Author: Malan, Daniel Francois
Title: Pinning down loosened prostheses : imaging and planning of percutaneous hip refixation
Issue Date: 2015-10-29

9

General Discussion

The research in this thesis represents part of a project to develop a new set of tools to support minimally invasive treatment for aseptically loose hip prostheses. In Chapter 1, Fig. 1.1 illustrated how the elements of this project relate to one another. These elements are intended to work together as an integrated whole, and we undertook our research as members of a multidisciplinary research group. Together, we aimed to contribute to an integrated solution that improves the planning and execution of minimally invasive stabilization of aseptically loosening hip prostheses. Each of the following topics plays a role: Image processing (this thesis), pre-operative planning and intra-operative guidance (this thesis), mechanical finite element simulations, and surgical instrument design.

9.1 The effect of fibrous interface tissue before cement injection

Fibrous interface tissue should ideally be removed before an aseptically loose prosthesis is stabilised, as osteolysis is progressive and self-propagating [Park et al., 2004]. Computer simulations have shown that interface tissue removal increases the mechanical stability of a hip prosthesis [Andreykiv et al., 2012]. Currently, there are no minimally invasive instruments with which to manually remove fibrous interface tissue before bone cement is injected. To examine the effect that removal of interface tissue had in actual patients, we analysed the post-operative CTs of patients that were treated at our institution between 2006 to 2011 (Chapter 3). These patients all eventually received percutaneous cement injections to stabilise their aseptically loose hip prostheses, yet their treatment differed in one important aspect: One patient group was pre-operatively treated using gene-directed enzyme pro-drug therapy (GDEPT) to eliminate fibrous interface tissue, while the other group received no pre-operative treatment. Our measurements indicated that pre-operative treatment increased the volume and efficacy of percutaneous cement injection into the periprosthetic targets. From our results we deduced that prior removal of periprosthetic fibrous interface tissue may enable better cement flow and cement penetration. This may in turn lead to better prosthesis stabilisation. The implication of this result is that new instruments by which fibrous interface tissue can be removed would benefit the patient. The development of such instruments fall outside the scope of this thesis, but is being further pursued by researchers at our institution [Kraaij et al., 2012, den Dunnen et al., 2013].

9.2 Image capture and image processing

Surgeons who perform hip surgery often judge traditional two dimensional (2D) X-ray radiographs to be sufficient for diagnosis and planning their procedures. X-rays are relatively cheap, fast, and deliver only a low radiation dose. For the diagnosis of arthritic damage to a hip, a high-resolution 2D X-ray image shows the associated narrowed joint space and uneven femoral surface. A suitable prosthesis size may be selected by overlaying templates directly on the printed X-ray plate. Hip arthroplasty is performed during open surgery that gives the surgeon direct visual and tactile access to the affected hip, providing him/her with the direct intra-operative insight that is required to perform the procedure. When instead working percutaneously through a minor incision, the surgeon cannot directly see the tips of his/her instruments. The surgeon therefore becomes dependent on medical imaging or guidance to provide the necessary intra-operative insight. Because the patient's body is a complex 3D object, 2D X-ray images may not be sufficient to meet the needs of minimally invasive surgery.

Three dimensional (3D) medical imaging technology, also known as cross-sectional imaging, has dramatically improved over the last two decades. Modern Computed Tomography (CT) and Magnetic Resonance (MR) imaging devices are now widely available and produce image volumes of the human body that can resolve details of 0.5mm or finer in any spatial direction.

