

Cover Page



Universiteit Leiden



The handle <http://hdl.handle.net/1887/36019> holds various files of this Leiden University dissertation.

Author: Malan, Daniel Francois

Title: Pinning down loosened prostheses : imaging and planning of percutaneous hip refixation

Issue Date: 2015-10-29

8

A fluoroscopy-based planning and guidance software tool for minimally invasive hip refixation by cement injection

Daniel F. Malan, Stéfan J. van der Walt,
Renata G. Raidou, Bas van den Berg, Berend C. Stoel,
Charl P. Botha, Rob G.H.H. Nelissen, Edward R. Valstar

Adapted from:
International Journal of Computer Assisted Radiology and Surgery,
First Online, August 2015.
doi: 10.1007/s11548-015-1252-8

Abstract

Purpose

In orthopaedics, minimally invasive injection of bone cement is an established technique. We present HipRFX, a software tool for planning and guiding a cement injection procedure for stabilizing a loosening hip prosthesis. HipRFX works by analysing a pre-operative CT and intraoperative C-arm fluoroscopic images.

Methods

HipRFX simulates the intraoperative fluoroscopic views that a surgeon would see on a display panel. Structures are rendered by modelling their X-ray attenuation. These are then compared to actual fluoroscopic images which allows cement volumes to be estimated. Five human cadaver legs were used to validate the software in conjunction with real percutaneous cement injection into artificially created peri-prosthetic lesions.

Results

Based on intraoperatively obtained fluoroscopic images, our software was able to estimate the cement volume that reached the pre-operatively planned targets. The actual median target lesion volume was 3.58 ml (range 3.17 to 4.64 ml). The median error in computed cement filling, as a percentage of target volume, was 5.3% (range 2.2%-14.8%). Cement filling was between 17.6% and 55.4% (median 51.8%).

Conclusions

As a proof-of-concept, HipRFX was capable of simulating intraoperative fluoroscopic C-arm images. Furthermore, it provided estimates of the fraction of injected cement deposited at its intended target location, as opposed to cement that leaked away. This level of knowledge is usually unavailable to the surgeon viewing a fluoroscopic image, and may aid in evaluating the success of a percutaneous cement injection intervention.

Keywords:

Fluoroscopy, DRR, pre-operative planning, hip arthroplasty, percutaneous, X-ray, simulation

Introduction

The long-term survival of hip prostheses are primarily limited by the occurrence of aseptic loosening [Agarwal, 2004]. This pathological process involves extensive resorption of bone adjacent to the prosthesis and its replacement by fibrous tissue that offers little mechanical stability. Minimally invasive cement injection, already an established technique in the orthopaedic field of vertebroplasty [Phillips, 2003], has in recent years been used experimentally to treat hip prosthesis loosening [de Poorter et al., 2008a, Raaijmakers and Mulier, 2010].

Injected cement stabilizes a loosened prosthesis by re-establishing rigid mechanical contact between bone and the existing peri-prosthetic cement mantle [Andreykiv et al., 2012]. Knowing the amount of cement that reaches its intended target may therefore be important in judging the success of a cement injection procedure.

Cement is injected through hollow needles that access the target lesions. Needle insertion may be guided by two dimensional (2D) X-ray fluoroscopic supervision. Using intraoperative Computed Tomography (CT), more accurate guidance of needles' out-of-plane displacement and rotation may be performed [Puri et al., 2006] but this presupposes the availability of CT hardware in the treatment room. More recently, systems have become available that use a combination of pre-operative CT and intraoperative fluoroscopy. One such example is XperGuide (Philips, Best, The Netherlands) [Leschka et al., 2012, Ruijters et al., 2008], where the accurate alignment of the patient, fluoroscopy images and the CT image volume depends on specialized hardware.

Once needles have been placed, cement is injected. During cement injection, the amount of deposited cement can directly be monitored by the position of the syringe's plunger which is marked in millilitres, as shown in Fig.8.1a. The spatial distribution of the injected cement can be estimated by looking at 2D X-ray fluoroscope projections as shown in Fig. 8.1b, where it appears as darkened image areas. The distribution of injected cement may also be monitored using three dimensional (3D) intraoperative CT as shown in Fig. 8.1c. Currently, relying on 2D fluoroscopy for monitoring the flow of injected cement during minimally invasive prosthesis stabilization is the clinical standard [de Poorter et al., 2008a, Raaijmakers and Mulier, 2010].

Cement injection is typically performed over the course of fewer than ten minutes due to polymethylmethacrylate (PMMA) bone cement hardening within




Figure 8.1 – Percutaneous cement injection can be monitored by looking at a) the syringe’s plunger, b) fluoroscopic projections, or c) intraoperative cross-sectional CT slices.

this time span. Of the three intraoperative monitoring methods shown in Fig. 8.1, only CT provides true 3D information. CT cannot be continuously active throughout the procedure due to its higher radiation, as well as the need for the surgeon to have access to the patient. This makes 3 D flow monitoring with CT scan difficult. Lastly, in practice, CT hardware is seldomly available in the treatment room. The patient needs to be slid into and out of the the CT tunnel, and medical staff need to stand clear while the X-ray source is active to reduce radiation exposure. 3D CT can be performed intermittently at best [de Poorter, 2010].

We present a proof-of-concept software tool that enables the volume of injected cement to be quantitatively estimated, using a single pre-operative CT image volume, a pre-operatively defined cement target, and one or more intraoperatively acquired 2D fluoroscopic images. This estimate may be performed for each image as it is read from the fluoroscope. The software is designed to augment specific aspects of a minimally invasive cement injection procedure including planning, execution and analysis.

In this paper we describe our proof-of-concept planning software and then proceed to analyse its use in a pre-clinical cadaver experiment that was performed at our institution. In this experiment, percutaneous cement injection was planned in HipRFX and subsequently performed as for real human patients by an experienced orthopaedic surgeon.

8.1 Materials and Methods

8.1.1 Workflow

The workflow in which our software is to be used is illustrated in Fig. 8.2. Our proof-of-concept software performs steps 3–4 and 9. Steps 2 and 8 are currently performed with human intervention, using external software.

Firstly, our software allows a surgeon to plan a cement injection procedure using a pre-operatively acquired CT as template. Secondly, our software can compute Digitally Reconstructed Radiographs (DRRs) from a pre-operative CT [Galvin et al., 1995]. DRRs are simulated 2D X-ray radiographs, computed from CT. DRRs can be used in guidance, or in estimating cement quantities. Thirdly, our software allows for a rough quantitative estimate of the resulting 3D cement distribution to be made. This is done by comparing the observed intraoperative fluoroscopic images to DRRs that are generated from the pre-operative CT volume.

