 2002 — VIS 2002*. Ed. by Moorhead, R, Gross, M, and Joy, KI. IEEE Computer Society, 2002:45–52.
30. Nelson, TR and Pretorius, DH. Interactive acquisition, analysis, and visualization of sonographic volume data. *International Journal of Imaging Systems and Technology* 1997;8:26–37.

31. Estépar, RSJ, Martín-Fernández, M, Caballero-Martínez, PP, Alberola-López, C, and Ruiz-Alzola, J. A theoretical framework to three-dimensional ultrasound reconstruction from irregularly sampled data. *Ultrasound in Medicine & Biology* 2003;29:255–269.
32. Estépar, RSJ, Martín-Fernández, M, Alberola-López, C, Ellsmere, J, Kikinis, R, and Westin, CF. Freehand ultrasound reconstruction based on ROI prior modeling and normalized convolution. In: *Medical Image Computing and Computer-Assisted Intervention — MICCAI 2003*. Ed. by Ellis, RE and Peters, TM. Vol. 2879. Springer, 2003:382–390.
33. Barry, CD, Allott, CP, John, NW, Mellor, PM, Arundel, PA, Thomson, DS, and Waterton, JC. Three-dimensional freehand ultrasound: image reconstruction and volume analysis. *Ultrasound in Medicine & Biology* 1997;23:1209–1224.
34. Gobbi, DG and Peters, TM. Interactive intra-operative 3D ultrasound reconstruction and visualization. In: *Medical Image Computing and Computer-Assisted Intervention — MICCAI 2002*. Ed. by Dohi, T and Kikinis, R. Vol. 2489. Lecture Notes in Computer Science. Springer, 2002:156–163.
35. Rohling, R, Gee, A, and Berman, L. A comparison of freehand three-dimensional ultrasound reconstruction techniques. *Medical Image Analysis* 1999;3:339–359.
36. Sanches, JM and Marques, JS. A multiscale algorithm for three-dimensional free-hand ultrasound. *Ultrasound in Medicine & Biology* 2002;28:1029–1040.
37. Nikolov, SI, González, JPG, and Jensen, JA. Real time 3D visualization of ultrasonic data using a standard PC. *Ultrasonics* 2003;41:421–426.
38. Eulenstein, S, Lange, T, Hünerbein, M, Schlag, PM, and Lamecker, H. Ultrasound-based navigation system incorporating preoperative planning for liver surgery. *International Congress Series* 2004;1268:758–763.

39. Huang, Sw, Kim, K, Witte, RS, Olafsson, R, and O'Donnell, M. Inducing and Imaging Thermal Strain Using a Single Ultrasound Linear Array. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 2007;54:1718–1720.
40. Cinquin, P, Bainville, E, Barbe, C, et al. Computer assisted medical interventions. *IEEE Engineering in Medicine and Biology Magazine* 1995;14:254–263.
41. Cleary, K, Ibáñez, L, Gobbi, D, Cheng, P, Gary, K, Aylward, S, Jomier, J, Enquobahrie, A, Zhang, H, Kim, Hs, and Blake, MB. *IGSTK: The Book. 2.0*. Gaithersburg, Maryland, USA: Signature Book Printing, 2007.
42. Enquobahrie, A, Cheng, P, Gary, K, Ibáñez, L, Gobbi, D, Lindseth, F, Yaniv, Z, Aylward, S, Jomier, J, and Cleary, K. The Image- Guided Surgery Toolkit IGSTK: An Open Source C++ Software Toolkit. *Journal of Digital Imaging* 2007;20:21–33.
43. Gary, K, Ibáñez, L, Aylward, S, Gobbi, D, Blake, MB, and Cleary, K. IGSTK: An open source software platform for image-guided surgery. *Computer* 2006;39:46–53.
44. Treece, GM, Gee, AH, Prager, RW, Cash, CJC, and Berman, LH. High-definition freehand 3-D ultrasound. *Ultrasound in Medicine & Biology* 2003;29:529–546.
45. Gobbi, DG, Lee, BKH, and Peters, TM. Correlation of preoperative MRI and intraoperative 3D ultrasound to measure brain tissue shift. In: *Medical Imaging 2001: Visualization, Display, and Image-Guided Procedures*. Ed. by Mun, SK. Vol. 4319. SPIE, 2001:264–271.
46. Angelsen, BAJ. *Ultrasound Imaging. Waves, signals and signal processing*. Trondheim, Norway: Emantec AS, 2000.
47. Wear, KA. The effects of frequency-dependent attenuation and dispersion on sound speed measurements: applications in human trabecular bone. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 2000;47:265–273.

48. Coupé, P, Hellier, P, Azzabou, N, and Barillot, C. 3D freehand ultrasound reconstruction based on probe trajectory. In: *Medical Image Computing and Computer-Assisted Intervention — MIC-CAI 2005*. Ed. by Duncan, J and Gerig, G. Vol. 3749. Springer, 2005:597–604.
49. Sanches, JM and Marques, JS. A Rayleigh reconstruction/interpolation algorithm for 3D ultrasound. *Pattern Recognition Letters* 2000;21:917–926.
50. Thijssen, JM, Weijers, G, and Korte, CL de. Objective Performance Testing and Quality Assurance of Medical Ultrasound Equipment. *Ultrasound in Medicine & Biology* 2007;33:460–471.
51. Lindseth, F, Langø, T, Bang, J, and Hernes, TAN. Accuracy evaluation of a 3D ultrasound-based neuronavigation system. *Computer aided surgery* 2002;7:197–222.
52. Hastenteufel, M, Vetter, M, Meinzer, HP, and Wolf, I. Effect of 3D ultrasound probes on the accuracy of electromagnetic tracking systems. *Ultrasound in Medicine & Biology* 2006;32:1359–1368.
53. Rousseau, F, Hellier, P, and Barillot, C. A novel temporal calibration method for 3-D ultrasound. *IEEE Transactions on Medical Imaging* 2006;25:1108–1112.
54. Lemke, HU. Summary of the white paper of DICOM WG24 'DICOM in Surgery'. In: *Medical Imaging 2007: PACS and Imaging Informatics*. Ed. by Horii, SC and Andriole, KP. Vol. 6516. SPIE, 2007:651603–651613.

## Paper B

# Accuracy of electromagnetic tracking with a prototype field generator in an interventional OR setting

Lars Eirik Bø<sup>1,2</sup>, Håkon Olav Leira<sup>1,3</sup>, Geir Arne Tangen<sup>1</sup>, Erlend Fagertun Hofstad<sup>1</sup>, Tore Amundsen<sup>1,3</sup> and Thomas Langø<sup>1</sup>

<sup>1</sup>SINTEF Technology and Society, Trondheim, Norway

<sup>2</sup>Norwegian University of Science and Technology, Trondheim, Norway

<sup>3</sup>St. Olavs Hospital, Trondheim, Norway

*Medical Physics*, vol. 39, no. 1, pp. 399–406, 2012

**Abstract**

We have studied the accuracy and robustness of a prototype electromagnetic window field generator WFG in an interventional radiology suite with a robotic C-arm. The overall purpose is the development of guidance systems combining real-time imaging with tracking of flexible instruments for bronchoscopy, laparoscopic ultrasound, endoluminal surgery, endovascular therapy and spinal surgery.

The WFG has a torus shape, which facilitates X-ray imaging through its centre. We compared the performance of the WFG to that of a standard field generator SFG under the influence of the C-arm. Both accuracy and robustness measurements were performed with the C-arm in different positions and poses.

The system was deemed robust for both field generators, but the accuracy was notably influenced as the C-arm was moved into the electromagnetic field. The SFG provided a smaller root-mean-square position error, but was more influenced by the C-arm than the WFG. The WFG also produced smaller maximum and variance of the error.

EM tracking with the new WFG during C-arm based fluoroscopy guidance seems to be a step forward, and with a correction scheme implemented it should be feasible.

**B.1 Introduction**

Electromagnetic-based navigation of flexible instruments has been explored in various applications. The advantage of this technology is that it does not require a clear line of sight, thus permitting the tracking of instruments in confined spaces and inside the body. This is especially useful for flexible instruments such as catheters, endoscopes and endoscopic ultrasound probes. In endovascular therapy, tracking of guide wires, catheters and needles has been tested, especially for treatment of abdominal aortic aneurisms. Here, tracking has been used to guide e.g. deployment of stent grafts. [1, 2] Similar technology has been applied to cardiac interventions, most notably to catheter-based ablation therapies for treatment of atrial fibrillation. [3] In pulmonology,

electromagnetic-based navigated bronchoscopy based on preoperative CT imaging has been introduced, particularly for the diagnosis and targeting (biopsy) of small, peripheral lesions. This has shown to increase the success rate from as low as 30 % to about 67 %. [4, 5] In laparoscopic surgery, electromagnetic (EM) tracking has been proposed, especially to guide flexible laparoscopic ultrasound probes and ablation probes in liver surgery. [6, 7] Preliminary tests combining small ultrasound probes and navigation for spinal surgery also indicate that the added flexibility provided by EM tracking may be advantageous. In addition, several commercial products exist, such as StealthStation Axiem (Medtronic Navigation, Louisville, USA), PercuNav (Philips Healthcare, DA Best, The Netherlands), iGuide CAPPa (Siemens AG, Munich, Germany), i-Logic (SuperDimension GmbH, Dusseldorf, Germany) and ig4 (Veran Medical Technologies, Inc., St. Louis, USA).

In a setup using EM tracking, the placement of the field generator, which is responsible for generating the EM measurement volume, could influence the setup of other equipment and potentially restrict the movement of the medical personnel. It may also influence how imaging can be performed, e.g. by obstructing X-rays at certain angles. Northern Digital Inc. (NDI) has designed a new prototype field generator, referred to as a window field generator, in an attempt to bypass some of these limitations. This field generator has a torus shape, and the central opening facilitates X-ray imaging with the field generator mounted directly underneath the operating table. This way, the field generator is out of the way and may stay in place throughout the procedure, even when X-ray imaging is required.

EM tracking is vulnerable to disturbances from ferromagnetic interference sources in the surroundings, which may influence the accuracy of the system. It is therefore important to assess the accuracy, not only for each system, but also for each new location where the system is to be used. If there are disturbances that are constant and may be properly characterized, they may be compensated using static correction schemes. [8–10] However, since the interference depends on the surroundings, it must be characterized for each new location and the correction scheme must be adapted accordingly. Also, if the environ-

ment changes during the procedure, e.g. by introduction of additional equipment, this must be taken into account.

Wilson et al. [11] have presented a protocol for accuracy evaluation of EM tracking and applied it to various operating room (OR) settings. A similar protocol was adapted by Yaniv et al. [12] in their overall assessment of EM tracking in the clinical environment. The two works report root-mean-square (RMS) errors in the range of 0.79 mm to 6.67 mm and 0.38 mm to 6.49 mm respectively for various combinations of tracking systems and OR environments. Yaniv et al. have also considered the robustness of the various systems. The robustness is a measure of the resistance to distortion upon introduction of additional equipment to the work volume. If the system is robust, this means that any disturbance is constant and fixed relative to the EM transmitter. This makes it possible to apply a static correction scheme.

In this article, we have adapted the protocol described by Yaniv et al. to study the accuracy and robustness of an EM tracking system within the setting of a new interventional radiology suite. The goal was to compare the performance of the new prototype field generator with that of the original field generator, and to study the influence of a new C-arm on the tracking system.

## **B.2 Materials and methods**

### **Experimental setup**

Our study was done in one of the new interventional radiology suites at St. Olavs University Hospital in Trondheim. The suite, which was opened in august 2010, is a part of the project The Operating Room of the Future (see <http://www.stolav.no/for>) and is used for both research projects and routine clinical procedures within interventional radiology. The OR is equipped with a robotically controlled cone beam CT imaging system (Siemens Artis zeego, Siemens Healthcare, Forchheim, Germany) referred to as a C-arm.

For position tracking, we used the Aurora Electromagnetic Measurement System (NDI, Waterloo, Canada). The system consists of

the utility software NDI ToolBox, a system control unit, four system interface units for position sensor inputs and a field generator that generates an EM tracking volume. This is shown in Figure B.1. In our setup, we used two different field generators: the standard, commercially available, rectangular field generator (SFG) and the prototype window field generator (WFG). The SFG operates with either a cube-shaped measurement volume with side lengths 0.5 m or an extended, dome-shaped volume with side lengths up to 0.96 m. In this study, we used the cube shaped measurement volume. The WFG operates only with a reduced, dome-shaped measurement volume with side lengths up to 0.65 m.

The main feature of the new prototype is its torus shape, which allows for X-ray images to be taken through the centre opening. However, since the C-arm potentially is a major source of EM interference, the common practice for EM-based navigated procedures has been to remove the C-arm unit from the operating field during navigation. But to a clinician such an approach is cumbersome, as it would be preferable to do imaging and navigation concurrently, or at least intermittently. In an operating room, there are many people, trolleys with equipment, wires and tubes, so that moving the large fluoroscopy unit back and forth is unpractical and time-consuming. In order to take full advantage of the new field generator, it would therefore be advantageous to be able to perform tracking with the C-arm in or close to the operating field. We have therefore analyzed the accuracy and robustness of the tracking in this setting.