At the time we performed our research, CT was seen as the modality of choice for quantifying periprosthetic osteolysis [Egawa et al., 2009], and may still be seen as such. We therefore based our image processing and planning software on CT. Medical CT scanners are designed to image the human body. As metal differs dramatically from biological tissues, CT image-reconstruction algorithms ecounter problems when a hip containing a hip prostheses is imaging. Metal-induced imaging artefacts ensue. To improve image quality, metal artefact reduction (MAR) algorithms have been developed that suppress the severe "beam hardening" seen in CT [Cahir and Toms, 2009]. Artefact reduction by projection interpolation is one of the oldest MAR methods [Glover and Pelc, 1981] and has continued to garner interest in recent years [Oehler et al., 2008, Veldkamp et al., 2010]. Alternative iterative and algebraic approaches to CT MAR have been developed [Bal and Spies, 2006, Lemmens et al., 2009, Boas and Fleischmann, 2011, Gondim Teixeira et al., 2014], yet to our knowledge projection interpolation remains the only method to have been clinically implemented by a major CT manufacturer [Watzke and Kalender, 2004]. The limited uptake of MAR can possibly be explained by the observation that metal artefacts are, in practice, not as destructive as one is led to believe. The severity of “star streak" artefacts shown in publications on the topic [Boas and Fleischmann, 2011, Lemmens et al., 2009, Kalender et al., 1987] depend on the contrast window chosen; some information that appears lost may be recovered if the complete intensity range of the data is considered. There is also the issue that MAR may in itself degrade image detail, especially directly adjacent to a metal prosthesis. In our research we found that MAR by linear sinogram interpolation reduced our ability to delineate peri-prosthetic lesions created in cadaver specimens [Malan et al., 2012b]. While MAR used in isolation has its drawbacks, we showed in Chapters 5 and 6 that MAR is useful when used to augment instead of replace standard CT images.

Metal-induced image artefacts not only occur in CT, but also in MRI and ultrasound [Sofka, 2007]; each by a different underlying mechanism. In our experience orthopaedic surgeons are well aware of the difficulties associated with imaging close to metal implants, yet less so of available mitigation strategies. This is especially true for MRI which research shows to be safe [Kumar et al., 2006] and useful [Cahir and Toms, 2009, Lee et al., 2007, Potter and Foo, 2006] for peri-prosthetic imaging of the hip. For this reason we believe that MRI may in future become an increasingly valuable tool for assessing osteolysis.

A 3D image volume provides a wealth of data that can be used for pre-operative planning and simulation. Common orthopaedic procedures for which CT-based pre-operative planning systems have been developed include hip arthroscopy [Audenaert et al., 2012], shoulder arthroplasty [Botha and Post, 2001], hip arthroplasty [Dick et al., 2009] and knee arthroplasty [Ellis, 2005]. Before being useful for modelling aseptic loosening, some crucial post-processing steps need to be taken on a medical image volume.

9.3 Transforming medical images to 3D models of the hip

3D image volumes may be used for powerful and beautiful medical visualizations [Preim and Botha, 2013b]. However beautiful these visualizations can be, it is bio-mechanical simulations that are often the desired tool for performing quantitative analyses. Such analyses usually require that intensity images with clearly defined tissue boundaries are segmented to form a multi-material 3D model.

The golden standard for 3D image segmentation is manual delineation by a trained human [Seim et al., 2008, Lim et al., 2006]. It is difficult to interact with a 3D volume via a computer's 2D display and therefore segmentation is often performed in a slice-by-slice manner. With the steady increase in resolution of modern MRIs and CTs, manual segmentation may already be prohibitively time consuming for routine clinical practice. Manual, or even semi-automated image segmentation is a tedious and labour-intensive task that benefits greatly from automation [Heimann et al., 2009]. In this thesis we preferred automated over semi-automatic methods whenever we could.

The research community has already invested much effort into designing and improving automated tissue classification algorithms. Advances in automatic segmentation include examples as diverse as for the liver [Heimann et al., 2009], brain [van der Lijn et al., 2008], heart [van Assen et al., 2006], lungs [van Rikxoort et al., 2009] and retinal lesions [Sánchez et al., 2011]. In this thesis we investigated whether one can automatically segment the peri-prosthetic tissues and structures of a hip that is affected by osteolysis [Malan et al., 2010, Malan et al., 2012a]. We wanted to discern tissues relevant to the creation of a finite element model of the aseptically loosening hip prosthesis [Andreykiv et al., 2012]. These include cortical bone, trabecular bone, fibrous interface tissue, bone cement, and the metal prosthesis. Our problem was especially challenging for two reasons: Firstly, the presence of a metal prosthesis degrades the CT image quality. Secondly, peri-prosthetic osteolysis does not have a well-defined normally distributed variation around an average shape [Scheerlinck et al., 2008, Howie et al., 2007]. The powerful class of "statistical shape models" and "active appearance models" that rely on the assumption of Gaussian shape distributions [Heimann and Meinzer, 2009] were therefore largely eliminated from our solution space. To make matters even more challenging, this thesis addressed a medical procedure which is itself still considered experimental. As such there exist few publications on this topic and there were no publicly available data-sets or competing algorithms with which to compare ours—something that is of great value in image segmentation research [Heimann et al., 2009, Martin et al., 2001].