When using HipRFX, the first step is to load a pre-operative CT volume of the affected hip. The user then specifies a 3D target “mask” that delineates the sub-volume(s) that should be filled with bone cement, i.e. osteolytic lesion(s). In our experiments we used the stand-alone MITK software [Maleike et al., 2009] to perform these segmentations, but any medical volume segmentation tool may be used for this purpose. The segmented cement mask is then read by HipRFX as a binary image volume input file.

Once cement injection targets have been defined, the user has the possibility of adding and positioning virtual cement injection needles that access the designated target(s). If desired, additional segmented anatomical structures may be concurrently loaded and displayed.

From the moment that a CT volume is loaded, our software is able to interactively simulate fluoroscopic images based on a virtual C-arm that may be freely rotated. The goal of this simulation is to subsequently compare it with the actual intra-operative fluoroscopic images. Images may be simulated with or without designated needles and injected cement. By comparing actual intraoperative fluoroscopic images to these simulated images—both in the complete absence, as well as in the complete presence of the intended cement filling—we generate 2D difference images between observed and simulated outcomes. These difference images are transformed to a 2D “cement filling” images that may be back-projected into the 3D volume to create a volumetric cement filling map. For each position in the cement filling image, the filling value’s sensitivity to noise may be computed. For example: regions

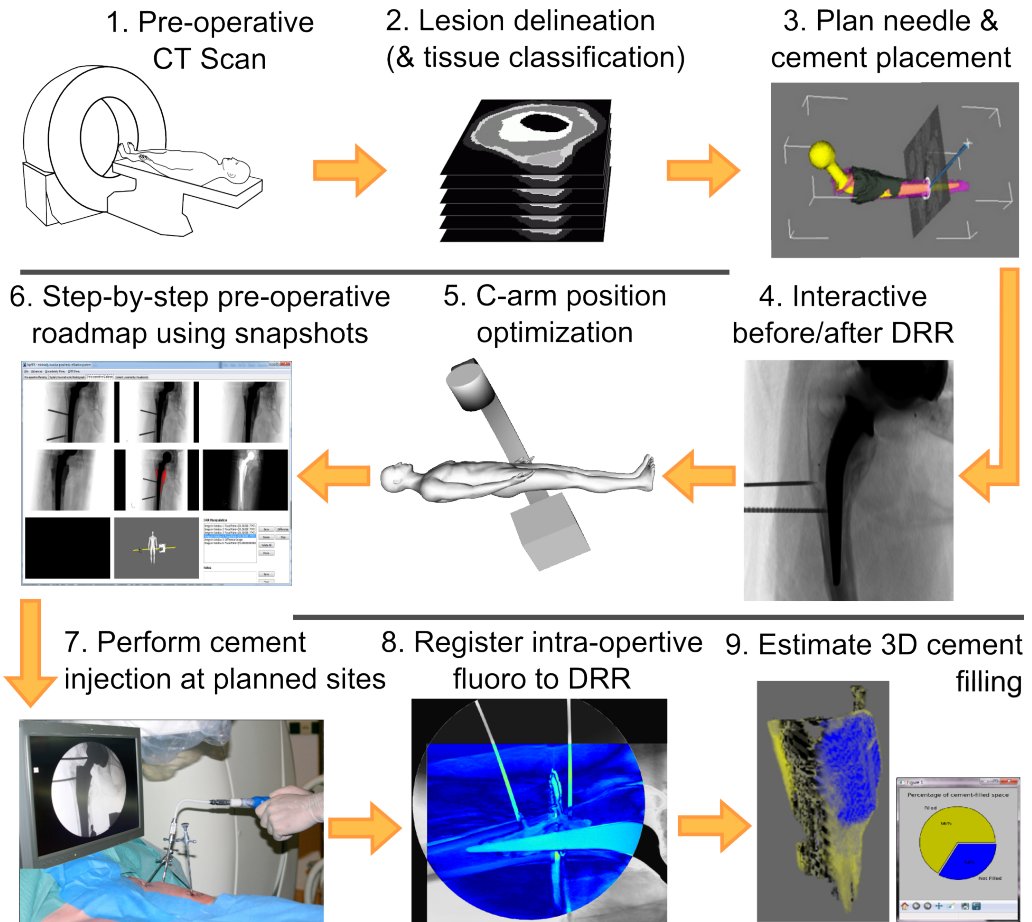


Figure 8.2 – The envisaged workflow of planning and executing a minimally invasive hip refixation using HipRFX. Steps 8 and 9 could conceptually be performed continually, in real time, during the execution of step 7.

that are occluded by a dense metal prosthesis show few image differences between a high and low degree of cement filling. Here, even a small amount of image noise may have an effect on the computed cement filling value. This is further described in Section 8.1.3.2.

8.1.2 Software User Interface

HipRFX is implemented as a module in the DeVIDE Runtime Environment [Botha and Post, 2008] and makes extensive use of the Visualization ToolKit (VTK), NumPy and SciPy [van der Walt et al., 2011].

The functionality contained in HipRFX is split across four separate panels, as explained in Fig. 8.3. Each panel has its own distinct purpose. We now describe the functionality and underlying methods of each panel separately. In its current proof-of-concept form, the image registration and intensity matching is performed by an external tool we built with scikit-image [van der Walt et al., 2014].