All data acquisition was performed using the NDI ToolBox. Four tools with position sensors were used, each having either five or six spatial degrees of freedom (DOF): a Traxtal Reference Tool with six DOF, a custom-designed catheter tool containing an Aurora Micro 6DOF Sensor, an Aurora Tracking Needle with five DOF and an Aurora 6DOF Cable Tool. The tools are shown in Figure B.1 (b). The catheter tool was made from an ordinary, one-lumen catheter with diameter 1.7 mm. The Aurora sensor was inserted into the lumen and fixed with epoxy glue near the tip of the catheter. A pivot calibration procedure was performed using the NDI ToolBox software during tool

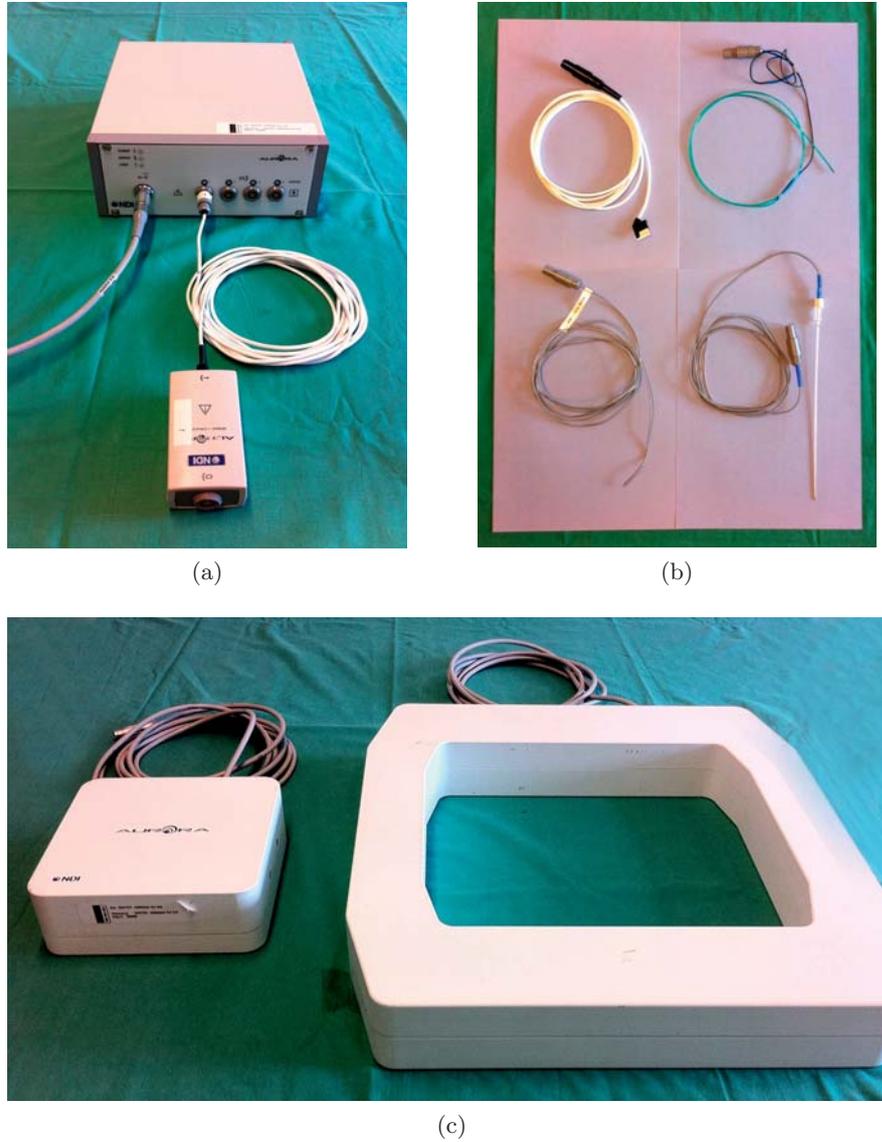


Figure B.1: The Aurora Electromagnetic Measurement System from NDI: (a) system control unit with interface unit for position sensor input, (b) four different tools with position sensors, (c) the standard field generator (left) and the prototype window field generator (right).

characterization to determine the position of the tip of the catheter relative to the embedded sensor.

The field generator was mounted underneath the operating table giving no restrictions to the movement of the robotically controlled C-arm. We used the same accuracy phantom as Yaniv et al. [12]: a Plexiglas cube with sides measuring 180 mm and with 225 parallel holes precisely machined from one side, each with diameter 1.9 mm and depth in the range of 10 mm to 150 mm. The phantom was equipped with a reference tool and placed on the operating table approximately in the centre of the tracking system's measurement volume as shown in Figure B.2. In this way, our experimental setup represented a navigation volume relevant for a clinical setup. This setup was used throughout all of the following experiments.

### **Tracking system accuracy analysis**

The catheter tool was manually inserted into each of the 225 holes in the phantom, all the way to the bottom, and 100 position samples were collected from each hole. Since the diameter of the catheter was only 0.2 mm less than the diameter of the hole, this provided an accurate measurement of the position of the bottom of the hole. The procedure was performed first with the C-arm inside the tracking system's measurement volume, and then repeated with the C-arm outside the measurement volume. These two measurement experiments were done for both the SFG and the WFG.

For data processing, we used a modified version of a MATLAB software application implemented by Wilson et al. [11]. For each of the 225 holes, a representative transformation was calculated from the 100 recorded position samples. The translational part of this transform was estimated as the arithmetic mean of the acquired translation data, and the rotational part as the arithmetic mean of the acquired rotation data, given in unit quaternions, followed by normalization. In addition, the distance from the origin of the tracking system to each of the 100 position samples was calculated, and the sample variability for the given hole, defined as the difference between the largest and the

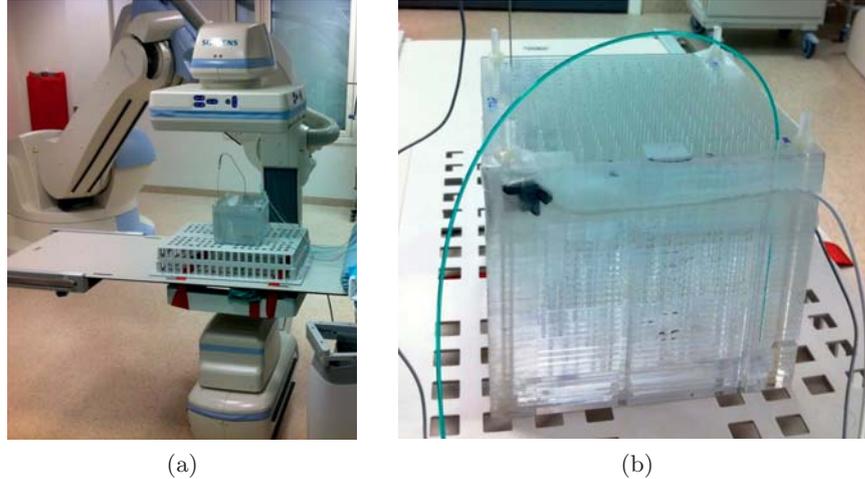


Figure B.2: The accuracy phantom and its setup. The field generators were mounted directly underneath the operating table (the photo shows the WFG). The phantom was equipped with the Traxtal Reference Tool, here seen in the front left corner, and placed on the operating table approximately in the centre of the tracking system's measurement volume. The catheter tool used for the accuracy measurements is seen inserted into the hole in the front right corner of the phantom. For the robustness measurements, two more tools were used: the Aurora Tracking Needle was inserted into the hole in the rear left corner of the phantom and the Aurora 6DOF Cable Tool was attached to the rear right corner.

smallest of these distances, was found. The sample variability is a simple measure of the precision of the performed measurements.

A paired-point rigid registration [13] was then performed between the tracking system and the phantom coordinate system using 9 of the 225 sampled positions. For each of the 225 holes, the distance between the known point coordinate and the estimated representative transformation, transformed by the registration matrix, was calculated. Also, the angular difference between the known orientation of the hole and the measured orientation of the catheter tool was found. The MAT-

LAB application provided the following descriptive statistics: maximal sample variability, RMS error, mean error, standard deviation, error range, maximum error and 95th percentile.

### **Tracking system robustness analysis**

The robustness of an EM tracking system was defined by Yaniv et al. as the system’s “resilience to distortions arising from tools and imaging apparatus that are introduced and removed from the work volume during the procedure” [12]. This can be equipment containing ferromagnetic metal or emitting EM fields. We have focused our work on the influence of the C-arm on the tracking system accuracy since this is the potentially biggest source of EM disturbance that can be introduced into the measurement volume of the tracking system.

In image-guided interventions, it is common to use a patient-mounted reference tool and track all other tools relative to this. Motivated by this practice, Yaniv et al. quantified the robustness by considering the distance between two stationary tracking sensors for a certain period of time. If the variability of the measured distance is low, it means that the system reports the position of one sensor relative to the other in a consistent manner, and the system is then regarded as robust. This does not guarantee that the reported position is correct, i.e. that the system is accurate, but it means that potential inaccuracies may be corrected using static correction schemes. The measurements were done with the imaging apparatus both in home position away from the operating field, and in imaging position during both X-ray fluoroscopy imaging and cone beam CT imaging.

We extended the analysis of Yaniv et al. by also including measurements with the C-arm in a number of intermediate positions and different poses and considering not only the variability of the distance measurements in each position, but also how the measurements change between different positions. If the variability is low in each position, but the measurements change from one position to another, this means that the system is robust and a static correction scheme adapted to the given C-arm position can be applied. We also used four tracking

sensors rather than just two in order to account for the possibility of anisotropic disturbances; in this way, distortions that are perpendicular to the axis between two of the sensors will influence the measured distance between two of the other sensors.

The Plexiglas phantom and the tracking system were both set up as for the accuracy analysis. The phantom here only served to ensure that the measurements were carried out within the same navigation volume that was used for the accuracy analyses. Inside and on the phantom we placed the three tracked tools in addition to the reference tool. As seen in Figure B.2(b), we tried to distribute the tool positions as much as possible both in the horizontal plane direction and in depth so that the distance between the sensors ranged from 160 mm to 230 mm. The C-arm was then moved stepwise relative to the operating table in various manners as illustrated in Figures B.4(a), (d), (g) and (j): the C-arm was translated along and perpendicular to the table, it was rotated around the table and the X-ray detector was lowered towards the table. For each step, position data for all four tools were recorded for a period of 30 s and stored. With the C-arm in imaging position, we also recorded position data during both X-ray fluoroscopy imaging and cone beam CT imaging. The distances from the reference tool to each of the other three tools were then calculated. This procedure was repeated for both field generators.

### B.3 Results

#### Tracking system accuracy

Statistical measurements of the position and angle error of the catheter tool from the four experiments are summarized in Table B.1. A comparison of the results from the SFG and the WFG, without the influence of the C-arm, demonstrates only minor differences. The SFG appears to provide a smaller RMS position error, while the WFG has less maximum and variance of the error. Also note that the maximal variability within the 100 samples of each phantom node is higher for the SFG.

The accuracy measurements show the strong disturbance caused by the C-arm on the EM tracking system. The influence is greater when using the SFG, increasing the RMS position error with about 7 mm, compared to 4 mm increase for the WFG. The results when using the SFG have a larger spread of the error as well. And, when performing the measurements, the Aurora system with the SFG was not able to track the catheter tool in three of the 225 phantom nodes, which thus had to be left out of the calculations.

The RMS angle error and spread of the angle of the catheter pointer were also increased with the C-arm in the field for both field generators. The WFG outperformed the SFG with respect to RMS angle error and spread of the angle, both with and without the influence of the C-arm.

### Tracking system robustness

The variability in measured sensor distance over the period of 30 s for which the C-arm was at a fixed position was relatively small: For all sensors and all positions, the standard deviation of the measurements was below 0.06 mm and the range was below 0.3 mm.

The robustness measurements made during image acquisition are shown in Figure B.3. The plots correspond to those presented by Yaniv et al. [12] and show the distance between the Aurora 6DOF Cable Tool and the reference tool demeaned using the mean of the first 2 s. The imaging sequences were started 10 s after the position measurements, and each sequence lasted approximately 10 s. The measurements show that both field generators are robust to X-ray fluoroscopy imaging, while they are severely influenced by the cone beam CT imaging. The latter caused a deviation of 14.2 mm for the SFG, while the WFG reached a deviation of 30.9 mm before tracking eventually was completely disrupted.

While the tracking system was shown to be robust, except for during cone beam CT imaging, the measured distances between the sensors were strongly affected by the position of the C-arm. In Figures B.4(b), (c), (e) and (f), the mean measured distance from the reference tool to each of the three other tools is plotted as a function

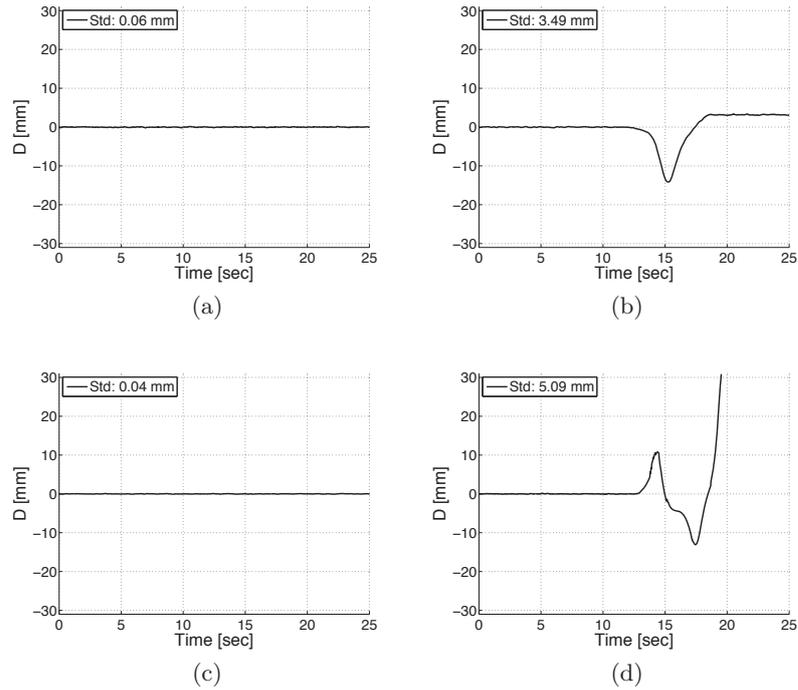


Figure B.3: The plots show the variation in measured distance between the Aurora 6DOF Cable Tool and the reference tool during X-ray fluoroscopy imaging (left column) and cone beam CT imaging (right column) measured with the SFG (top row) and WFG (bottom row) respectively. The distance between the sensors is demeaned using the mean of the first 2 s.

Table B.1: Results of the accuracy measurements performed in the radiology suite. The values in the table are in millimeters and degrees.

|     |                         | C-arm<br>out of field |       | C-arm<br>in field |       |
|-----|-------------------------|-----------------------|-------|-------------------|-------|
|     |                         | Pos.                  | Angle | Pos.              | Angle |
| WFG | Max. sample variability | 0.37                  |       | 0.40              |       |
|     | RMS error               | 1.16                  | 1.11  | 5.09              | 4.06  |
|     | Mean error              | 1.11                  | 0.63  | 4.91              | 3.66  |
|     | Standard deviation      | 0.38                  | 0.73  | 1.63              | 1.69  |
|     | Error range             | 1.75                  | 3.41  | 7.37              | 7.48  |
|     | Maximum error           | 1.87                  | 3.42  | 8.47              | 7.81  |
|     | 95th percentile         | 1.69                  | 2.96  | 7.50              | 6.53  |
| SFG | Max. sample variability | 0.67                  |       | 1.09              |       |
|     | RMS error               | 0.79                  | 1.57  | 7.59              | 9.57  |
|     | Mean error              | 0.56                  | 1.00  | 4.89              | 9.91  |
|     | Standard deviation      | 0.42                  | 0.91  | 4.62              | 2.72  |
|     | Error range             | 2.55                  | 4.82  | 40.07             | 15.42 |
|     | Maximum error           | 2.60                  | 4.90  | 41.35             | 16.14 |
|     | 95th percentile         | 1.44                  | 3.26  | 14.15             | 12.72 |

of the C-arm's displacement along and perpendicular to the operating table. The distances have been demeaned with the mean of the first recording in each series. This shows that the difference between the minimum and the maximum measurement as the C-arm is moved along the operating table into the EM field is up to 8 mm for the SFG and 13 mm for the WFG. The same deviation as a function of the C-arm's rotation around the operating table is shown in Figures B.4(h) and (i). Rotation of the C-arm caused a sensor distance alteration of maximum 7 mm. Finally, the effect on the measured sensor distance resulting from moving the X-ray detector closer to the operating table is shown in Figures B.4(k) and (l). Moving the detector 13 cm caused a reduction of the sensor distance of maximum 4 mm.

