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

Summary

In this thesis we examined how computer software can be used to analyse medical images of an aseptically loosening hip prosthesis and subsequently to plan and guide a minimally invasive cement injection procedure to stabilize the prosthesis. The research in this thesis was undertaken as a member of a research group where, since 2005, a novel treatment method to this end was developed and tested on a small group of selected patients.

We specifically endeavoured to develop automated or semi-automated algorithms that could enable patient-specific planning and treatment within the time constraints of clinical practice. It was also our goal to make use of existing medical imaging hardware, such as Computed Tomography (CT) and C-arm fluoroscopy, so that solutions arising from this research could eventually be adopted into existing practice.

The first topic we addressed was the detection and measurement of periprosthetic bone lesions from CT image volumes. In Chapter 2 experimentally examined whether we could accurately measure the volumes of periprosthetic bone lesions in the presence of image artefacts caused by the presence of a metal hip prosthesis. The answer to this question was important as many of our subsequent experiments and algorithms depended on the visibility of these lesions in three dimensional (3D) CT image volumes. In the presence of a large metallic prosthesis, severe "beam hardening" image artefacts degrade image quality. We found that we were able to delineate these periprosthetic lesions, although metal artefact reduction (MAR) by linear sinogram interpolation reduced our ability to delineate lesions. This is counter-intuitive as the goal of MAR is to improve image quality. While MAR used in isolation has its drawbacks, we showed in Chapters 5 and 6 that MAR is useful when it is combined with unprocessed CT images to improve tissue classification results.

In Chapter 3 we examined post-operative CTs of patients that were treated at our institution in the period spanning 2006 to 2011. These patients all

eventually received percutaneous cement injection to refixate their aseptically loose hip prosthesis, but their treatment differed in one important aspect: one patient group was pre-operatively treated using gene-directed enzyme prodrug therapy (GDEPT) to kill fibrous interface tissue, while the other group received no pre-operative treatment. We measured the volumes of each patient's periprosthetic osteolytic lesions, as well as the volume of cement that was injected during the procedure. Our results showed that patients who underwent pre-operative GDEPT to kill fibrous tissue exhibited larger injected cement volumes than those who did not. From this we deduced that prior removal of periprosthetic fibrous interface tissue enabled better cement flow and cement penetration that might in turn lead to better prosthesis refixation. We propose that new instruments by which fibrous interface tissue may be removed would aid the success of the procedure.

In our experiments in Chapter 4 we briefly examined whether radio-dense materials such as those encountered in the hip can significantly alter the CT image intensity values observed in adjacent structures. We found that this is indeed the case, and cautioned that this effect may disrupt algorithms that rely strongly on absolute image intensity values for tissue classification or for deriving numerical material properties such as elasticity. We concluded the chapter by providing some estimates of errors that may be introduced into derived finite element models if such inaccuracies are allowed to propagate to the computational stage.

In Chapter 5 we developed a tissue classification algorithm that may automatically label periprosthetic bone, cement and fibrous interface tissue. The ability to label these tissues is important for mechanically simulating an in-vivo hip prosthesis. We first explored the use of statistical classifiers that treat 3D image cubes (voxels) as individual entities. Every voxel was represented by a set of numerical values corresponding to image features. Our algorithm combines several complementary image processing outputs, including MAR, to arrive at its output.

In Chapter 6 we added 3D constrained graph cuts to the voxel classification technique developed in the previous chapter. Graph Cuts enabled us to enforce geometrical knowledge in our tissue classification that could not be achieved by treating each image voxel individually. We proceeded to test our classification pipeline on a set of twelve clinical CT volumes. We were able to classify the periprosthetic tissues we were interested in with sensitivities exceeding 73% and specificities exceeding 96%. While this result was not as good as a human performing the task manually, our algorithm could perform its task on a standard personal computer in half an hour–even in its unoptimised state. This is comparable to the amount of time that would be required by a trained human expert.

In Chapter 7 we improved an existing particle-based multi-material meshing algorithm. This allowed sparser multi-material meshes to be created without having to sacrifice geometric surface accuracy by excessive smoothing. As an example, we showed how our adjusted algorithm was able to generate a tetrahedral multi-material mesh from the six million segmented image voxels of a femur dataset, the output of which consisted of 200,000 tetrahedra that retained the shape of all segmented tissue regions with an average surface error of less than 0.1 mm and never exceeding 3.9 mm at any point.

In Chapter 8 we turned our attention to how one can use the information and models gained from algorithms such as those in the preceding chapters to help a surgeon plan and execute a minimally invasive cement injection operation. 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 feedback we received from the orthopaedic surgeon who used HipRFX during this emulated procedure was positive, citing useful per-operative information that complemented his workflow without taking up unnecessary time or space in the operating room, nor required new operating room hardware.

We conclude this thesis in Chapter 9 with a general discussion and future directions for continuing the work we presented.