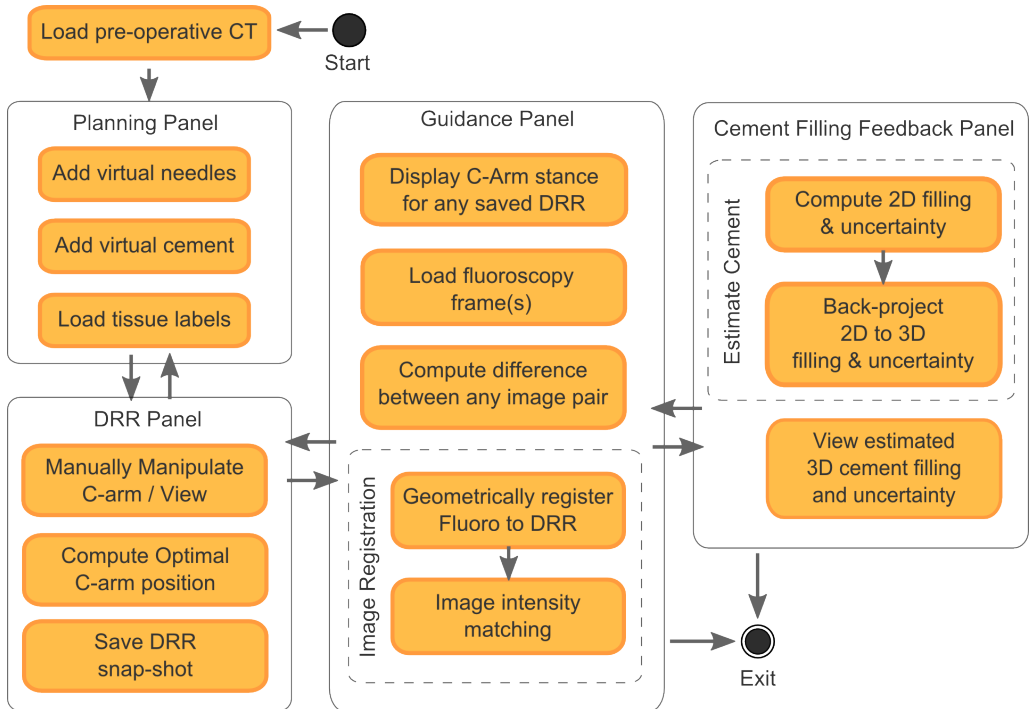


Figure 8.3 – Overview of HipRFX’s main activities, split among its four panels. Components are grouped as functional units.

8.1.2.1 Planning panel

The first view the user is presented with is the planning panel, shown in Fig. 8.4. As soon as a CT volume is loaded, it is rendered in a slice view and as a 3D volume. The panel contains two viewports showing 3D renderings of the hip’s bony structures from anterior-posterior (A) and medio-lateral (B) perspectives. To facilitate the positioning of the needles, the isovalues used for 3D rendering may be adjusted to provide the right amount of context. The central viewport contains a slice-based viewer that can be manipulated interactively (C). The user may load 3D segmented structures of interest. Examples include osteolytic lesions, or arteries and nerves that need to be avoided dur-

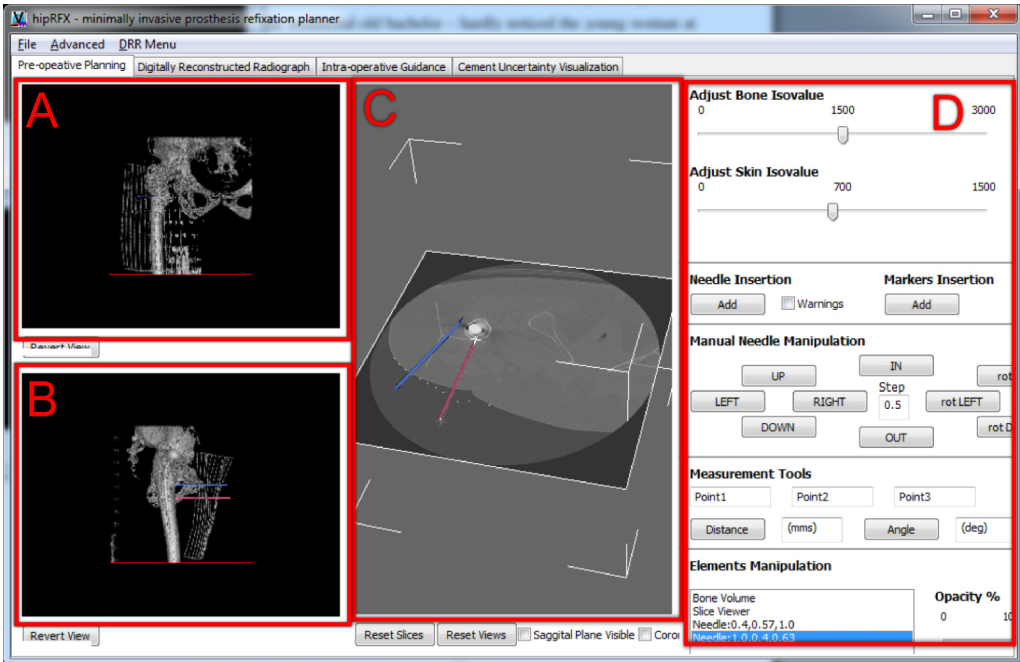


Figure 8.4 – The planning panel provides 3D context renderings (A, B), an interactive slice-view, and control buttons (D). Cement injection needles (shown here in blue and red) can be added to reach desired cement injection targets.

ing surgery. Any number of cement injection needles can be virtually inserted into the CT volume and manipulated. The right-hand panel (D) contains controls to adjust the viewports’ display parameters, to perform distance and angle measurements, and to add, manipulate and remove needles.

8.1.2.2 DRR panel

The Digitally Reconstructed Radiograph (DRR) panel shown in Fig. 8.5 allows a simulated fluoroscope to be interactively viewed and manipulated. The majority of the panel is taken up by the interactively rendered fluoroscopic image (A). Each image frame is created by a ray casting algorithm that directly simulates the propagation of X-rays through the patient as shown in Fig. 8.8 [Bifulco et al., 2002]. The user may manually adjust the brightness and contrast to match the operating room fluoroscope’s settings (B), either to predetermined values obtained from calibration, or visually.

An interactive 3D representation of the patient and the fluoroscopic C-arm is provided (C). The orientation of the C-arm and accompanying fluoroscopic

view may be interactively manipulated to match the operating room's setup.

The image may be inverted to either emulate fluoroscopy or X-ray radiographs—the simulated X-ray attenuation algorithm is identical between these modes. This is an aesthetic choice to be made by the modality that the surgeon is most familiar with.

Needles or segmented cement targets that have been loaded in the Planning View (Fig. 8.4) are realistically overlaid in the DRR image. In this way the operator sees a simulation that corresponds to the fluoroscopic view he/she would see intraoperatively.