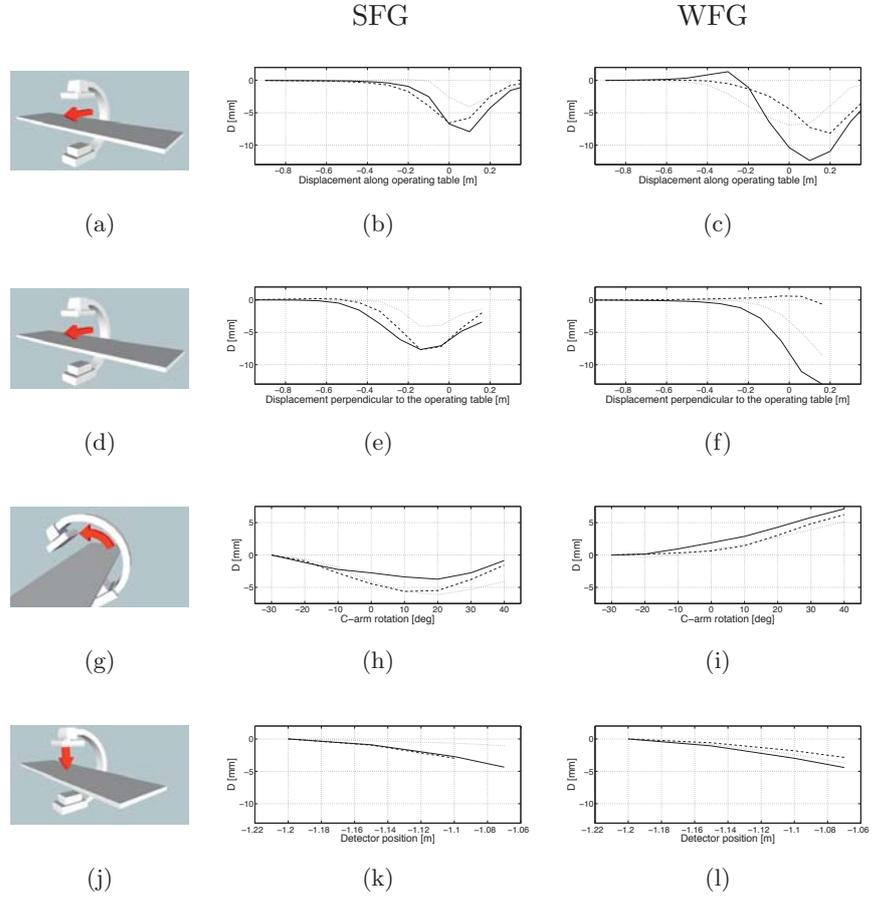


Figure B.4: The figures in the left column indicate the stepwise movement of the C-arm relative to the operating table during the robustness measurements. The plots in the middle column show the mean measured distance from the reference tool to each of the three other tools as a function of the C-arm's displacement using the SFG. The plots in the last column show the same results for the WFG. The distances have been demeaned with the mean of the first recording in each series. The actual distances between the tools were between 160 and 230 mm. In the two upper rows, 0 displacement indicates that the C-arm is straight above the measurement volume, and we see that the deviation has a peak close to this point.

## B.4 Discussion

When the C-arm was placed far from the measurement volume, the measured accuracy was comparable to that reported by the manufacturer, for both field generators. The results for the SFG are also very close to the values found by Yaniv et al. in a similar interventional radiology suite using the same field generator (see Yaniv et al. [12], Table II, columns 7 and 8). They present a lower maximal sample variability, but slightly higher position errors. The angle errors are almost identical.

When comparing the two field generators, the maximal sample variability was highest for the SFG, which indicates a greater need for smoothing of the measurements, e.g. by averaging a certain number of samples. This kind of noise reduction may, however, reduce the frame rate of the system, or at least introduce a certain time lag, which should be avoided for navigation.

As the C-arm was moved close to the measurement volume and placed directly above the field generator, the EM tracking was strongly influenced, and the measured accuracy went down. The sample variability was relatively unaffected, indicating that the measurements were quite stable, but the measurement error was considerably increased. The SFG performed worse than the WFG: it produced more outliers in the accuracy measurements, with a 95th percentile nearly twice that of the WFG, and it was also unable to track the sensors in certain positions within the measurement volume. The WFG did not present any of these problems, and it thus appears to be more stable. However, with 95th percentile of 7.50 mm and 14.15 mm respectively, the error is considerable in both cases.

The robustness measurements showed that the distances between the various sensors varied little as long as the C-arm stayed in one position: the standard deviation of the measured distance was below 0.06 mm and the range was below 0.3 mm for all positions. This is similar to the results of Yaniv et al., who present a standard deviation of 0.05 mm for Aurora in their interventional radiology suite. The system may thus be said to be robust with respect to the OR and C-arm influ-

ence on the EM field. However, the measured distance varied greatly as the C-arm was moved or rotated, which is consistent with the poor accuracy that was measured with the C-arm close to the measurement volume. The variation was largest for the WFG, which might be due to the larger measurement volume potentially exposing it to more ferromagnetic interference sources in the surroundings. However, for both field generators the variation was notable, and it is thus clear that a correction scheme is required in order to use the tracking system with the C-arm in this position.

Since the system is robust, a static correction scheme may be adapted. In its simplest form, such a scheme involves measuring the position and orientation of a position sensor at a number of fixed reference points throughout the measurement field. As the C-arm is introduced, the same measurements are repeated. This will provide us with deformation data for the reference points for the given position of the C-arm. Deformations between these reference points can be determined by different interpolation schemes as described by Kindratenko [9]. By mapping the deformation field with this calibration procedure we are then able to correct any further position data readings. This will however only be valid for the given system setup and C-arm position, meaning that a precalibration process must be performed for all relevant C-arm positions. An improved solution could be to place several position sensors throughout the measurement field forming a set of reference points. Distortions detected by these sensors as the C-arm is moved closer and into the measurement field could be used to characterize the deformation field. This opens the possibility of doing calibration and correction of position data in real time.

Hybrid solutions have been investigated, using a combination of optical and electromagnetic tracking. The optical tracking data are not influenced by metal objects in the environment and can be used as reference points. [8, 10] With this technique we may be able to map the deformation field by performing an intraoperative calibration sequence, moving the hybrid tool through the volume of interest before starting navigation.

Other solutions suggested in the literature include merging posi-

tion data generated from live fluoroscopy images with electromagnetic position tracking data and in this way increase accuracy. [14] Recent work has also shown how statistical models for the tool movement can be used to estimate the tool position. [15–17] In our future work, we will look into and experiment with these techniques to improve the accuracy and performance of navigation in the clinical environment.

## B.5 Conclusion

EM tracking with the new WFG during C-arm based fluoroscopy guidance seems to be a step forward, and with a correction scheme implemented it should be feasible. With navigation technology, these procedures may be performed with less imaging, i.e. less X-ray exposure in total. We believe that such a system could be valuable for numerous clinical applications, such as endovascular therapy and navigated bronchoscopy, but also for experimental surgery in e.g. laparoscopy, where the C-arm is used for verification and comparison purposes, and spinal surgery. We will continue to develop EM based tracking integrated in minimal access therapy applications and follow up this study with clinical experiments to demonstrate the potential value in combination with a C-arm.

## Acknowledgements

We would very much like to thank Emmanuel Wilson at the Georgetown University Medical Center for providing us with the phantom and MATLAB code used in this work.

This study was supported by: SINTEF (Trondheim, Norway), The Ministry of Health and Social Affairs of Norway, through the National Centre for 3D Ultrasound in Surgery (Trondheim, Norway), project 196726/V50 eMIT (Enhanced minimally invasive therapy, FRIMED program), and the Operating Room of the Future project at St. Olavs Hospital (Trondheim, Norway).

## References

1. Abi-Jaoudeh, N, Glossop, N, Dake, M, Pritchard, WF, Chiesa, A, Dreher, MR, Tang, T, Karanian, JW, and Wood, BJ. Electromagnetic Navigation for Thoracic Aortic Stent-graft Deployment: A Pilot Study in Swine. *Journal of Vascular and Interventional Radiology* 2010;21:888–895.
2. Manstad-Hulaas, F, Ommedal, S, Tangen, GA, Aadahl, P, and Hernes, TAN. Side-Branched AAA Stent Graft Insertion Using Navigation Technology: A Phantom Study. *European Surgical Research* 2007;39:364–371.
3. Wilson, K, Guiraudon, G, Jones, DL, and Peters, TM. Mapping of Cardiac Electrophysiology Onto a Dynamic Patient-Specific Heart Model. *IEEE Transactions on Medical Imaging* 2009;28:1870–1880.
4. Eberhardt, R, Anantham, D, Herth, F, Feller-Kopman, D, and Ernst, A. Electromagnetic Navigation Diagnostic Bronchoscopy in Peripheral Lung Lesions. *Chest* 2007;131:1800–1805.
5. Makris, D, Scherpereel, A, Leroy, S, Bouchindhomme, B, Faivre, JB, Remy, J, Ramon, P, and Marquette, CH. Electromagnetic Navigation Diagnostic Bronchoscopy for Small Peripheral Lung Lesions. *The European Respiratory Journal* 2007;29:1187–1192.
6. Hildebrand, P, Kleemann, M, Schlichting, S, Besirevic, A, Roblick, U, Bürk, C, and Bruch, HP. Prototype of an Online Navigation System for Laparoscopic Radiofrequency Ablation of Liver Tumors. In: *World Congress on Medical Physics and Biomedical Engineering*. Ed. by Dössel, O and Schlegel, WC. Vol. 25/6. IFMBE Proceedings. Springer Berlin Heidelberg, 2009:352–354.
7. Solberg, OV, Langø, T, Tangen, GA, Mårvik, R, Ystgaard, B, Rethy, A, and Hernes, TAN. Navigated Ultrasound in Laparoscopic Surgery. *Minimally Invasive Therapy & Allied Technologies* 2009;18:36–53.

8. Chung, AJ, Edwards, PJ, Deligianni, F, and Yang, GZ. Freehand Cocalibration of Optical and Electromagnetic Trackers for Navigated Bronchoscopy. In: *Medical Imaging and Augmented Reality*. Ed. by Yang, GZ and Jiang, T. Vol. 3150. Lecture Notes in Computer Science. Springer Berlin/Heidelberg, 2004:320–328.
9. Kindratenko, V. A Survey of Electromagnetic Position Tracker Calibration Techniques. *Virtual Reality* 2000;5:169–182.
10. Nakamoto, M, Nakada, K, Sato, Y, Konishi, K, Hashizume, M, and Tamura, S. Intraoperative Magnetic Tracker Calibration Using a Magneto-Optic Hybrid Tracker for 3-D Ultrasound-Based Navigation in Laparoscopic Surgery. *IEEE Transactions on Medical Imaging* 2008;27:255–270.
11. Wilson, E, Yaniv, Z, Zhang, H, Nafis, C, Shen, E, Shechter, G, Wiles, AD, Peters, T, Lindisch, D, and Cleary, K. A Hardware and Software Protocol for the Evaluation of Electromagnetic Tracker Accuracy in the Clinical Environment: a Multi-Center Study. In: *Medical Imaging 2007: Visualization and Image-Guided Procedures*. Ed. by Kevin, RC and Michael, IM. Vol. 6509. Proceedings of SPIE 1. SPIE, 2007:65092T.
12. Yaniv, Z, Wilson, E, Lindisch, D, and Cleary, K. Electromagnetic Tracking in the Clinical Environment. *Medical Physics* 2009;36:876–892.
13. Horn, BKP. Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America A* 1987;4:629–642.
14. Azizian, M and Patel, R. Data fusion for catheter tracking using Kalman filtering: applications in robot-assisted catheter insertion. In: *Medical Imaging 2011: Visualization, Image-Guided Procedures, and Modeling*. Ed. by Wong, KH and III, DRH. Vol. 7964. 1. SPIE, 2011:796413.

15. Wolf, R, Duchateau, J, Cinquin, P, and Voros, S. 3D Tracking of Laparoscopic Instruments Using Statistical and Geometric Modeling. In: *Medical Image Computing and Computer-Assisted Intervention — MICCAI 2011*. Ed. by Fichtinger, G, Martel, A, and Peters, T. Vol. 6891. Lecture Notes in Computer Science. Springer Berlin/Heidelberg, 2011:203–210.
16. Luo, X, Kitasaka, T, and Mori, K. Bronchoscopy Navigation beyond Electromagnetic Tracking Systems: A Novel Bronchoscope Tracking Prototype. In: *Medical Image Computing and Computer-Assisted Intervention — MICCAI 2011*. Ed. by Fichtinger, G, Martel, A, and Peters, T. Vol. 6891. Lecture Notes in Computer Science. Springer Berlin/Heidelberg, 2011:194–202.
17. Feuerstein, M, Reichl, T, Vogel, J, Traub, J, and Navab, N. Magneto-Optical Tracking of Flexible Laparoscopic Ultrasound: Model-Based Online Detection and Correction of Magnetic Tracking Errors. *IEEE Transactions on Medical Imaging* 2009;28:951–967.