In our approach to tissue classification we used several complementary image volumes simultaneously—including unprocessed, spatially filtered, and metal-artefact-reduced image volumes. Graph cuts is a powerful segmentation technology that has, up to now, not been widely applied in medical imaging outside of selected settings [van der Lijn et al., 2008, Merickel et al., 2007]. The addition of Graph Cuts to our classifier enabled us to enforce geometrical constraints that would not have been achievable by treating each image voxel individually, and it improved our classification results notably.

Before mechanical simulation of a bone can be performed, even a segmented 3D tissue map must be converted into an appropriate simulation model [Taddei et al., 2006, Andreykiv et al., 2012]. The Finite Element Method (FEM) [Cody et al., 1999] is a well-established and validated method of modelling stresses, strains and deformations in continuum models and is used in countless applications including the automotive world, civil engineering, and modelling the human body. In its standard form the FEM requires the simulated model to be represented as a mesh of connected elements—typically tetrahedra [Bryan et al., 2010]. The number of mesh elements need to be kept low enough for the model to be tractable, while in a multi-material model nodes and edges must conform on all material interfaces.

In simulations of man-made structures, the input to the meshing algorithm are often models created in computer aided design (CAD) software. These models are analytically defined and their surfaces are threfore noise-free and easy to discretize cleanly using a limited number of elements. In contrast, the segmented tissue map obtained from a 3D medical image volume is discretized as a dense grid of rectangular voxels that can be seen as a noisy approximation of the underlying structures. Converting this dense grid to a volumetric mesh can create an intractable number of mesh elements. Single-material meshing methods exist that can create well-behaved meshes with a limited element-count [Labelle and Shewchuk, 2007]. Doing the same for multi-material volumes is more challenging, and the lack of widely available solutions have provided an opportunity for some companies to build commercial software to perform this task [Pettersen and Høgetveit, 2011]. In Chapter 7 we built upon prior work [Meyer et al., 2008] and extended their particle-based algorithm for multi-material meshing. Particle based methods have an advantage over direct mesh refinement in that it abstracts the shape that need to be meshed as an implicit surface, which allows for scale invariance. We presented a proof-of-concept software mesher that, compared to the Delaunay-based algorithm on which it builds, allows sparser multi-material meshes to be created without having to sacrifice geometric surface accuracy by excessive smoothing.

CT has its shortcomings in the presence of a large metal prosthesis, yet has one distinct advantage over e.g. MRI and ultrasound: it quantifies the human body in Hounsfield units. The Hounsfield Unit is an objective and reproducible measure that directly corresponds to the physical rate at which X-rays are absorbed by the body's tissues. This in turn allows mechanical material properties to be derived from a CT image. Mechanical properties may then be assigned to patient-specific models of, for example, the femur [Bessho et al., 2009]. We caution that dense bone and metallic prostheses have a measurable effect on image intensity values and may disrupt algorithms that rely strongly on absolute image intensity values for tissue classification. We briefly looked into this topic in Chapter 4.

9.4 Planning and guiding a minimally invasive procedure

Software by which the execution of orthopaedic operations can be planned are already in clinical use. Examples include software for needle biopsies [Braak et al., 2010] and for planning hip arthroscopy [Audenaert et al., 2012]. 3D planning software commonly use CT or MRI image volumes as input, and can manipulate and visualize data either on standard desktop computers or dedicated hardware stations. Planning may be performed prior to an operation and outside of the operating room. Ideally, such software should also help the surgeon to operate according to the predetermined plan that is constrained by the operating room environment. Limitations may include the amount of time allowed for treating the patient and the need to maintain sterility, i.e. that computer interfaces may not be physically touched.

The stabilizing effect of injected cement on a loosened hip prosthesis depends on how much contact is re-established between the bone and cement surrounding the prosthesis [Andreykiv et al., 2012]. In our specific application of minimally invasive cement injection, knowing the amount of cement that reached its intended target may therefore be important in judging the success of the procedure.