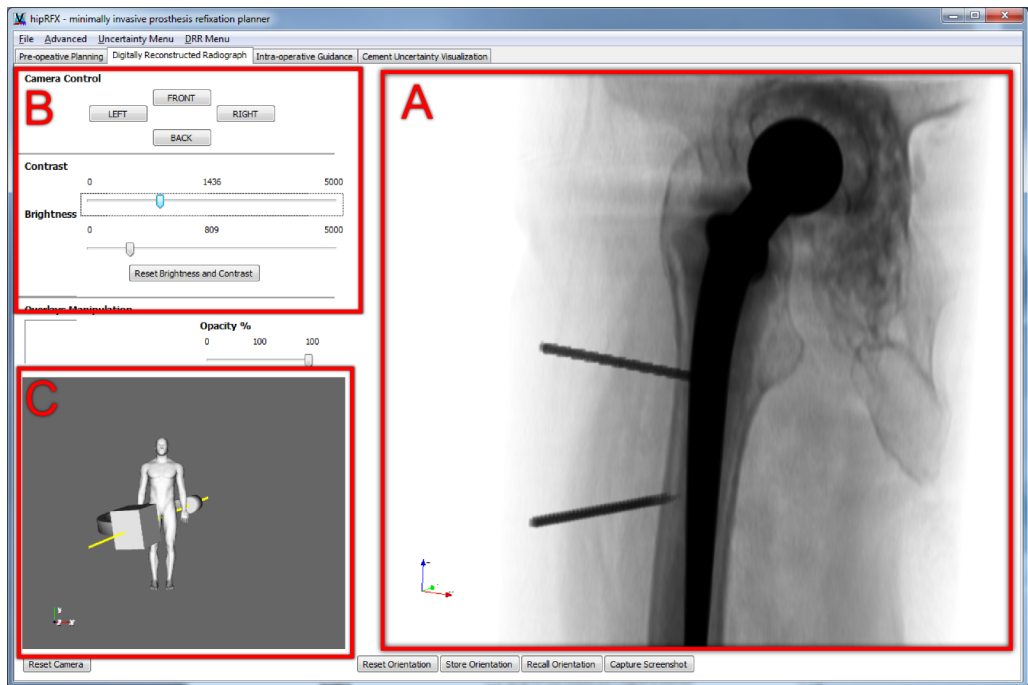


Figure 8.5 – The DRR panel interactively simulates the fluoroscopic view a surgeon would see (A), including the effect of adding needles or bone cement. Snapshots may be saved at the press of a button, for later reference.

Snapshots may be saved to the Guidance Panel along with the accompanying C-arm orientation as shown in Fig. 8.6. These snapshots of expected fluoroscopic views may then be used as a roadmap to guide the surgeon along planned steps during the procedure.

8.1.2.3 Step-by-step Guidance panel

In the DRR Panel shown in Fig. 8.5, the operator is allowed to store a number of DRR snapshots for later reference. Snapshots are displayed in a grid view, and clicking on any snapshot recalls the corresponding C-arm orientation that was used. These snapshots can act as an intraoperative road-map since at any time the operator may visually compare them to the live fluoroscopy image.

Changes in fluoroscopic images may be subtle between steps, e.g. between partial and complete cement filling. The software can highlight the difference between any pair of snapshots—this creates a copy of the image with differences overlaid in blinking red (see Fig. 8.6).

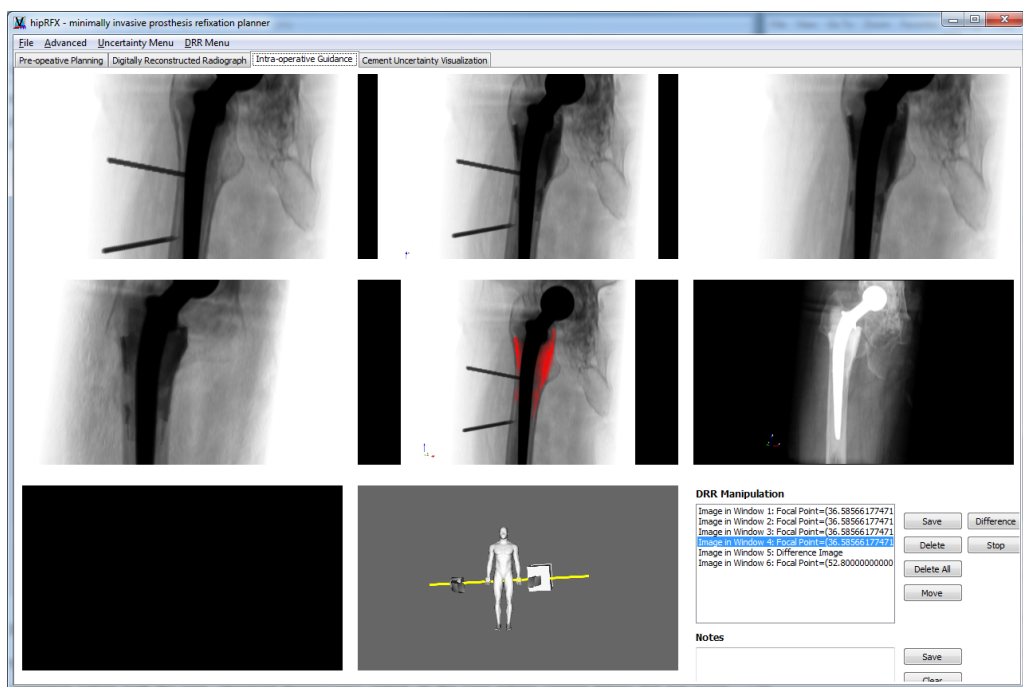


Figure 8.6 – The guidance panel records simulated fluoroscopic snapshots. Differences between frames may be computed and displayed as overlays (shown in red in the centre snapshot). An example of an X-ray-mode DRR is shown in the centre-right. In the centre of the bottom row the fluoroscope C-arm’s position is shown that corresponds to the selected snapshot.

8.1.2.4 Cement filling feedback panel

Differences between real intraoperative fluoroscopic images and DRRs may be analysed to yield estimates of cement filling. Along with filling estimates,

sensitivity to image noise, i.e. “certainty” may be computed. The purpose of the cement filling feedback panel is to visualize these computed values. Using a bivariate colormap similar to those of [Moreland, 2009], we represent certainty with luminance and filling with hue—see Fig. 8.7.




Figure 8.7 – The cement filling feedback panel. A slice through the CT volume is shown on the left, with a 3D rendering of the whole cement target on the right. Yellow indicates target areas that are filled with cement, whereas blue areas indicate a lack thereof. Luminance increases with certainty.

8.1.3 Algorithm

We are able to estimate the volume of injected cement by comparing an observed fluoroscopic image to simulated fluoroscopic images. First we generate a set of two DRRs. The first DRR simulates the fluoroscopic view in the absence of percutaneously injected cement. The second DRR simulates the fluoroscopic view when the target is completely filled with cement, as illustrated in Fig. 8.11. During the cement-injection procedure, real observed fluoroscopic images represent scenarios that lie in between these two extremes. By numerically computing the image intensity differences between the observed

fluoroscopic image and the two DRRs, our software numerically estimates the amount of cement that was injected.