## Paper C

# Versatile robotic probe calibration for position tracking in ultrasound imaging

Lars Eirik Bø<sup>1,2,3</sup>, Erlend Fagertun Hofstad<sup>1</sup>, Frank Lindseth<sup>1,2</sup> and Toril A. N. Hernes<sup>2</sup>

<sup>1</sup>SINTEF Technology and Society, Trondheim, Norway

<sup>2</sup>Norwegian University of Science and Technology, Trondheim, Norway

<sup>3</sup>The Central Norway Regional Health Authority, Trondheim, Norway

*Physics in Medicine and Biology*, vol. 60, no. 9, pp. 3499–3513, 2015

**Abstract**

Within the field of ultrasound-guided procedures, there are today a number of methods for ultrasound probe calibration. While these methods are usually developed for a specific probe, they are in principle easily adapted to other probes. In practice, however, the adaptation often proves tedious, and this is impractical in a research setting, where new probes are tested regularly. Therefore, we developed a method which can be applied to a large variety of probes without adaptation. The method used a robot arm to move a plastic sphere submerged in water through the ultrasound image plane, providing a slow and precise movement. The sphere was then segmented from the recorded ultrasound images using a MATLAB programme, and the calibration matrix was computed based on this segmentation in combination with tracking information. The method was tested on three very different probes demonstrating both great versatility and high accuracy.

**C.1 Introduction**

When using ultrasound in image-guided therapy, it is essential to know the position of the ultrasound images in space; i.e. where in space are the objects that appear in the images located. This is necessary both to create three-dimensional volumes from the images, and subsequently to navigate on those volumes. To determine its position, the ultrasound probe is usually equipped with a position sensor whose position is measured in real time. The most common types of position sensors are optical sensors consisting of multiple infrared light sources which are tracked by infrared cameras, and electromagnetic sensors consisting of small coils whose positions can be determined by setting up a controlled, varying magnetic field and measuring the voltages induced in the coils. Since the position sensor is fixed to the probe, the position of the ultrasound images relative to the sensor is constant and can be found in a process referred to as probe calibration. By combining the result of this calibration with the real-time measurements of the sensor, the position of each ultrasound image can be determined.

Substantial research has been done in the field of probe calibration the last two decades, and this has resulted in a number of fast, automatic and accurate calibration methods. Thorough reviews of this work are given in Mercier et al. [1] and Hsu et al. [2]. While these methods are usually developed for a specific probe, they can in principle easily be adapted to other probes. In practice, however, this is not so. Ultrasound probes are more and more specialized and tailored to an increasing number of applications. The differences between probes are large with respect to properties such as shape, field of view, resolution, contrast and noise. As a result, the adaptation of calibration methods often proves tedious, requiring modifications to several central components, especially phantoms and image processing algorithms. A side-looking probe may, for instance, not be able to get to the surface of a phantom made for an ordinary forward-looking probe; a probe with a small field of view may not be able to image all of the wires in a typical wire phantom; and a high-resolution probe imaging a bead phantom may produce quite different reverberation artefacts than an ordinary probe.

This tedious adaptation is impractical in a research setting, where ultrasound guidance regularly is applied to new medical fields, all of which have their own more or less specialized ultrasound equipment. Especially for initial trials and small-scale feasibility studies, the need for a customized phantom is an obstacle. Therefore, we have developed a method which is not particularly fast, nor completely automatic, but which is accurate and robust and, most importantly, can be applied to a large variety of probes without adaptation.

## Theoretical background

To fully understand what a probe calibration does, it is helpful to know the various coordinate systems that are involved. These are illustrated in Figure C.1. In the following,  $M_{r \leftarrow s}$  denotes a 4-by-4 transformation matrix that, through multiplication, converts the coordinate vector  $p_s$  of a given point in coordinate system  $s$  into the coordinate vector  $p_r$

of the same point in coordinate system  $r$ , i.e.

$$p_r = M_{r \leftarrow s} \cdot p_s.$$

First, a position sensor is usually attached to the patient or the operating table to act as a fixed reference, and this sensor defines the reference coordinate system  $r$ . This means that the positions of all other objects, such as instruments, images and the patient itself, are given in this system. Then, the position sensor attached to the ultrasound probe has its own coordinate system  $s$ . The tracking system continuously measures the positions of these two first coordinate systems and calculates the spatial relationship between them, i.e. the rigid transformation  $M_{r \leftarrow s}$  converting position sensor coordinates into reference coordinates. Finally, the two-dimensional ultrasound images are defined in a third coordinate system  $i$ . The spatial relationship between this system and the one defined by the sensor on the probe is fixed, and the goal of the probe calibration is to find this relationship in the form of a rigid transformation  $M_{s \leftarrow i}$ .

The basic principle of most probe calibration methods is to image some object whose appearance in the ultrasound images makes it easy to accurately measure its position within the image. This is referred to as an imaging phantom, and it usually includes one or more imaging targets, which are features that are easily identified in the images. To enable the ultrasound imaging, the phantom must be built in or submersed in an acoustic coupling media such as water. It is also equipped with a reference sensor  $r$ , and the position  $p_r$  of the imaging target relative to this reference is measured, usually using some kind of tracked pointer. As ultrasound images of the phantom are recorded, the positions of both the phantom and the ultrasound probe are captured by the tracking system and saved together with the images. The imaging target's position  $p_i$  in the coordinate system  $i$  is also extracted from the ultrasound images, and the calibration matrix can then be found by minimizing the distance between the target positions in the

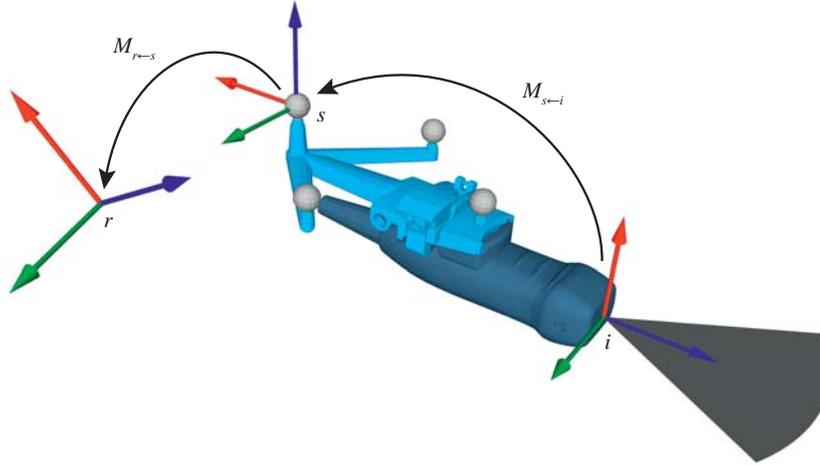


Figure C.1: The various coordinate systems involved in probe calibration: the reference system  $r$ , the position sensor system  $s$  and the image system  $i$ . The transformation  $M_{r \leftarrow s}$ , which changes as the probe is moved, is continuously measured by the tracking system. The transformation  $M_{s \leftarrow i}$ , on the other hand, is fixed, and it is the goal of the probe calibration to find this.

two coordinate systems, i.e. as

$$M_{s \leftarrow i} = \arg \min_{M'_{s \leftarrow i}} \sum_j^n \|M'_{s \leftarrow i} \cdot p_i^j - M_{r \leftarrow s}^{-1} \cdot p_r^j\|.$$

A lot of different phantoms have been proposed, but most of them belong to one of the following five groups [1]:

**Point target phantoms:** Here, a single, small object serves as imaging target, e.g. a small bead [3–6], a wire cross [7–9] or the tip of a stylus [10–12]. These phantoms are simple and thus easy to

make, but they only provide one datapoint per ultrasound image and must thus be imaged many times for each calibration. The small size of the bead or wire also means that its appearance in the ultrasound image varies a lot between probes; what appears as a focused and distinct spot in one image, may appear blurred and noisy when imaged with a different system. A target that is suitable for one probe may therefore not be suitable for another.

**Multiple point targets phantoms:** These are similar to the point target phantoms, except they include several small objects that are to be imaged either successively [13, 14] or simultaneously [4, 15]. This limits the amount of recordings that is required for each calibration. In the last case, however, there is the added problem of properly aligning the ultrasound plane with the targets in order to image them all at the same time.

**Z-fiducial phantoms:** These phantoms include one or more Z-fiducials [4, 16, 17], which each consists of three thin wires stretched between the walls of the phantom forming a Z in the axial plane. When imaged from above, the three wires appear as three bright points in the ultrasound image, and by measuring the relative distance between these points, the line of intersection between the ultrasound plane and the wires can be determined. With this approach, a large number of independent datapoints can be collected in one recording. However, again the appearance of the wire in the ultrasound image varies a lot between probes. In addition, the geometry of the Z-fiducials must be adapted to the size of the image plane.

**2-D shape alignment phantoms:** These phantoms include a membrane with an irregular, jagged edge [18, 19]. The ultrasound plane is aligned with the membrane so that this edge appears as a jagged line in the ultrasound image, and the corners of this line can then be located in the image. As with the multiple point targets phantoms, the proper alignment of the ultrasound plane with the membrane is a challenge.

**Wall phantoms:** Here, one or more plane surfaces, such as walls [20] or membranes [19, 21], are imaged, producing bright lines in the ultrasound image. The most basic versions just use the bottom of a water tank as an imaging target. The lines are easily identified in the ultrasound images, and the technique thus lends itself to automatic image processing. However, the visibility of the imaged surface in the images is very dependent on the angle between the probe and the surface. Reverberation artefacts caused by multiple reflections of the sound is also a challenge.

The references given here are only meant as examples. More exhaustive references and even more phantom variants can be found in Mercier et al. [1] and Hsu et al. [2].

## C.2 Materials and methods

Our method uses a point target phantom with a plastic sphere in a water tank as an imaging target. This is similar to the approach taken by Sauer et al. [14]. The setup includes an optical tracking system (Polaris Spectra, NDI, Waterloo, ON, Canada) that measures the positions of small, retroreflective plastic spheres with a diameter of 11.5 mm. We designed a calibration arm made mainly from plastic incorporating seven such spheres: one sphere was mounted at the tip of the arm to function as an imaging target, while the other six were mounted in a particular pattern on a plastic plate at the other end of the arm to function as a tracking reference. The arm is shown in Figure C.2(a). The position  $p_r$  of the centre of the imaging target relative to the reference was measured using the tracking system. The arm was then attached to a six-axis robot arm (UR5, Universal Robots, Odense, Denmark) mounted next to a water tank on a bench and positioned so that the end of the arm with the imaging target reached into the water. The robot had a repeatability of  $\pm 0.1$  mm. The ultrasound probe that was to be calibrated, which was also equipped with a position sensor, was positioned straight above the imaging target and rigidly attached to the bench. Finally, a computer running an in-house navigation sys-

tem (CustusX, SINTEF, Trondheim, Norway)[22] was connected both to the ultrasound scanner and to the tracking system. The complete setup is shown in Figure C.2(c).

The robot now moved the imaging target to a starting position just outside the ultrasound image plane. It then moved it slowly, at the speed of 1 mm/s, first straight through the image plane, and then in the opposite direction until it was back at the starting position again. This is illustrated in Figure C.3. While the target was moving, the ultrasound images produced by the ultrasound scanner were recorded by the navigation system. For each image  $k$ , the system also recorded the position  $M_{r \leftarrow s}^k$  of the ultrasound probe relative to the reference sensor on the calibration arm. The imaging target was then moved to a new starting position. This process was repeated nine times so that the imaging target was passed through the ultrasound image plane altogether 18 times at nine different positions. The distance between the positions were chosen so that they were evenly distributed throughout the plane.

This process was repeated nine times so that the imaging target was passed through the ultrasound image plane altogether 18 times at nine different positions evenly distributed throughout the plane.

The next step was to find the position of the imaging target within each of these 18 recordings. In practice, this meant that for recording  $j$  we needed to find the image  $k_j$  that passed through the centre of the sphere, and then determine the image coordinates  $p_i^{k_j}$  of the sphere centre within this image. This is again illustrated in Figure C.3. To achieve this, the recorded images were read into the software MATLAB (Mathworks, Natick, MA, U.S.A.) and displayed. As shown in Figure C.4, the images changed as the imaging target moved through the image plane. However, since the target was spheric and moved with constant speed and direction, the appearance of the images was symmetric around the centre of the sphere. By drawing a rectangular box around the sphere, summing the intensities of all the pixels within the box for each ultrasound image and plotting this sum against the image number, the result was therefore a graph that was also sym-

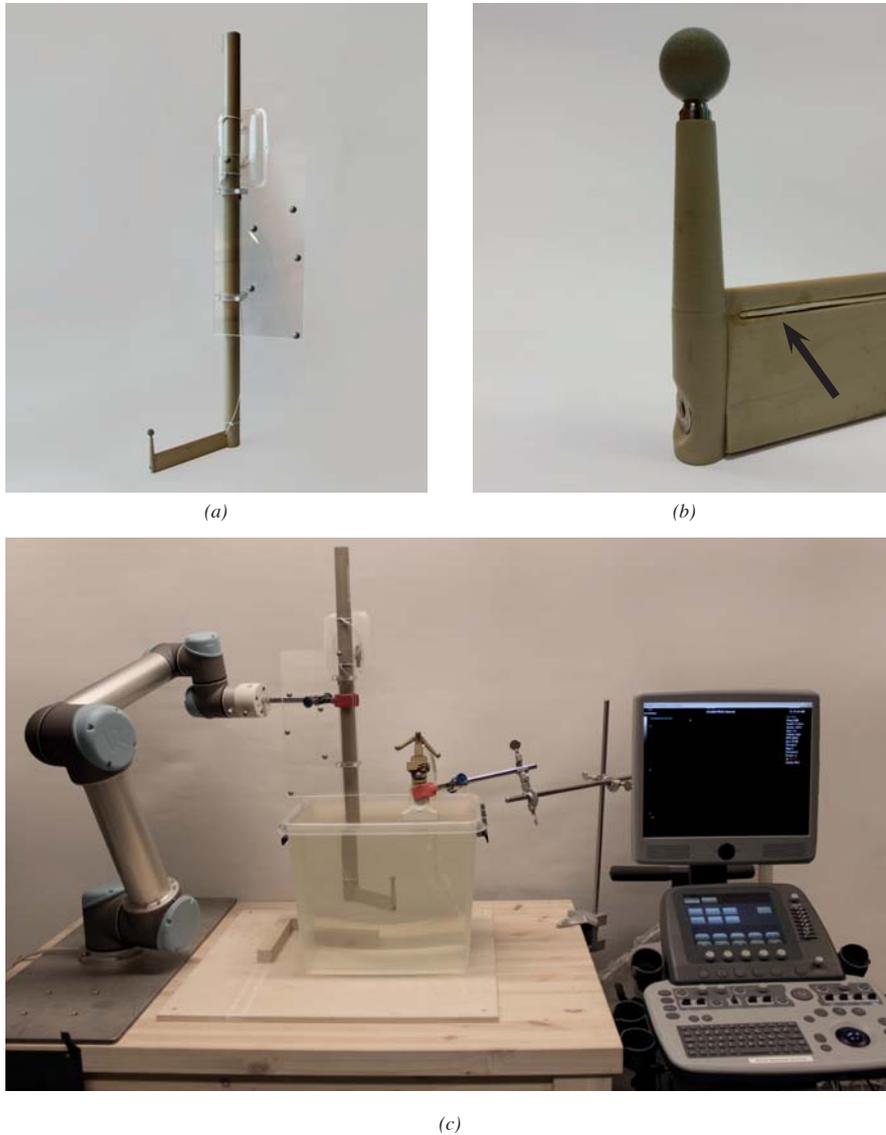


Figure C.2: (a) The calibration arm, (b) a close up of the imaging target and the electromagnetic reference sensor (white cable indicated by arrow) and (c) the complete setup with the robot arm holding the arm in the water tank and the ultrasound probe positioned directly above the imaging target.

metric. This can be seen in the bottom panel of Figure C.4. Using this graph for support, it was easy to flip through the displayed images and find the one going through the centre of the imaging target. The resulting image  $k_j$  only showed the circular surface of the plastic sphere, but knowing its diameter and the pixel size of the ultrasound image, a circle of the same size as the sphere was drawn in the image. By zooming in, this circle could manually be moved so that its circumference corresponded with the surface of the imaged sphere, and its centre thus corresponded with the centre of the sphere as seen in Figure C.5. This was repeated for each of the 18 recordings of the imaging target.