A well-established intra-operative guidance approach that is often used in orthopaedics is to attach reflective optical markers to the patient's body and the surgeon's instruments, and then to track the movement of these with a stereoscopic camera system [Bäthis et al., 2004]. This is a useful approach in positioning rigid instruments. In our application we are, however, also interested in the progressive distribution of radiopaque cement fluid before it hardens—something that optical markers cannot help us with. The amount of deposited cement can traditionally only be quantified by the displacement of the cement syringe's plunger, on which calibrated millilitre marks are indicated along the side.

We believed that—if at all possible—it is preferable to design a guidance system that does not require novel imaging equipment to be present in the operating room. Re-use of existing technology may ease adoption of a new planning system and will reduce costs. An X-ray C-arm fluoroscope is commonly available in the operating room due to its compact size, multiple uses and relatively low cost, whereas CT and MRI are not. In Chapter 8 we therefore presented HipRFX, a proof-of-concept guidance system that requires only an intra-operative X-ray fluoroscope to be present during a cement-injection operation, assuming that a CT image was obtained beforehand. We are not the first to use this approach—The Philips XperGuide guidance system for needle biopsies [Braak et al., 2010] also makes use of pre-operatively acquired CT for planning and intra-operative X-ray fluoroscopy for guidance. Our system does, however, deliver some innovations that we have not yet come across elsewhere. These include the ability to quantitatively estimate injected cement volumes based on fluoroscopy images only.

In pre-clinical tests our software was able to generate reasonable estimates of the percutaneously injected cement volume that reached the intended target—information that would otherwise not be quantitatively known from a single post-operative fluoroscopic image. Our system can furthermore record a step-by-step series of simulated fluoroscopy images. These simulated fluoroscopy snapshots may guide a surgeon during the execution of a minimally invasive cement injection procedure by alerting him/her of deviations from the plan, or indicating when the intended goal has been achieved. The promising results we obtained in Chapter 8 were despite the sub-optimal conditions under which we performed our experiments, namely that the hardware we used was uncalibrated and required several estimation steps that may be avoided in future.

Our guidance software may be further developed along several lines. One idea is to improve cement estimates by integrating the information obtained from multiple 2D fluoroscopic views that were taken simultaneously or in quick succession from different angles. Another idea is to use our software to automatically compute the optimal C-arm view angles to use for an intended intervention. The need for needle placement guidance may also in future be addressed by our system, as is done in the Philips XperGuide [Racadio et al., 2007, Leschka et al., 2012]. The planning and guidance approach we showcased with HipRFX may eventually also have applications beyond the procedure for which is was developed. One such possible application could be vertebroplasty [Phillips, 2003], which also relies on the injection of radiopaque bone cement.

9.5 Perspectives on the Future

At the time of writing, percutaneous hip prosthesis refixation is still considered novel and even experimental, yet it has already helped several patients at our institution who would otherwise—because of their frail health—not have qualified for surgery. There are currently no purpose-made tools for this novel minimally invasive orthopaedic procedure. While we addressed several aspects of imaging a patient's hip and pre-operative planning, a multi-faceted approach that goes wider than the content of this thesis is needed to advance this procedure towards greater maturity.

Despite CT's established role pre-operative imaging of peri-prosthetic loosening [Cahir et al., 2007], several recent studies have shown that Magnetic Resonance Imaging (MR) has great potential in imaging the peri-prosthetic region of the hip, even in the presence of large metallic prostheses [Cahir and Toms, 2009, Sofka, 2007, Walde et al., 2005]. The additional advantage that MRI delivers no ionising radiation will be especially advantageous when one considers younger patients. We believe that the tissue classification algorithm described in Chapters 5 and 6 of this thesis will allow adaptation to MRI. Our algorithm teaches itself to classify tissues from heterogenous data, and does not depend on any inherent quality of CT.

Throughout this thesis we used the example of a loose femoral stem to illustrate and test our work. While a loosening prosthesis stem is indeed often the reason for a failing hip prosthesis, the acetabular cup component also fails [Howie et al., 2007] at rates reported to be even higher than for the femoral stem [Keener et al., 2003, Swe, 2013, Dut, 2013]. We hypothesize that the methods we developed are also adaptable to the acetabular cup, as well as to other joints where aseptic loosening can occur such as the knee.