8.1.3.1 Computing the amount of cement

The DRR image formation process is illustrated in Fig. 8.8. Every element in the CT volume has at least one X-ray path that passes through it to the DRR image plane. The DRR formation is a discrete approximation of the physical X-ray imaging process, and is explained in Appendix A.




Figure 8.8 – DRRs are formed by simulating the propagation of X-rays through a CT image volume

We can estimate the total volume of cement in the 3D CT by estimating the distance that each ray travelled through cement on its path from X-ray source to the image. The distance that each ray travelled through cement can be expressed as the distance that it travelled through the target, multiplied by the fraction of this distance that was filled with cement. We call this fraction the “fill fraction”, $F = s_c/s_t$ as shown in Fig. 8.9.

In Appendix A we explain the linear relationship between the length of an X-ray path through a uniform substance and the cumulative absorption coefficient μ_{ray} by which the ray is attenuated. We can therefore rewrite the fill fraction as $F = \mu_c/\mu_t$, where μ_c is the cumulative attenuation coefficient contributed by the cement-filled portion of the target. μ_t is the cumulative absorption coefficient that the target would have contributed if it was completely filled with cement. From the superposition principle—also explained

in Appendix A—we can express μ_c as the difference between the cumulative attenuation coefficients of the entire ray and of the ray portion that passes through no cement, i.e. $\mu_c = \mu_{ray} - \mu_{none}$. Combining all of the above, we can rewrite the fill fraction as

$$F = \frac{\mu_{partial} - \mu_{none}}{\mu_{full} - \mu_{none}}, \tag{8.1}$$

where the subscript “partial” refers to the actual partially filled observed volume, as may be observed intraoperatively. The subscript “none” refers to the case where cement is completely absent, and “full” refers to the case where complete cement filling of the target is achieved.

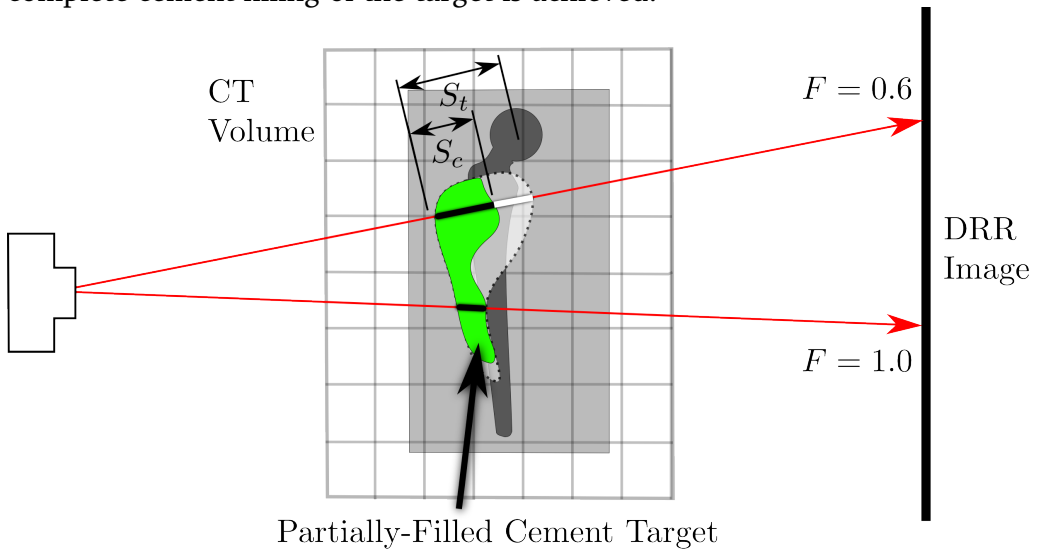


Figure 8.9 – The filling fraction $F = s_c/s_t$ is the distance that each ray passes through cement, expressed as a fraction of the target area’s thickness along that ray.

From Appendix A we know how to relate attenuation coefficients μ_{ray} to fluoroscopy image intensity I , allowing us to rewrite Eq. 8.1 in terms of image intensities:

$$F = \frac{-\log(I_{partial}) + \log(I_{none})}{-\log(I_{full}) + \log(I_{none})}$$

which becomes

$$F = \frac{\log(I_{partial}/I_{none})}{\log(I_{full}/I_{none})}. \quad (8.2)$$

I_{none} and I_{filled} refer to the fluoroscopy images seen when either no cement is injected and when the target volume is completely cement-filled. In an intra-operative setting one never has access to I_{filled} while I_{none} could be recorded at the start of the procedure. We chose to simulate both I_{none} and I_{filled} by using our DRR algorithm on the pre-operative CT volume. Note that $I_{partial}$ corresponds to the observed intraoperative fluoroscopy image that is available continuously during the cement injection procedure. An example of I_{none} , I_{filled} and $I_{partial}$ is shown in Figure 8.11. This triad of images allows us to compute F for all pixels in the 2D fluoroscopy image where these values are defined, as shown in Fig. 8.13. The value of F for an arbitrary pixel in the image is denoted by $F_{x,y}$.

To obtain a quantitative cement volume estimate, we need to weigh each pixel's $F_{x,y}$ with the thickness s_t of the cement target at the point through which the associated ray passed. We call these the "projection weights"—illustrated in Fig. 8.10. The projection weights correspond directly to s_t in Fig. 8.9, albeit appropriately scaled to yield a millilitre-valued output.

The projection weights W are directly proportional to the attenuation coefficients that a completely cement-filled target would contribute to the rays terminating on each pixel. This is illustrated in Fig. 8.10 and can be mathematically expressed by referencing Eq. 8.7 in Appendix A. Here I_{cement} is the image that would result from X-rays passing only through the completely filled cement target and nothing else, as shown in Fig. 8.12. Mathematically:

$$W \propto s_t \propto \mu_t = -\frac{1}{k} \log I_{cement}.$$

As $F = s_c/s_t$, it follows that $W \cdot F \propto s_c$. Looking at Fig. 8.9 we note that summing s_c for every ray through the CT volume provides an approximation to the amount of cement contained in the partially filled target.