Now, the image  $k_j$  showing the centre of the imaging target and the image coordinates  $p_i^{k_j}$  of the centre within this image had been found for each of the 18 recordings, i.e. for  $j = 1, \dots, 18$ . The transformations  $M_{r \leftarrow s}^{k_j}$  representing the position and orientation of the ultrasound probe corresponding to each of these images were then extracted from the navigation system.

The positions  $M_{r \leftarrow s}^{k_j}$  of the ultrasound probe corresponding to each of these images were then extracted from the navigation system. Combining this information with the previously measured coordinates  $p_r$  of the imaging target relative to the reference on the calibration arm, its position relative to the position sensor on the ultrasound probe was found as

$$p_s^{k_j} = M_{s \leftarrow r}^{k_j} \cdot p_r = (M_{r \leftarrow s}^{k_j})^{-1} \cdot p_r.$$

The result was two sets of coordinates for each of the 18 selected images, both of which described the position of the imaging target: one in the image coordinate system and one in the probe's coordinate system. The calibration matrix was then given as

$$M_{s \leftarrow i} = \arg \min_{M'_{s \leftarrow i}} \sum_{j=1}^{18} \|M'_{s \leftarrow i} \cdot p_i^{k_j} - p_s^{k_j}\|,$$

and this was calculated using a closed-form method given by Horn [23].

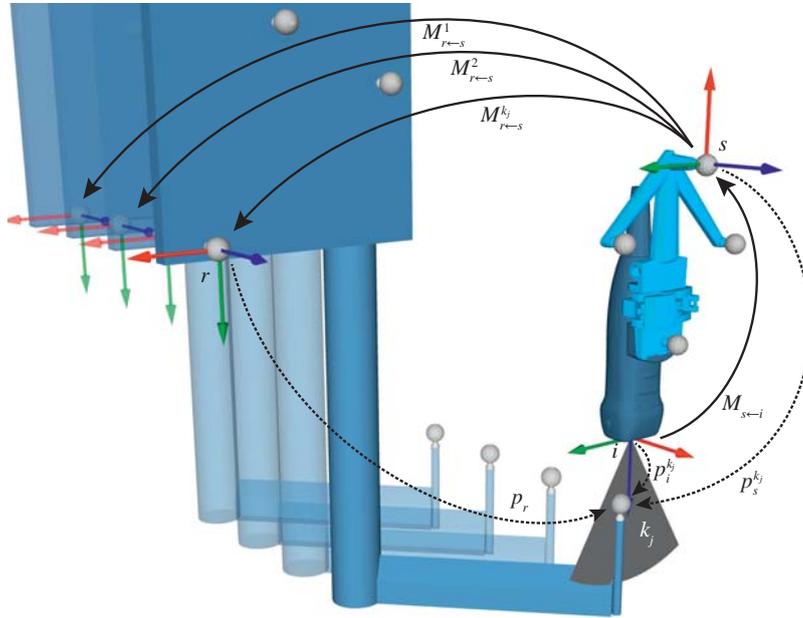


Figure C.3: The imaging target being moved slowly through the ultrasound image plane. Image  $k_j$  is the ultrasound image recorded exactly as the image plane cuts through the centre of the sphere,  $M_{r \leftarrow s}$  and  $M_{s \leftarrow i}$  are 4-by-4 transformation matrices, here illustrated with solid lines, and  $p_r$ ,  $p_i^{k_j}$  and  $p_s^{k_j}$  are coordinate vectors, here illustrated with dotted lines.

### Adaption to other tracking systems

Our navigation system interfaces not only with the optical tracking system, but also with an electromagnetic tracking system (Aurora, NDI, Waterloo, ON, Canada). This is used in settings where either it is hard to achieve a clear line of sight, or the optical position sensors are too bulky to be integrated properly with the instruments in question. This is typically the case for small or flexible instruments. To enable the calibration of ultrasound probes equipped with such electromagnetic position sensors, we mounted an electromagnetic reference sensor on the calibration arm close to the imaging target. This

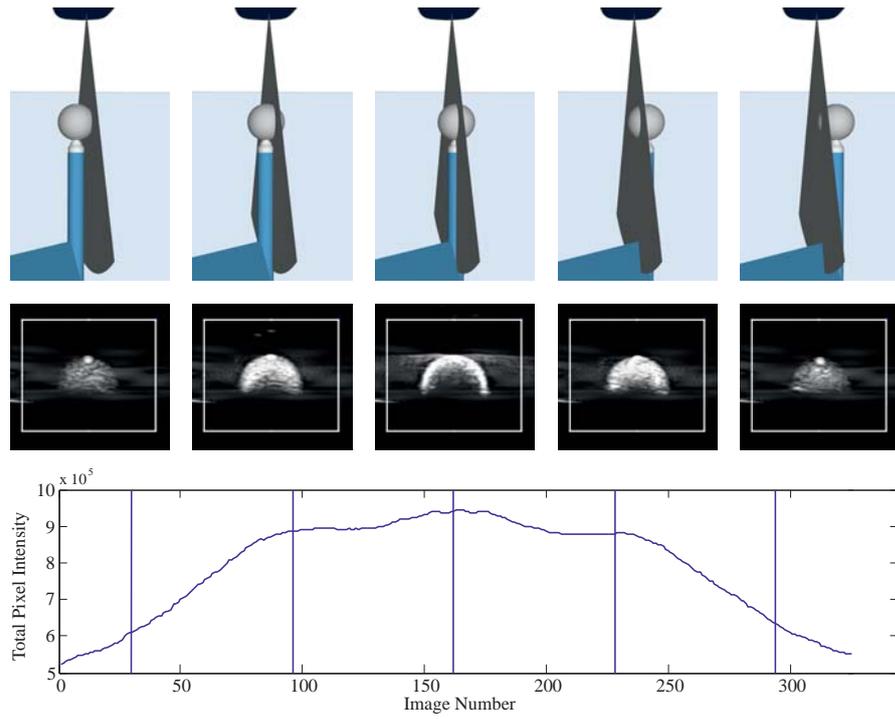


Figure C.4: The imaging target at five different positions on its way through the ultrasound image plane (top panel) and the corresponding ultrasound images (middle panel). The bottom panel shows the sum of the pixel intensities within the white box for each ultrasound image plotted against the image number, and the vertical lines indicate the above images.

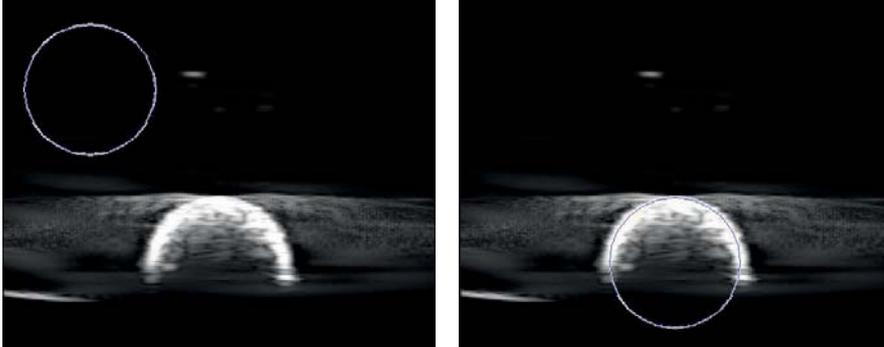


Figure C.5: A typical ultrasound image cutting through the centre of the imaging target with a circle of the same size as the sphere drawn on top. The left image shows the circle at its initial position, and the right image shows the same circle after manual alignment with the sphere surface.

can be seen in Figure C.2(b). Since this tracking system could not measure the position of the target directly, we instead found the spatial relationship between the optical and the electromagnetic reference sensors. This was done by rigidly attaching the calibration arm to a cubic plastic box with each side measuring 25 cm. On each of the four side walls there were drilled four small holes 16 cm apart. The position of each of these 16 holes was then measured two times: first using an optically tracked pointer, and then using an electromagnetically tracked pointer. Finally, the rigid transformation minimizing the distance between these two point sets was found, again using the method by Horn [23]. The position of the imaging target relative to the electromagnetic reference sensor was then found by applying the resulting transformation to the optically measured position.

## Evaluation

To evaluate the calibration method, we used the ultrasound scanner SonixMDP (Analogic, Boston, MA, USA) and chose three ultrasound probes that differed a lot with respect to both shape, size and image

resolution. The probes are listed in Table C.1 and shown in Figure C.6. One of them was equipped with an optical position sensor, and the other two were equipped with electromagnetic position sensors. Each probe was calibrated five times, producing altogether 15 different calibrations. Since each calibration was based on 18 recordings of the imaging target, a total of 270 recordings were made.

Table C.1: The probes used to evaluate the calibration method.

| Probe  | Application     | Depth<br>(mm) | Frequency<br>(MHz) | Tracking        |
|--------|-----------------|---------------|--------------------|-----------------|
| C5-2   | Abdomen         | 50–300        | 2–4                | Optical         |
| L13-7  | Pituitary gland | 20–90         | 6.6–10             | Electromagnetic |
| LAP9-4 | Laparoscopy     | 20–90         | 5–9                | Electromagnetic |

To provide a measure of the overall accuracy that was independent of the calibration setup, a precisely engineered accuracy phantom was also used. This phantom consisted of a water tank with two nylon wires crossing each other at its centre. The tank was equipped with a tracking reference, and the position  $p_r$  of the wire cross relative to the reference was measured using a mechanical stylus with an accuracy of 0.01 mm. The tank is shown in Figure C.7(a). The wire cross was imaged four times for each calibration, resulting in 60 accuracy recordings altogether.

Based on these calibrations and recordings, we calculated four different quality measures [4]: the leave-one-out cross-validation error (LooCvE), the point reconstruction accuracy (PRAc), the calibration reproducibility (CR) and the three-dimensional navigation accuracy (3-D NAc).

As described previously, each calibration was based on 18 recordings of the imaging target from which the position of the imaging target both in the image coordinate system and in the probe’s coordinate system was found. To determine the LooCvE, we left out the  $j$ th record-

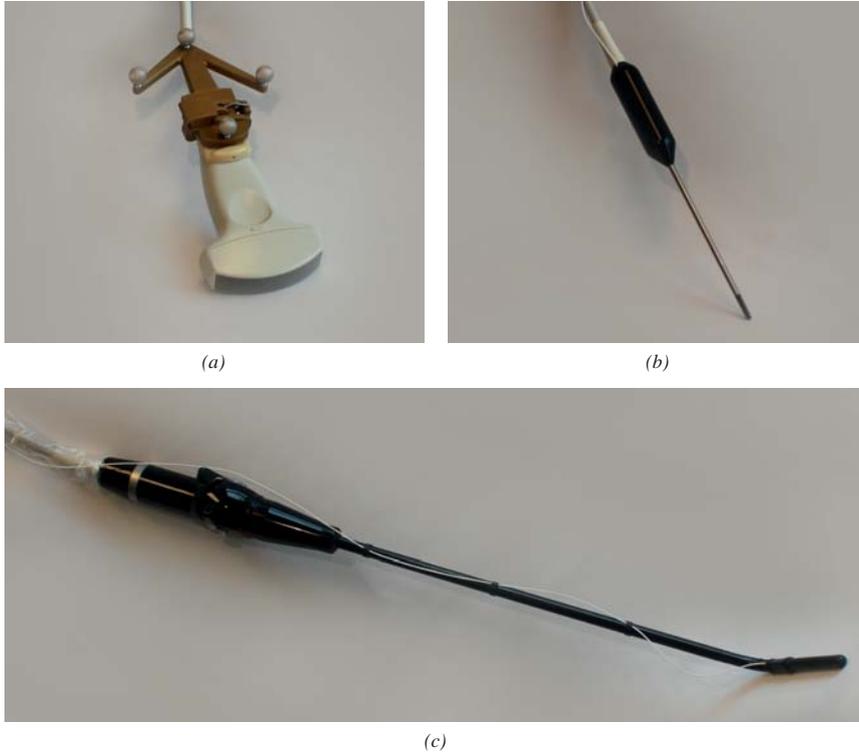


Figure C.6: The probes used to evaluate the calibration method: (a) C5-2 equipped with an optical position sensor, and (b) L13-7 and (c) LAP9-4 equipped with small, electromagnetic position sensors.

ing and calculated a new calibration  $M_{s \leftarrow i}^{-j}$  based on the remaining 17. The resulting calibration was applied to the image coordinates  $p_i^{k_j}$  of the sphere extracted from the chosen recording, transforming them to the probe's coordinate system. A partial error was then calculated as the euclidean distance between these transformed coordinates and the coordinates  $p_s^{k_j}$  of the sphere measured by the tracking system. This process was repeated for all 18 recordings, and the LooCvE was calculated as the average taken over the resulting 18 partial errors, i.e.