The algorithms and software we developed can all be described as proof-of-concept. None of the software has been developed to industrial-level robustness, and can therefore not yet be clinically validated. The scope of our research left some resolvable yet important bottlenecks unaddressed. One of these is the infrastructure for enabling effortless data-flow between the different components described in the thesis. This infrastructure includes file format conversions, network communication with picture archiving and communication (PACS) systems to obtain CT image data, reading images from intra-operative C-arm fluoroscopic devices, and communicating models to finite element modelling software. Image registration and intensity matching between intra-operative fluoroscopic images and simulated fluoroscopic images was demonstrated, but not fully automated nor integrated in our software platform. Steps as these are not necessarily interesting in a research context, yet are extremely important from an operational standpoint and may require substantial effort to implement properly.

The ability to simulate the biomechanics of the affected hip may need to be improved in order to compute optimal cement placement [Andreykiv et al., 2012]. The development and validation of new surgical instruments is desired so that peri-prosthetic fibrous interface tissue may removed prior to cement injection [Malan et al., 2014]. Research on the aforementioned topics was conducted at our institute concurrently with the work in this thesis, but much work remains.

Patient-specific bio-mechanical modelling may currently still be a niche technique due its combined requirements for advanced 3D imaging, software than can transform such images into 3D models, as well as suitable simulation software. Success stories include unique pathologies that preclude the use of standardised prostheses, such as creating prostheses for missing parts of a patient's skull [Esses et al., 2011, Metzger et al., 2008]. When fast and reliable automatic image volume segmentation techniques become available, routine clinical applications to patient-specific modelling may follow. A recent example where patient-specific simulations made the transition to a mainstream orthopaedic application is that of femoro-acetabular impingement [Audenaert et al., 2012]. We foresee that routine patient-specific 3D modelling will eventually become the norm for a full gamut of orthopaedic applications, as capabilities of 3D modelling tools improve and their cost lowers.

Given the ever growing number of patients that receive hip prostheses and prostheses of other major joints, we expect that ever more patients will be affected by the clinical problem of prosthesis loosening caused by aseptic periprosthetic osteolysis. We believe that the analysis and treatment of aseptic prosthesis loosening will therefore require continued attention and we look forward to developments in this field.

9.6 Conclusion

In this thesis we examined how computer software can be used to assist a minimally invasive cement injection procedure that is intended to stabilize an aseptically loose hip prosthesis. We chose to base our methods on the output of widely available medical imaging hardware such as Computed Tomography (CT) and X-ray fluoroscopy. We believe that solutions using widely available imaging hardware has a better chance of being widely applicable and adoptable into existing clinical practice.

Taken as a whole, this thesis addresses several aspects that constitute a medical imaging workflow to support percutaneous hip prosthesis refixation by cement injection. Throughout our research we focused on CT as the source of 3D medical images. Through the analysis of clinical data from our institution, we conclude that prior removal of periprosthetic fibrous interface tissue may enable better cement flow and penetration.

One of our goals was to enable the generation of finite element mesh of the hip that include the complexities of aseptic loosening. This has to be done without resorting to a solution that generates a superfluous number of mesh elements. To this end, we presented an automated pipeline for segmenting periprosthetic tissues in clinical CT volumes of patients with hip prostheses. We subsequently improved on a viable approach to multi-material conformal tetrahedral meshing.

We also worked towards a pre-operative planning and intra-operative guidance tool. For this we developed HipRFX, a proof-of-concept software module. In pre-clinical tests HipRFX was able to simulate intra-operative fluoroscopic C-arm images, providing the surgeon with a step-by-step series of snapshots with which to compare the operation as it is performed. We were furthermore able to generate estimates of the percutaneously injected cement volume that actually reached the intended targets.

The algorithms and software that we developed can all be described as proof-of-concept. Throughout this thesis we have used the example of a loose femoral stem to illustrate and test our work, and we believe that the methods we developed may also be applicable to the femoral cup, as well as other joints where aseptic loosening occurs. We have proposed some paths along which our initial steps may be taken further, and we look forward to see where these paths will lead.