We can now compute an absolute quantitative estimate of the volume of the filled portion of the cement target. This is done by calculating the weighted




Figure 8.10 – The projection weights scale linearly with the thickness of the cement through which rays pass to reach the image plane.

sum of the fill fractions over all image pixels, and multiplying the result with the volume of the complete cement target, in millilitres:

$$V_{filled} = \frac{\sum_{x,y} (W_{x,y} \cdot F_{x,y})}{\sum_{x,y} W_{x,y}} \cdot V_{target} \quad (8.3)$$

A necessary prerequisite for computing Eqs. 8.2 and 8.3 is that all relevant DRR and fluoroscopy images are registered, i.e. that their field of view are the same, that they have the same brightness and contrast, and that their subject is in the same position. In our proof-of-concept system and in our experimental set-up, calibration was approximated visually. In a practical clinical system such pre-operative registration should instead be automatically and robustly implemented, and orientation recorded via angle encoders attached to the imaging hardware. To correct for mismatches in our experiments, we applied suitable rigid similarity transformation that corrects for scale, rotation and translation in each fluoroscopic image. These parameters were estimated from four manually selected point correspondences in each image pair. Brightness and contrast were corrected by linearly adjusting the intensity of the fluoroscopic image to match that of the DRR in a least squares sense. Additional caveats apply. The cement fill fraction F in Eq. 8.2 is only defined in image regions onto which the cement target project, and only where the




Figure 8.11 – a) Pre-operative DRR showing no injected cement. b) Post-operative DRR showing complete cement filling of the intended cement target. c) Intraoperative fluoroscopic image showing partial cement filling.




Figure 8.12 – a) The “cement-only” image I_{cement} . b) The subset of the image over which F can be computed using Eq. 8.2.

image brightness is not fully saturated. This is illustrated in Fig. 8.12. All other areas in the image are marked as non-computable. Non-computable areas include those where the thick metal prosthesis completely absorbs the X-ray beam, as this would result in a zero-valued denominator I_{none} in Eq. 8.2. In non-computable areas falling inside the domain of I_{cement} , F is interpolated using thin plate spline radial basis functions, as shown in Fig. 8.13. In our experiments we manually delineated suitable regions within the domain of




Figure 8.13 – a) The cement filling image F from Eq. 8.2. b) F , after its interpolation to the area spanned by I_{cement} . c) The certainty image C computed using Eq. 8.4.




Figure 8.14 – A mask were used to designate a valid subset of the computed cement filling projection in our experimental data. a) Manually selected regions are superimposed in white, and exclude confounding objects such as metal needles. b) Input to the interpolation algorithm. Non-computable regions are shown with a diagonally hatched pattern. c) The interpolated output.

I_{cement} to serve as input to the interpolation algorithm, thereby excluding areas where needles were present or where cement leaked into the incision created during preparation of the test femurs. This is shown in Fig. 8.14.

Once F has been computed over the whole 2D projected cement filling domain, it may be back-projected along the ray path to 3D and visualized in in the cement filling feedback panel as shown in Fig. 8.7.

8.1.3.2 Sensitivity of the output to image noise

Fluoroscopy images contain some degree of image noise that may distort the derived cement filling values. In areas occluded by radio-dense materials such as the prosthesis, little or no information about cement filling can be deduced. The uncertainty inherent in the resulting cement filling computation can be explicitly analysed. This uncertainty is essentially equal the degree that the result is affected by noise in the input image. Where dense metal objects like the prosthesis occlude the image, the useful signal is attenuated while image noise remains constant. This results in a very low signal to noise ratio, resulting in unreliable cement filling estimates. For each pixel, the sensitivity of the filling fraction F to noise can be approximated as its derivative with respect to the observed fluoroscopic image intensity:

$$\begin{aligned} S &= \frac{\partial F}{\partial I_{observed}} \\ &= \frac{\partial}{\partial I_{observed}} \frac{\log(I_{observed}/I_{none})}{\log(I_{full}/I_{none})} \\ &= \frac{1}{I_{observed} \log(I_{full}/I_{none})}. \end{aligned}$$

S is only defined for the regions where F is directly computable. The reciprocal of the sensitivity function defines what we call the “certainty image” C :

$$C = I_{observed} \log(I_{full}/I_{none}). \quad (8.4)$$

All non-computable pixels in C are set to zero. An example of the certainty image is shown in Fig. 8.13.

As was the case for F , we may re-project C along the projection rays and visualize it in three dimensions on the cement filling feedback panel. F and C are then combined in a single bi-variate colour mapping shown in Fig. 8.7. The sensitivity of the overall cement volume estimate V_{filled} may be computed by applying the weights W to each individual pixel’s filling value sensitivity S in the same way as it was done in Eq. 8.3:

$$S_{filled} = \frac{\sum_{x,y} (W_{x,y} \cdot S_{x,y})}{\sum_{x,y} W_{x,y}} \cdot V_{target}. \quad (8.5)$$

8.1.4 Validation Experiment on cadaver specimens

Substituting real fluoroscopic images with simulated DRRs allowed for perfect image registration, as all the required images could be generated using the exact same projection parameters. However, to test our software in a physically representative workflow using realistic clinical hardware, we enacted minimally invasive hip re-fixation procedures on five ex-articulated cadaver legs. The workflow we followed mirrors that shown in Fig. 8.2. An experienced orthopaedic surgeon was asked to perform the minimally invasive cement injection—a procedure with which he was familiar.

8.1.4.1 Preparation of cadaver specimens

Five ex-articulated cadaver legs were obtained. Each leg included the whole femoral region, from the superior to the inferior epiphysis. In order to simulate peri-prosthetic lesions that could be filled with cement, we used a similar approach as that of [Kraaij et al., 2012]. Two experienced orthopaedic surgeons placed cemented polished Exeter stems (Stryker, Limerick, Ireland) in each femur. Exeter stems of sizes 3 and 4 (both with an offset of 44 mm) were used, with the most appropriate size chosen for each leg.

Since we used cadaveric legs which were exarticulated in the hip joint, they provided free access to the femoral head and neck. The soft tissues of each leg were kept intact—this was important so as to provide a realistic target for percutaneous cement injection.

After each prosthesis was placed and the cement hardened, the prosthesis was again removed. The straight-edged and smoothly polished wedge-shaped Exeter prosthesis' surface enables such removal without damaging the cement mantle [Malan et al., 2012b]. After removing the prosthesis, a single cut through the skin and muscle tissue onto the bone surface was made with a dissecting knife, perpendicularly to the femur shaft. The femur itself was then sawed through with a bone saw. The leg was then tilted open along the cut line to allow access to the femur shaft where it was bisected. Lesions were created using an abrasive drill bit in an identical way as was previously performed by [Malan et al., 2012b]. The muscle tissue on the opposite side of the incision was kept intact, preventing the two semi-bisected halves of the leg to separate completely. After lesions were created, the prostheses were re-inserted. The snug fit between prosthesis and cement mantle ensured that the original alignment of the two partially bisected leg halves was restored in each case. The soft tissue around the incision perimeter was sewn closed so that the soft tissue could recover some degree of conformity as well.