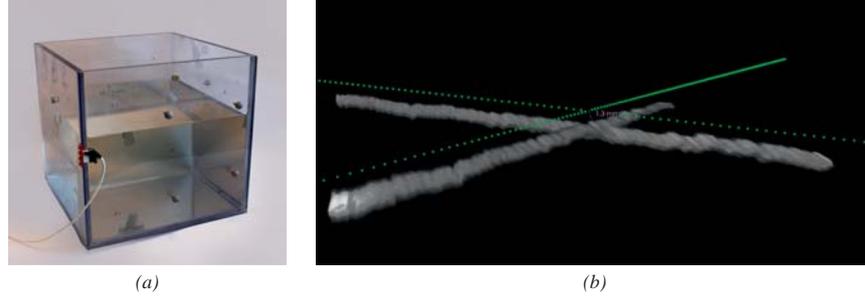


Figure C.7: (a) The accuracy phantom used to measure the three-dimensional navigation accuracy (3D-NAc), and (b) a three-dimensional visualization of the mechanically measured wire cross (in green) and the ultrasound volume for one of the 60 accuracy recordings.

as

$$\Delta_s^{\text{LooCvE}} = \frac{1}{18} \sum_{j=1}^{18} \|M_{s \leftarrow i}^{-j} \cdot p_i^{k_j} - p_s^{k_j}\|.$$

The PRAc is similar to the LooCvE. The difference is that rather than using the same recordings for evaluation that are used to calculate the calibration, we use a separate set of recordings; since we had produced altogether five calibrations for each probe, we used the 72 recordings originally used to produce the other four calibrations. The given calibration was applied to the image coordinates of the sphere extracted from each of these 72 recordings, and again the partial error was calculated as the euclidean distance between these transformed coordinates and the sphere coordinates measured by the tracking system. The PRAc for the  $m$ th calibration was then found as the average of these 72 partial errors, i.e. as

$$\Delta_s^{\text{PRAc}} = \frac{1}{72} \sum_{n \neq m} \sum_{j=1}^{18} \|M_{s \leftarrow i}^m \cdot p_i^{n,k_j} - p_s^{n,k_j}\|,$$

where  $M_{s \leftarrow i}^m$  is the  $m$ th calibration matrix and the coordinates  $p_s^{n,k_j}$  and  $p_i^{n,k_j}$  are extracted from the  $j$ th recording of the  $n$ th calibration.

The CR is a measure of the precision of the calibration method. Each of the five calibrations were applied to a virtual image point  $p_i^{\text{virtual}}$ , which in this case was chosen to be the lower right corner of the recorded images in accordance with Lindseth et al. [4]. This resulted in five points in the probe's coordinate system. The CR for the  $m$ th calibration was then calculated as the mean of the euclidean distances from the point transformed by this calibration to the points transformed by the other four calibrations, i.e. as

$$\Delta_s^{\text{CR}} = \frac{1}{4} \sum_{n \neq m} \|M_{s \leftarrow i}^m \cdot p_i^{\text{virtual}} - M_{s \leftarrow i}^n \cdot p_i^{\text{virtual}}\|.$$

The last quality measure, the 3-D NAc, is based on the ultrasound recordings of the accuracy phantom, and it is thus the only one that is independent of the calibration setup. For each recording, the navigation system created a three-dimensional ultrasound volume based on the given probe calibration, the recorded images and the corresponding tracking data. This was done using the reconstruction algorithm Pixel Nearest Neighbour [24]. The two wires were then segmented from the ultrasound volume using a fast, automatic method for centre line extraction [25]. Finally, these centre lines were registered to a model of the wire cross that was based on the mechanical measurements of the wires. This was done using the Iterative Closest Point algorithm [26]. The translational part  $t_r$  of the resulting registration was used as a measure of the distance between the wire cross in the ultrasound volume and the mechanically measured wire cross. The 3-D NAc was then found as the mean of this distance taken over all four volumes, i.e. as

$$\Delta_s^{3\text{-D NAc}} = \frac{1}{4} \sum_{l=1}^4 \|t_r^l\|,$$

where  $t_r^l$  is the translational part of the registration matrix based on the  $l$ th recorded ultrasound volume. Figure C.7(b) shows both the mechanically measured wire cross and the ultrasound volume for one of the 60 accuracy recordings.

The 3-D NAc is in fact a measure of the overall accuracy of the navigation system when navigating on a reconstructed three-dimensional ultrasound volume. Thus, it includes multiple error sources in addition to the probe calibration, such as sensor attachment repeatability, position sensor tracking, synchronisation between position data and images and reconstruction algorithm [27]. This should therefore be regarded as an upper bound on the accuracy of the probe calibration.

### C.3 Results

A total of 15 different calibrations were performed. Each calibration took approximately 60 minutes, out of which setting up the equipment, acquiring data and processing data took around 20 minutes each. The time spent on setting up the equipment is of course reduced when multiple calibrations are performed at the same time.

The quality measures for the three different probes that were calibrated are shown in Table C.2. Both accuracy and precision are good with PRAc below 1.07 mm and CR below 0.89 mm for all 15 calibrations. These results are further supported by the 3-D NAc, which shows that the overall accuracy of the system is below 1.45 mm for all probes and calibrations, and as low as 1.09 mm for the C5-2 probe with optical tracking.

### C.4 Discussion

The motivation for this work was the need for a calibration method which could be applied to any kind of ultrasound probe, regardless of shape, field of view, resolution, contrast or noise level, without any need for adaptation. This is an important requirement in a research setting where new and specialized probes are tested on a regular basis. As described in the introduction, a multitude of calibration methods and phantoms have already been presented, but after having tried a number of them during the last 15 years, we still had not found one

Table C.2: Leave-one-out cross-validation error (LooCvE), point reconstruction accuracy (PRAc), calibration reproducibility (CR) and three-dimensional navigation accuracy (3-D NAc) for three different probes. Each probe was calibrated five times, and the table shows the mean, the standard deviation (SD), the minimum value and the maximum value among these five calibrations. All numbers are given in mm.

| Probe                       |         | LooCvE | PRAc | CR   | 3-D NAc |
|-----------------------------|---------|--------|------|------|---------|
| C5-2<br>(optical)           | Mean    | 0.73   | 0.78 | 0.36 | 0.95    |
|                             | SD      | 0.07   | 0.04 | 0.07 | 0.13    |
|                             | Minimum | 0.65   | 0.73 | 0.31 | 0.74    |
|                             | Maximum | 0.80   | 0.82 | 0.49 | 1.09    |
| L13-7<br>(electromagnetic)  | Mean    | 0.46   | 0.96 | 0.61 | 0.98    |
|                             | SD      | 0.06   | 0.10 | 0.17 | 0.28    |
|                             | Minimum | 0.38   | 0.87 | 0.49 | 0.80    |
|                             | Maximum | 0.52   | 1.07 | 0.89 | 1.45    |
| LAP9-4<br>(electromagnetic) | Mean    | 0.61   | 0.76 | 0.38 | 1.13    |
|                             | SD      | 0.03   | 0.05 | 0.12 | 0.18    |
|                             | Minimum | 0.58   | 0.72 | 0.31 | 0.89    |
|                             | Maximum | 0.67   | 0.84 | 0.60 | 1.35    |

that fulfilled this requirement. The method presented here is therefore not very novel, but it is a practical solution adapted to our needs.

We chose a point-based method mainly due to its simplicity: a point target is easy to get at with the ultrasound probe regardless of its size and shape, and since it is spherically symmetric it can be imaged from any angle. This means that the angle of the ultrasound image plane does not have to be aligned with the calibration arm, which greatly simplifies the setup. We found that the imaging target should be relatively large compared to the resolution of the scanner. While small bead-like targets tend to be smeared out in the ultrasound images, the larger plastic spheres appear as a well-defined, semicircu-

lar shape which is easily segmented from the images. The accuracy of the manual segmentation that we propose is hard to determine, as the ground truth is not known. However, the semicircular shape is usually only between one and two millimetres thick, and in these cases the segmentation can probably be assumed to have submillimetre accuracy.

The use of a robot for moving the imaging target has both advantages and disadvantages. It facilitates automation of the calibration process, and it performs the movements very accurately and repeatably. Slow, steady motions of the target produces better image quality, which again makes the subsequent segmentation easier and more accurate. The main advantage, however, is that it solves the problem of aligning the centre of the target with the ultrasound plane, which is one of the major problems of point-based methods [11]. Since the motion of the imaging target is performed at a constant speed and follows a linear trajectory, the acquired ultrasound images are completely symmetric around the centre of the target. It is therefore straight forward to identify the image corresponding to this centre. Moreover, due to the low speed of the motion, the error made by missing the centre with a few frames is negligible.

One problem when operating with a research system in combination with a large number of different ultrasound scanners, is the time lag between images collected from the scanner and position data collected from the tracking system. If the lag is constant, which it often is, this can be measured and compensated for. However, depending on the hardware being used, the lag may vary, making accurate compensation difficult. The calibration method presented here goes a long way towards eliminating this problem during calibration. This is because for each position where the imaging target was passed through the ultrasound image plane, it was moved both back and forth with a constant speed along the same linear trajectory. Assuming the time lag of the system remained constant during this movement, which lasted only 40 s, the position errors introduced by the lag would be exactly opposite for the two movements and thus cancel each other out.

It is challenging to measure the exact accuracy of a probe calibration. The LooCvE and the PRAc are both based on data from the same

phantom and setup that is used for the calibration itself. They can therefore provide information about the consistency of the collected data, but systematic errors, e.g. caused by inaccurate characterization of the calibration phantom (in our case this is the calibration arm), will not be detected. The 3-D NAc, on the other hand, is measured using a separate phantom, and it will therefore also reveal systematic errors in the calibration method. However, as previously mentioned, this is a measure of the overall accuracy of the navigation system, which includes several other error sources. A high 3D-NAc (i.e. high value, not high accuracy) does therefore not necessarily mean that the calibration accuracy is poor, but a low 3D-NAc (i.e. low number, not low accuracy) means that the the calibration accuracy is good.

We achieved a 3-D NAc below 1.45 mm for all probes, which is sufficient for most clinical uses. This also means that the calibration accuracy is sufficient. The results for PRAc were similar to those presented by Lindseth et al. [4] (see their Table 6), and slightly better than those presented by Hsu et al. [2] (see their Table 1). Lindseth et al. reported considerably higher maximum values, but this may be due to the fact that their results were based on 15 calibrations for each probe, which is three times as many as ours. The CR numbers were similar for all three studies. It is, however, important to note that the variation between different probes and different tracking systems often is larger than the variation between different calibration methods when it comes to accuracy. While both of the cited studies used fairly standard probes and optical tracking systems, two of the probes used in this study (the L13-7 and the LAP9-4) had a shape which would make them difficult to calibrate with most other calibration methods. For these probes we also used an electromagnetic tracking system which is much more vulnerable to disturbances from the surroundings than the optical systems. The fact that we achieve comparable results also with these probes is a testimony to its high performance.

The main disadvantage of this method is that the robot is both expensive and space-demanding. However, industrial robots are becoming both smaller, cheaper and more easily programmable, with prices starting around €15,000, or about half of that of a traditional

industrial robot. This may still be somewhat expensive if probe calibration is the only task that is to be performed, but a robotic arm like the UR5 is a very flexible tool which can be used to automate a wide range of laboratory tasks. Together, it may justify the investment. One should also note that the robot may be replaced by any mechanical device capable of providing a slow, linear movement at a relatively constant speed. This would, however, reduce the flexibility of the method and also the potential of automation.

Another disadvantage is that both data acquisition and data processing is relatively time consuming compared to other methods. However, we found that our method provides a reasonable compromise between time, flexibility and reliability. It also has a great potential for further automation. By integrating the robot with the navigation system, a simple program could perform the entire data acquisition process without any manual interaction. The data processing, on the other hand, is harder to automate as different probes can produce very different images of the same object. In our case, the image of the plastic sphere can vary from a thin line with a clear, semicircular shape, to a fuzzy dot which is somewhat flatter on one side than on the other. Still, tools could be made to aid the segmentation of the sphere from the images, e.g. by automatic symmetry detection, and this could further speed up the process.

## **C.5 Conclusion**

The proposed probe calibration method can be used to calibrate a wide range of different probes without any adaptation and with high accuracy and repeatability. It is thus especially suitable in a research setting where new and specialized probes are tested on a regular basis. Though the method involves some manual steps today, it has the potential to be made fully automatic.

## Acknowledgements

The setup presented in this paper was based upon a similar probe calibration setup developed by SonoWand AS in Trondheim, Norway. The authors would like to thank SonoWand for sharing their knowledge with us.

## References

1. Mercier, L, Langø, T, Lindseth, F, and Collins, DL. A review of calibration techniques for freehand 3-D ultrasound systems. *Ultrasound in Medicine & Biology* 2005;31:449–471.
2. Hsu, PW, Prager, RW, Gee, AH, and Treece, GM. Freehand 3D Ultrasound Calibration: A Review. In: *Advanced Imaging in Biology and Medicine*. Ed. by Sensen, CW and Hallgrímsson, B. Springer, 2009:47–84.
3. State, A, Chen, DT, Tector, C, Brandt, A, Chen, H, Ohbuchi, R, Bajura, M, and Fuchs, H. Case Study: Observing a Volume Rendered Fetus within a Pregnant Patient. In: *IEEE Conference on Visualization, 1994*. Vol. 41. IEEE, 1994:364–368.
4. Lindseth, F, Tangen, GA, Langø, T, and Bang, J. Probe calibration for freehand 3-D ultrasound. *Ultrasound in Medicine & Biology* 2003;29:1607–1623.
5. Brendel, B, Winter, S, and Ermert, H. A simple and accurate calibration method for 3D freehand ultrasound. *Biomedizinische Technik* 2004;49:872–873.
6. Wang, J, Wang, T, Hu, L, Liao, H, Luan, K, and Sakuma, I. A Fully Automatic 2D Ultrasound Probe Calibration Method Using a Simple Phantom. In: *2011 4th International Conference on Biomedical Engineering and Informatics (BMEI)*. Ed. by Ding, Y, Peng, Y, Shi, R, Hao, K, and Wang, L. Vol. 4. IEEE, 2011:588–592.

7. Detmer, PR, Bashein, G, Hodges, T, Beach, KW, Filer, EP, Burns, DH, and Jr., DS. 3D ultrasonic image feature localization based on magnetic scanhead tracking: In vitro calibration and validation. *Ultrasound in Medicine & Biology* 1994;20:923–936.
8. Meairs, S, Beyer, J, and Hennerici, M. Reconstruction and Visualization of Irregularly Sampled Three- and Four-Dimensional Ultrasound Data for Cerebrovascular Applications. *Ultrasound in Medicine & Biology* 2000;26:263–272.
9. Yaniv, Z, Foroughi, P, Kang, HJ, and Boctor, E. Ultrasound Calibration Framework for the Image-Guided Surgery Toolkit (IGSTK). In: *Medical Imaging 2011: Visualization, Image-Guided Procedures, and Modeling*. Ed. by Wong, KH and III, DRH. Vol. 7964. Proceedings of SPIE. SPIE, 2011.
10. Khamene, A and Sauer, F. A Novel Phantom-Less Spatial and Temporal Ultrasound Calibration Method. In: *Medical Image Computing and Computer-Assisted Intervention — MICCAI 2005*. Ed. by Duncan, JS and Gerig, G. Vol. 3750. Lecture Notes in Computer Science. Springer, 2005:65–72.
11. Hsu, PW. Freehand three-dimensional ultrasound calibration. PhD thesis. 2007.
12. Lang, A, Parthasarathy, V, and Jai, A. Calibration of EM Sensors for Spatial Tracking of 3D Ultrasound Probes. In: *Data Acquisition Applications*. Ed. by Karakehayov, Z. New York, NY, USA: InTech, 2012.
13. Trobaugh, JW, Richard, WD, Smith, KR, and Bucholz, RD. Frameless stereotactic ultrasonography: method and applications. *Computerized Medical Imaging and Graphics* 1994;18:235–264.
14. Sauer, F, Khamene, A, Bascle, B, Schimmang, L, Wenzel, F, and Vogt, S. Augmented Reality Visualization of Ultrasound Images: System Description, Calibration, and Features. In: *IEEE and ACM International Symposium on Augmented Reality, 2001*. IEEE, 2001:30–39.

15. Leotta, DF. An efficient calibration method for freehand 3-D ultrasound imaging systems. *Ultrasound in Medicine & Biology* 2004;30:999–1008.
16. Comeau, RM, Sadikot, AF, Fenster, A, and Peters, TM. Intra-operative ultrasound for guidance and tissue shift correction in image-guided neurosurgery. *Medical Physics* 2000;27:787–800.
17. Chen, T, Abolmaesumi, P, Thurston, A, and Ellis, R. Automated 3D Freehand Ultrasound Calibration with Real-Time Accuracy Control. In: *Medical Image Computing and Computer-Assisted Intervention — MICCAI 2006*. Ed. by Larsen, R, Nielsen, M, and Sporring, J. Vol. 4190. Lecture Notes in Computer Science. Springer Berlin/Heidelberg, 2006:899–906.
18. Sato, Y, Nakamoto, M, Tamaki, Y, Sasama, T, Sakita, I, Nakajima, Y, Monden, M, and Tamura, S. Image Guidance of Breast Cancer Surgery Using 3-D Ultrasound Images and Augmented Reality Visualization. *IEEE Transactions on Medical Imaging* 1998;17.
19. Langø, T. Ultrasound Guided Surgery: Image Processing and Navigation. PhD thesis. 2000.
20. Prager, RW, Rohling, RN, Gee, AH, and Berman, LH. Rapid calibration for 3-D freehand ultrasound. *Ultrasound in Medicine & Biology* 1998;24:855–869.
21. Baumann, M, Daanen, V, Leroy, A, and Troccaz, J. 3-D Ultrasound Probe Calibration for Computer-Guided Diagnosis and Therapy. In: *Computer Vision Approaches to Medical Image Analysis: Second International ECCV Workshop – CVAMIA 2006*. Ed. by Beichel, RR and Sonka, M. Vol. 4241. Lecture Notes in Computer Science Volume. Springer, 2006:248–259.
22. Askeland, C, Lindseth, F, Solberg, OV, Bakeng, JBL, and Tangen, GA. CustusX: A Research Application for Image-Guided Therapy. *The MIDAS Journal* 2011;Systems and Architectures for Computer Assisted Interventions 2011:1–8.

23. Horn, BKP. Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America A* 1987;4:629–642.
24. Solberg, OV, Lindseth, F, Bø, LE, Muller, S, Bakeng, JBL, Tangen, GA, and Hernes, TAN. 3D ultrasound reconstruction algorithms from analog and digital data. *Ultrasonics* 2011;51:405–419.
25. Smistad, E. GPU-Based Airway Tree Segmentation and Centerline Extraction. PhD thesis. 2012.
26. Besl, PJ and McKay, ND. A Method for Registration of 3-D Shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 1992;14:239–256.
27. Lindseth, F, Langø, T, Bang, J, and Hernes, TAN. Accuracy evaluation of a 3D ultrasound-based neuronavigation system. *Computer aided surgery* 2002;7:197–222.

## Paper D

# Registration of MR to percutaneous ultrasound of the spine for image-guided surgery

Lars Eirik Bø<sup>1,2,3</sup>, Rafael Palomar<sup>4</sup>, Tormod Selbekk<sup>1</sup> and Ingerid Reinertsen<sup>1,2</sup>

<sup>1</sup>SINTEF Technology and Society, Trondheim, Norway

<sup>2</sup>Norwegian University of Science and Technology, Trondheim, Norway

<sup>3</sup>The Central Norway Regional Health Authority, Trondheim, Norway

<sup>4</sup>Oslo University Hospital, Oslo, Norway

In J. Yao, T. Klinder and S. Li (eds.) 2014, *Computational Methods and Clinical Applications for Spine Imaging*, Springer, pp. 209–218

**Abstract**

One of the main limitations of today's navigation systems for spine surgery is that they often are not available until after the bone surface has been exposed. Also, they lack the capability of soft tissue imaging, both preoperatively and intraoperatively. The use of ultrasound has been proposed to overcome these limitations. By registering preoperative magnetic resonance (MR) images to intraoperative percutaneous ultrasound images, navigation can start even before incision. We therefore present a method for registration of MR images to ultrasound images of the spine. The method is feature-based and consists of two steps: segmentation of the bone surfaces from both the ultrasound images and the MR images, followed by rigid registration using a modified version of the Iterative Closest Point algorithm. The method was tested on data from a healthy volunteer, and the data set was successfully segmented and registered with an accuracy of  $3.67 \pm 0.38$  mm.

**D.1 Introduction**

In spinal surgery today, many procedures are performed with no or only minimal image guidance. Preoperative computed tomography (CT) or magnetic resonance (MR) images are used for diagnosis and planning, but during surgery, two-dimensional C-arm fluoroscopy is widely used both for initial detection of the correct spinal level and for intra-operative imaging. Navigation systems exist, but mainly for placement of pedicle screws. These usually first come to use when the bone surface has been exposed. Using a simple landmark or surface registration method the preoperative CT image is then aligned with the patient and can be used for planning and guidance of the screws. A number of groups have evaluated the use of navigation for this purpose, and a review of the topic was presented by Tjardes et al. [1]. They conclude that the benefits of image-guidance in terms of accurate placement of the screws and reduced exposure to ionizing radiation have been proven, in particular for the cervical and lum-

bar procedures. In other areas of spine surgery, navigation and image guidance are still on the experimental stage.

One of the main limitations of today's navigation systems for spine surgery is that they often are not available until after the bone surface has been exposed. The use of ultrasound has been proposed to overcome this limitation. By registering preoperative images to intraoperative percutaneous ultrasound images, navigation can start before incision and therefore be used for both level detection and planning at an early stage of the procedure. Thus, the use of X-ray fluoroscopy can possibly be reduced.

In order to make a navigation system based on intraoperative ultrasound clinically useful, the greatest challenge is to achieve accurate and robust registration between the preoperative images and the ultrasound images with minimal user interaction. Registration of CT images of the spine to corresponding ultrasound images has been investigated by several groups, and two main approaches have been explored: feature-based registration and intensity-based registration. In the first case, corresponding features are extracted from the two datasets to be registered prior to registration. In the case of spine surgery, the feature of choice is the bone surface as this is the only feature that can be reliably detected in the ultrasound images. Segmentation of the bone surface from ultrasound images of the spine is still a challenging topic due to noise, artifacts and difficulties in imaging surfaces parallel to the ultrasound beam. A few methods have been described in the literature, ranging from simple ray tracing techniques [2] to more advanced methods based on probability measures [3–5] or phase symmetry [6]. Following surface extraction, the segmented bone surfaces are registered using the Iterative Closest Point (ICP) algorithm [7] or the unscented Kalman filter [5].

In intensity-based registration, a similarity metric based on the image intensities is optimized to find the spatial transformation that best maps one image onto the other. [2, 8–10] As MR/CT and ultrasound images present very different intensity and noise characteristics, a common approach is to create simulated ultrasound images from the pre-operative data and register the simulated image to the real

ultrasound image. In these simulations, the direction of sound wave propagation, transmission, reflection and noise can be modelled in order to obtain images that can be reliably registered to real ultrasound images based on image intensities.

While these studies show a lot of promise, they focus almost exclusively on the registration of preoperative CT images. However, many spinal procedures, such as the treatment of disc herniations and intraspinal tumours, rely on the soft-tissue imaging capabilities of MR. Thus, by combining ultrasound imaging with preoperative MR, navigation could be extended to a variety of spinal procedures that do not benefit from image guidance today. In these procedures, the ultrasound could also be used for intraoperative imaging, reducing the use of fluoroscopy even further. As a first step towards this end, we present a method for registration of preoperative MR images to percutaneous ultrasound images of the spine, including a preliminary assessment of its performance.

## **D.2 Methods and experiments**

Our registration method is feature-based and consists of two steps: First, the bone surfaces are segmented from both the ultrasound images and the MR images, and then the two surfaces are registered using a modified version of the ICP algorithm.

### **Ultrasound acquisition and segmentation**

The ultrasound images were acquired using a Vivid E9 scanner with an 11 MHz linear probe (GE Healthcare, Little Chalfont, UK). Some groups have used lower frequencies, which enable good imaging of deeper structures such as the transverse processes of the spine. [2, 5, 6, 8, 10] However, this makes imaging of superficial structures, such as the spinous processes and the sacrum, challenging. As these structures represent important features for the registration algorithm, we found that a relatively high frequency gave a better compromise between depth penetration and resolution. The ultrasound probe

was tracked with the Polaris optical tracking system (NDI, Waterloo, ON, Canada), and both images and corresponding tracking data were recorded using the navigation system CustusX [11] with a digital interface to both the ultrasound scanner and the tracking system. The two-dimensional ultrasound images were also reconstructed to a three-dimensional volume using the Pixel Nearest Neighbor (PNN) reconstruction algorithm. [12]

While the reconstructed, three-dimensional ultrasound volume is useful for navigation, the reconstruction process tends to introduce a certain blurring. The volume usually also has a lower resolution than the original, two-dimensional ultrasound images. We therefore used the latter as input to our segmentation method. In order to extract the bone surfaces from these images, we used a combination of the bone probability maps introduced by Jain et al. [3] and Foroughi et al. [4], and the backward scan line tracing presented by Yan et al. [2]. In ultrasound images, reflections from bone surfaces are seen as bright ridges perpendicular to the ultrasound beam. To calculate the probability of each pixel  $a^{i,j}$  of the image  $A$  being part of such a ridge, the image was smoothed with a Gaussian filter, before calculating the Laplacian of Gaussian (LoG), i.e.

$$A_G = \{a_G^{i,j}\} = A * G \quad \text{and} \quad A_{LoG} = \{a_{LoG}^{i,j}\} = A_G * L, \quad (D.1)$$

where  $G$  and  $L$  are the convolution kernels of the Gaussian filter and the LoG filter respectively. This is a common operation in blob detection and usually produces a strong positive response for dark blobs and a strong negative response for bright blobs. To enhance the bright reflections, the positive values were therefore set to zero before taking the absolute value of the rest. The result was then added to the smoothed version of the original image to produce an initial bone probability map  $P_1 = \{p_1^{i,j}\}$ , i.e.

$$p_1^{i,j} = a_G^{i,j} + |\max\{a_{LoG}^{i,j}, 0\}|. \quad (D.2)$$

The other feature to be considered was the intensity profile in the propagation direction of the ultrasound. For a bone surface, this is

typically characterized by a sudden, sharp peak followed by a dark shadow. To calculate the probability of a given pixel representing the maximum of such a profile, each scan line was considered separately. Assuming  $p_1^m$  is the  $m$ th pixel of the initial bone probability map  $P_1$  along a given scan line, the secondary bone probability of this pixel was found as

$$p_2^m = p_1^m - \frac{p_1^{m-\delta} + p_1^{m+\delta}}{2} - \frac{\omega}{\lambda} \sum_{n=1}^{\lambda} p_1^{m+\delta+n} , \quad (\text{D.3})$$

where  $2\delta$  is the width of a typical intensity peak and  $\lambda$  is the length of a typical bone shadow, both given in pixels. In our case, these were set to  $\delta = 24$  and  $\lambda = 322$ , which corresponds to 1.5 mm and 20 mm respectively.  $\omega$  is a weight that can be adjusted according to the overall noise level of the bone shadows in the image, and in our case this was set to 10.

The first term in (D.3) is simply the intensity of the  $m$ th pixel. At a bone reflection, this will be high and lead to a high bone probability. The second term combines the intensities at the distance  $\delta$  behind and in front of the  $m$ th pixel. At a sharp peak of width  $2\delta$ , both of these will be low and have little impact on the bone probability. On the other hand, if there is no such peak, at least one of these will be high and lead to a reduced bone probability. The last term is the average intensity of the pixels in the shadow region behind the peak. If there is a lot of signal in this area, this term will be high and thus reducing the bone probability

Finally, we applied a variant of the backward scan line tracing to the resulting probability map: For each scan line, starting at the bottom of the image, the first local maximum above a certain threshold was deemed part of a bone surface. This was repeated for all the recorded images, and based on the corresponding tracking data, all points were transformed into the three-dimensional reference space of the tracking system. A typical example of both the probability maps and the final segmentation of an image is shown in Figure D.1. The method was implemented in MATLAB (MathWorks, Natick, MA, USA).

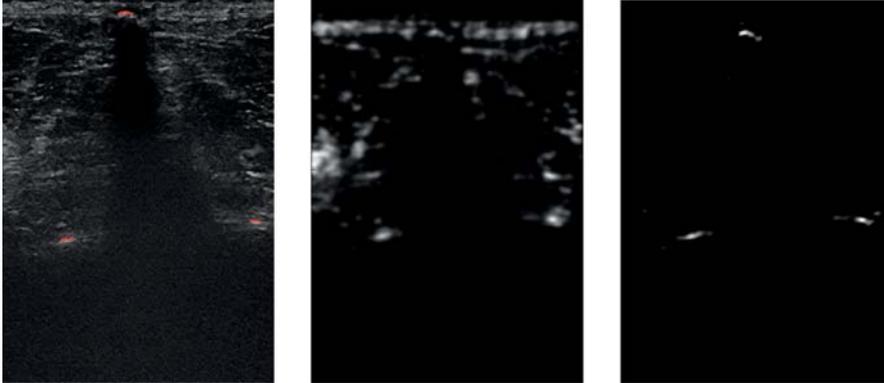


Figure D.1: An ultrasound image of a vertebra with the segmentation overlaid in red (left), the initial bone probability map (centre) and the final bone probability map after applying the threshold (right).