8.1.4.2 Pre-operative planning

After preparation, each femur was CT scanned with its prosthesis in place. This CT image corresponded to the diagnostic pre-operative CT scan that a patient would routinely have performed if he/she was eligible for minimally invasive hip prosthesis refixation [de Poorter et al., 2008a, Egawa et al., 2009].

The peri-prosthetic lesions of each leg were manually segmented in 3D from the CT volume using the Medical Imaging Interaction Toolkit (MITK 0.12.2), an interactive medical image segmentation software tool [Maleike et al., 2009]. These segmented lesions defined the cement injection targets that we wished to fill in the subsequent minimally invasive cement injection step.

Following the workflow described in Fig. 8.2, we used HipRFX to read the CT image volume and the manually delineated segmentations, and then to virtually place vertebroplasty needles to reach these lesions. We manually chose applicable angles for orienting the C-arm fluoroscope, after which we saved DRR snapshots of the expected view, both before and after cement injection.

8.1.4.3 Cement injection

An experienced orthopaedic surgeon who has previously performed percutaneous cement injection in more than thirty patients, performed the procedure on each of the prepared cadaver legs. Standard clinical grade needles (Biomet VerteShark Access, 11Gx15cm), vertebroplasty cement (Biomet Bone Cement V) and mixing sets (Optivac Procedure Set) were used—all supplied by Biomet Europe BV (Dordrecht, The Netherlands). Intraoperative fluoroscopy was performed using a standard clinical C-arm fluoroscope (Philips BV Pulsera , Best, The Netherlands).

As opposed to a traditional percutaneous procedure where no pre-operative guidance was available, an assistant held a 10.1 inch tablet-computer (Asus TF700T, Taipei, Taiwan) in view of the surgeon on which the precomputed DRR snapshots were displayed. A photo of one of our experiments is shown in Fig. 8.15.

For each cadaver, the C-arm was moved to the appropriate position as indicated by the planning snapshot. Vertebroplasty needles were subsequently inserted as per pre-operative plan. As opposed to patients, the bones of the cadavers were not affected by osteolysis and necessitated insertion of the needles by first drilling into the bone using a thin drill bit. Even with pre-drilling, inserting the needles into the bone required a surgical hammer to supply sufficient force. In four of the five legs both planned needles could be inserted,



Figure 8.15 – An assistant shows a simulated fluoroscopic image to the surgeon. Intraoperative fluoroscopic images are visible in the background. The cadaver leg is situated on the table, underneath the C-arm fluoroscope. The faces of the surgeon and assistant were anonymized for publication.

while needle breakage limited use to inserting a single needle in the remaining leg.

Polymethylmethacrylate (PMMA) bone cement was then mixed and injected under fluoroscopic C-arm guidance. The cement flow was fluoroscopically monitored in real-time, as is usually done in this kind of procedure [de Poorter et al., 2008a, Raaijmakers and Mulier, 2010]. Cement injection was continued until the peri-prosthetic space appeared filled or until the cement started leaking into the surrounding tissue.

8.1.4.4 Acquisition of ground truth data

Fluoroscopic images obtained during our cadaver experiment were captured and saved to file. HipRFX was then set to simulate an equivalent projection

viewpoint as used with the intraoperative C-arm fluoroscopy. Matching was performed manually, interactively and according to visual similarity, as the C-arm fluoroscope we used did not record the exact projection angles—neither visually nor in accompanying metadata. Our only guidelines to ensure consistency were the apparent agreement between the virtual C-arm positions and those used for the experiment, subjective agreement between the planning DRRs and those observed during the procedure, and the knowledge that HipRFX's theoretical focal length and field of view matched those of the C-arm..

Two simulated DRRs were created for each real fluoroscopic image: one representing the pre-operative state, and a second representing ideal cement filling of the target lesions that were segmented by hand in the pre-operative CT volume. We used the procedure described in section 8.1.2.4 to estimate cement filling.

After the cement had hardened, the cadaver legs were all returned to storage. A post-operative CT scan of each treated cadaver leg was performed. This step is not part of the workflow depicted in Fig. 8.2 but allowed for accurate and independent post-operative cement volume measurements to be performed, using manual delineation in the MITK software [Maleike et al., 2009, Malan et al., 2012a].

8.2 Results

Fig. 8.16 shows a side-by-side comparison between a fluoroscopy image obtained during the cadaver experiment and a corresponding HipRFX-generated DRR. Both of these images were generated using the actual experimental attained cement distribution – the fluoroscopic image created at the end of the cement injection, and the DRR generated using the post-operative CT volume. When the fluoroscopic image was made the cement had not yet hardened. Some changes to the cement distribution are visible between the images, especially where cement leaked outside the femoral shaft. The physical incision that resulted from our experimental preparation method can be seen as the radiolucent band into which the cement leaked.

We observed that much of the leaked cement did so into the incisions created by our experimental method. This is clearly visible in Fig. 8.16. We compared all computed values with ground truth values measured in post-operative CT. With a median of 81% (range 55%-84%), the cement leakage in our experiment was higher than observed in real patients [Malan et al., 2014]. The filling percentage (median 52%, range 18%-55%) was compa-




Figure 8.16 – a) Real fluoroscopy image during cadaver experiment vs. b) HipRFX-generated DRR from post-operative CT with overlaid virtual needles. A region of cement leakage is indicated by white arrows. The incision created during preparation is visible as a light-coloured line extending into the zone of cement leakage.

able to that observed in real patients. The properties of the experimentally created cadaver leg lesions and the subsequent injected cement volumes are shown in Table 8.1.

Target volume (ground truth)	Injected cement (incl. leaked)	Cement filling of target (ground truth)	Filling estimate using HipRFX	Cement vol. estimation error as % of target
3.58 ml	11.64 ml	51.8%	46.2% 50.3%	5.6% 1.6%
3.20 ml	8.90 ml	52.0%	46.8%	5.2%
3.17 ml	3.75 ml	33.4%	31.2%	2.3%
4.64 ml	4.69 ml	17.6%	32.5% 21.8%	14.9% 4.2%
4.09 ml	5.03 ml	55.4%	65.3%	9.8%

Table 8.1 – Cement filling estimates for the five femora using HipRFX. Ground truth values were obtained by directly measuring them from independent post-operative CT volumes. .