### MR acquisition and segmentation

The MR images were acquired using an Achieva 3.0 T scanner (Philips Healthcare, Amsterdam, Netherlands). In order to facilitate both the segmentation of the spine and the subsequent navigation, we customized a full, three-dimensional MR protocol which enhanced the contrast between the bone and the surrounding soft tissue. This had a field of view of  $80 \times 560 \times 560$  voxels and a voxel size of  $1 \times 0.48 \times 0.48 \text{ mm}^3$ . The lumbar vertebrae were segmented using a semiautomatic method based on active contours implemented in the segmentation software ITK-SNAP. [13] However, in the area of the sacrum, the contrast between the bone and the surrounding soft tissue was lower, and here active contours driven by robust statistics resulted in more accurate segmentations. For this part, we therefore employed the Robust Statistics Segmentation (RSS) module [14] included in the medical imaging analysis and visualization software 3D Slicer [15].

The use of active contours for segmentation may lead to oversegmentation of certain anatomical structures, known as leaks. In MR images, such leaks are especially prominent in areas with motion artifacts caused by the patient not lying completely still during the image

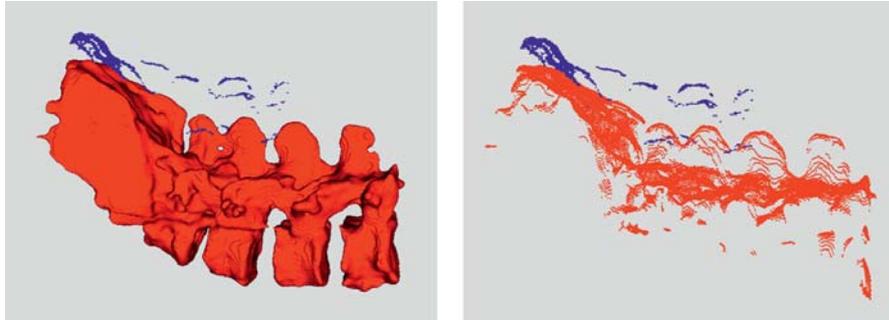


Figure D.2: The segmented ultrasound (blue) and MR (red) surfaces (left) and the same surfaces after reducing the MR surface with ray tracing (right).

acquisition. This is often a problem, especially for patients in need of spine surgery. To compensate for this, minor corrections of the segmentation results were performed manually for both the lumbar area and the sacrum.

The surface segmented from the MR volume represented the entire lumbar spine, and consisted therefore of a large number of points. However, only the surfaces facing the ultrasound probe were visible in the ultrasound images. Thus, a significant portion of the surface points in the segmented MR were irrelevant to the registration, as there were no corresponding points in the ultrasound images. To reduce the amount of data, and thus the work load of the registration algorithm, we therefore used a simple ray tracing method (posterior to anterior) to extract those points that were facing the ultrasound probe. An example of the resulting reduced surface can be seen in Figure D.2.

## Registration

Following segmentation, the segmented surfaces from ultrasound and MR were imported into the navigation system for registration. Like all automatic registration methods, the ICP algorithm requires an initialization or a reasonable starting point in order to converge to the correct solution. This was provided by assuming that the two vol-

umes covered approximately the same volume, that the first recorded ultrasound image was positioned at the sacrum and that the probe trajectory was from the sacrum upwards. The two image volumes were then aligned by first rotating the MR volume in order to align the x, y and z axes in the two volumes, and then translating the MR volume in order to align the points corresponding to the voxels  $(n_x/2, 0, 0)$  in both volumes, where  $n_x$  is the number of voxels in the x-direction (patient left-to-right).

After this initial alignment, we used the ICP algorithm to rigidly register the reduced MR surface to the ultrasound surface. In order to reduce the influence of possible outliers on the registration result, the algorithm was modified by incorporating the Least Trimmed Squares (LTS) robust estimator as described by Reinertsen et al. [16].

## Experiments

In order to evaluate our method, we acquired both ultrasound and MR images of the spine of a healthy volunteer. The only structures that were clearly discernible in both of these images were the top points of the spinous processes of three lowest vertebrae (L3, L4 and L5). These were therefore selected as control points and manually identified in both the original ultrasound volume and the MR volume. The surfaces were then registered to each other using the method described above, and the distances between the landmarks both after initial alignment and after final registration were computed.

## D.3 Results

Through careful optimization of the acquisition protocols, both MR and ultrasound images of high quality were achieved. The data sets were successfully segmented and registered using the methods described above. Figure D.3 shows the extracted surfaces both after the initial alignment and after rigid registration. The match can also be seen in Figure D.4, which shows transverse and sagittal views of corresponding ultrasound and MR volumes after registration. Finally,

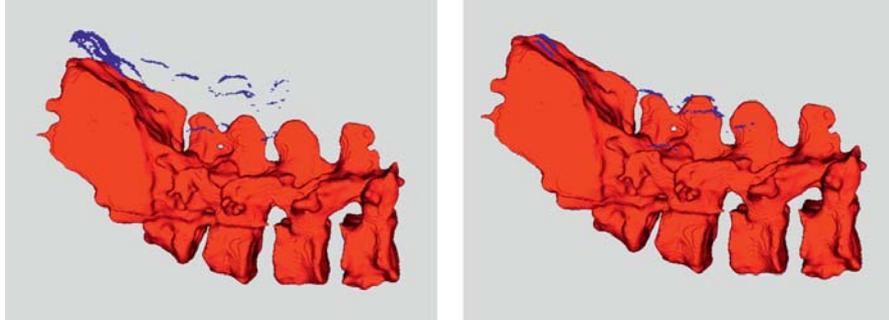


Figure D.3: The ultrasound (blue) and MR (red) surfaces after the initial alignment (left) and after the final registration (right).

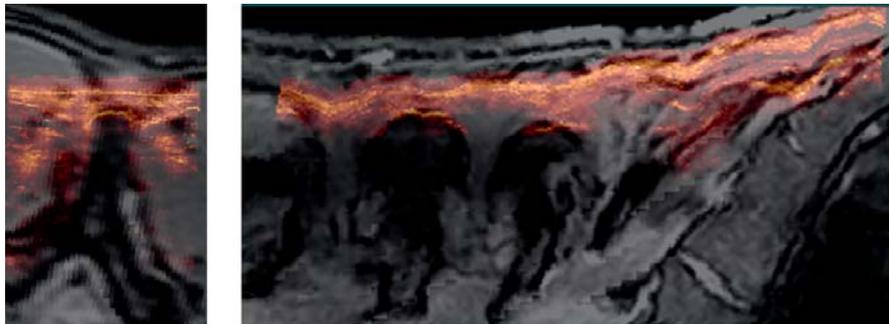


Figure D.4: A transverse slice (left) and a sagittal slice (right) from the ultrasound volume overlaid on top of the corresponding slices from the registered MR volume. The ultrasound data is shown in red and yellow and the MR data is shown in grey tones.

the distances between the control points before and after registration are given in Table D.1.

## D.4 Discussion

We have demonstrated that registration between MR and ultrasound images is feasible. The accuracy of  $3.67 \pm 0.38$  mm is clinically relevant

Table D.1: Distance between the control points in mm.

|                          | L3    | L4    | L5    | Mean±STD   |
|--------------------------|-------|-------|-------|------------|
| After initial alignment  | 23.29 | 21.27 | 22.40 | 22.32±1.01 |
| After final registration | 3.86  | 3.93  | 3.23  | 3.67±0.38  |

as it is sufficient to ensure that we are on the correct level. It is also comparable to that of many of the studies mentioned in the introduction. Still, this is a work in progress, and the results shown here are only preliminary.

It has been pointed out that intensity-based registration has an advantage over feature-based methods in that it makes use of all the information in the image, rather than just that of the bone surfaces. [10] In the case of spine imaging, however, other structures that are visible in the ultrasound images, such as muscle fibres and fat layers, are not imaged very well by neither CT nor MR. Their contribution to the registration procedure is therefore questionable.

The ultrasound images that we have acquired vary considerably in appearance from subject to subject. At the moment, this means that the parameters of the segmentation method, such as the width  $\delta$  of the reflections, the length  $\lambda$  of the shadows and the weight  $\omega$  must be manually adjusted to the particular data set. In the future, these adjustment should be done automatically, e.g. based on overall image statistics.

The MR segmentation methods that we presented here are only semiautomatic and quite time consuming. However, the result of this was a complete segmentation of the lumbar spine, and as we have already pointed out, only a small part of this information was actually relevant to the registration. We are therefore investigating methods to segment only the part of the anatomy that is most critical to the registration, i.e. the sacrum and the spinous and transverse processes. The results are promising, and it should be possible to perform this

segmentation both quickly and with minimal user interaction.

The last component of the method is the registration. Here, we have shown that a reasonable rigid registration can be achieved using the ICP algorithm. However, the spine is flexible, and the change in curvature from the MR scanner, where the patient is lying in a supine position, to the operating room, where the patient is placed in a prone position, can be large. A group-wise rigid registration method, like the one proposed e.g. by Gill et al. [10] where only the space between the vertebrae is deformed, would be more appropriate.

Finally, our method needs more extensive testing, both with respect to robustness to anatomical variations and with respect to accuracy. The distance measure that we have used here, based on manual identification of landmarks, gives a good indication of the registration accuracy, but we should include a measure of inter- and intra-observer variability. Such measures could therefore be complimented with other assessment methods, such as phantom studies where the exact geometry is known and a reliable ground truth thus can be established. All of the above are currently addressed in our research.

## **D.5 Conclusion**

The presented method is capable of registering MR images to percutaneous ultrasound images of the spine. The registration accuracy is clinically relevant, and with minor improvements the user interaction can be reduced to a minimum. This method is thus an important step towards the realisation of a system for MR- and ultrasound-guided spine surgery.

## **Acknowledgements**

The work was funded through the user-driven research-based innovation project VIRTUS (The Research Council of Norway grant no. 219326, SonoWand AS) and through a PhD grant from the Liaison

Committee between the Central Norway Regional Health Authority (RHA) and the Norwegian University of Science and Technology.

## References

1. Tjardes, T, Shafizadeh, S, Rixen, D, Paffrath, T, Bouillon, B, Steinhausen, ES, and Baethis, H. Image-guided spine surgery: state of the art and future directions. *European Spine Journal* 2010;19:25–45.
2. Yan, CXB, Goulet, B, Tampieri, D, and Collins, DL. Ultrasound–CT registration of vertebrae without reconstruction. *International Journal of Computer Assisted Radiology and Surgery* 2012;7:901–909.
3. Jain, AK and Taylor, RH. Understanding Bone responses in B-mode Ultrasound Images and Automatic Bone Surface extraction using a Bayesian Probabilistic Framework. In: *Medical Imaging 2004: Ultrasonic Imaging and Signal Processing*. Ed. by Walker, WF and Emelianov, SY. Vol. 5373. Proceedings of SPIE. SPIE, 2004:131–142.
4. Foroughi, P, Boctor, E, Swartz, MJ, Taylor, RH, and Fichtinger, G. Ultrasound Bone Segmentation Using Dynamic Programming. In: *2007 IEEE Ultrasonics Symposium*. Ed. by Yuhas, MP. IEEE, 2007:2523–2526.
5. Rasoulia, A, Abolmaesumia, P, and Mousavi, P. Feature-based multibody rigid registration of CT and ultrasound images of lumbar spine. *Medical Physics* 2012;39:3154–3166.
6. Tran, D and Rohling, RN. Automatic Detection of Lumbar Anatomy in Ultrasound Images of Human Subjects. *IEEE Transactions on Biomedical Engineering* 2010;57:2248–2256.
7. Besl, PJ and McKay, ND. A Method for Registration of 3-D Shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 1992;14:239–256.

8. Winter, S, Brendel, B, Pechlivanis, I, Schmieder, K, and Igel, C. Registration of CT and Intraoperative 3-D Ultrasound Images of the Spine Using Evolutionary and Gradient-Based Methods. *IEEE Transactions on Evolutionary Computation* 2008;12:284–296.
9. Lang, A, Mousavi, P, Gill, S, Fichtinger, G, and Abolmaesumi, P. Multi-modal registration of speckle-tracked freehand 3D ultrasound to CT in the lumbar spine. *Medical Image Analysis* 2012;16:675–686.
10. Gill, S, Abolmaesumi, P, Fichtinger, G, Boisvert, J, Pichora, D, Borshneck, D, and Mousavi, P. Biomechanically constrained groupwise ultrasound to CT registration of the lumbar spine. *Medical Image Analysis* 2012;16:662–674.
11. Askeland, C, Lindseth, F, Solberg, OV, Bakeng, JBL, and Tangen, GA. CustusX: A Research Application for Image-Guided Therapy. *The MIDAS Journal* 2011;Systems and Architectures for Computer Assisted Interventions 2011:1–8.
12. Solberg, OV, Lindseth, F, Bø, LE, Muller, S, Bakeng, JBL, Tangen, GA, and Hernes, TAN. 3D ultrasound reconstruction algorithms from analog and digital data. *Ultrasonics* 2011;51:405–419.
13. Yushkevich, PA, Piven, J, Hazlett, HC, Smith, RG, Ho, S, Gee, JC, and Gerig, G. User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability. *NeuroImage* 2006;31:1116–1128.
14. Gao, Y, Kikinis, R, Bouix, S, Shenton, M, and Tannenbaum, A. A 3D interactive multi-object segmentation tool using local robust statistics driven active contours. *Medical Image Analysis* 2012;16:1216–1227.
15. Fedorov, A, Beichel, R, Kalpathy-Cramer, J, Finet, J, Fillion-Robin, JC, Pujol, S, Bauer, C, Jennings, D, Fennessy, F, Sonka, M, Buatti, J, Aylward, S, Miller, JV, Pieper, S, and Kikinis, R.

- 3D Slicer as an image computing platform for the Quantitative Imaging Network. *Magnetic Resonance Imaging* 2012;30:1323–1341.
16. Reinertsen, I, Descoteaux, M, Siddiqi, K, and Collins, DL. Validation of vessel-based registration for correction of brain shift. *Medical Image Analysis* 2007;11:374–388.