For each real fluoroscopic image, HipRFX was used to estimate the amount of cement that was injected into the target region. Two of the five legs were each imaged from two different angles. The median estimation error, expressed as

a percentage of the cement target volumes, was 5.2% (range 1.6%-14.9%). These results are also summarized in Table 8.1.

To examine the role that experimental inaccuracies and assumptions had on our results, we further differentiated the absolute per-pixel errors made in the 2D “filling images” like those in Fig. 8.13. We compared results when using either a) our real experimentally obtained fluoroscopic images that include unaccounted-for cement leakage, image noise and possible geometric-and-image-intensity registration errors, b) simulated fluoroscopic images that include unaccounted-for cement leakage but with perfect registration, and c) simulated fluoroscopic images with no cement leakage. For each of these scenarios we furthermore distinguished between the estimation error made in the “computable” and the “non-computable” interpolated regions that were described and illustrated in Fig 8.14. Results are shown in Fig. 8.17.

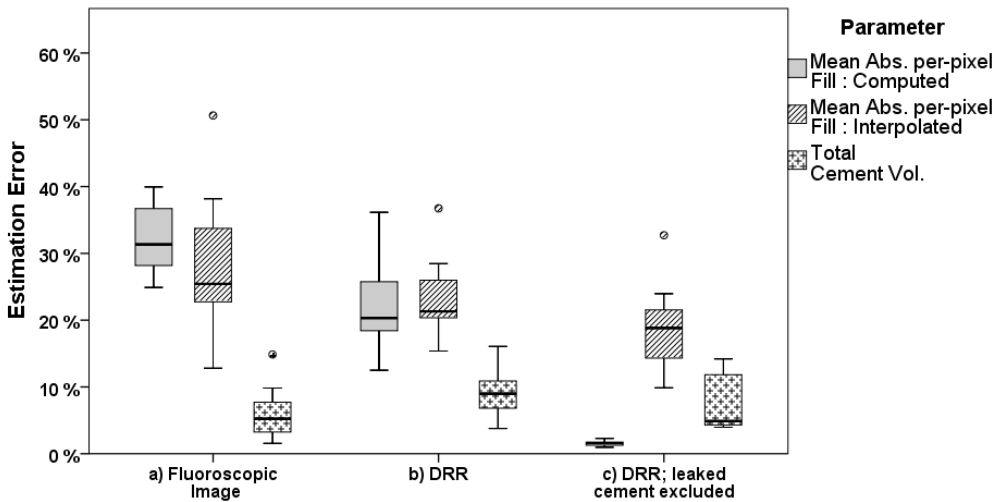


Figure 8.17 – Experimentally obtained cement filling estimation errors using HipRFX. a) errors when analysing real fluoroscopic data compared to b) the equivalent errors when substituting real fluoroscopic images with simulated fluoroscopic images. In c) cement leakage was omitted from the simulated images. Mean absolute per-pixel errors as well as the final integrated cement volume estimate errors are shown.

The simulated fluoroscopic images for (b) were DRRs that used the complete post-operative 3D cement distribution, segmented in the post-operative CTs. The simulated fluoroscopic images for (c) were identical to (b) except that it excluded all leaked cement. As our cement filling algorithm does not model

cement leakage this modification meets the assumption inherent in our cement estimation algorithm, which is that all injected cement projecting onto the image of the cement target must also be located within the cement target. After performing the experiment, we discussed our proof-of-concept software with the orthopaedic surgeon who participated. He valued the pre-operative planning capabilities that the software offers without necessitating additional or new clinical hardware. He stated that HipRFX represented a promising development in image-based guidance, and also represents a step towards reducing the need for intraoperative CT imaging.

8.3 Discussion

Minimally invasive hip refixation by cement injection is a novel and experimental technique for refixing loosened orthopaedic implants. To our knowledge there currently exists no task-specific tool for planning or performing this procedure. Our experimental HipRFX system has been purpose-designed for this task and uses existing fluoroscopic hardware to provide intraoperative guidance and to analyse the distribution of injected cement. HipRFX can simulate fluoroscopic images and computes cement volume estimates and a combined filling-and-uncertainty distribution. Should it be used intraoperatively, our system could provide a surgeon with valuable extra information that currently is either absent, or can only be obtained after the fact. HipRFX focuses on simulating and analysing the fluoroscopic images a surgeon would see — including the realistic fluoroscopic appearance of injected bone cement. This allows us to progress beyond needle guidance and to also simulate and assess the cement filling.

By visually judging the agreement between a simulated desired outcome and real intraoperative images, a surgeon could be aided in judging whether sufficient cement penetration has been achieved. An obvious deviation between simulated and observed fluoroscopic images may alert him to unwanted outcomes, such as cement leakage. In addition, we experimentally showed that HipRFX can compute numerical estimates of the injected cement volume.

Our pre-operative software tool is built around the assumption that C-arm fluoroscopy is the dominant intraoperative technology for monitoring percutaneous cement injection, and will most likely remain so for the foreseeable future. Additionally, we assume that at least a single pre-operative CT image volume will be available before commencement of surgical planning. These two assumptions are compatible with existing radiological protocol for mini-

mally invasive cement injections [de Poorter et al., 2008a, Raaijmakers and Mulier, 2010].

The Philips XperGuide [Fichtinger et al., 2005, Leschka et al., 2012] system shows some similarities with our approach to intraoperative guidance as it also bases its guidance on a prior CT volume, and also combines this with live single-plane C-arm fluoroscopy. However, XperGuide is focused on guiding needle biopsies only — and achieves this by overlaying glyphs onto radiological images. Unlike XperGuide our system is, in principle, compatible with all existing CT imagers and fluoroscopes. This was demonstrated in our cadaver experiment where we used a commercial C-arm fluoroscope that precluded any any direct calibration between our software and the imaging hardware. Here we showed that HipRFX is capable of estimating the volume of cement to have reached pre-operatively defined target lesions. While we used post-operative CT as ground truth to compare our results to, post-operative CT does not form any part of our workflow as shown in Fig. 8.2.

Where they project to the same pixel positions, our algorithm cannot distinguish leaked cement from cement inside the target region. A large amount of leaked cement may therefore negatively affect the accuracy of the cement filling estimation, whereby the algorithm interprets leaked cement to be located in the target lesion. This view is supported by Fig. 8.17 that shows much reduced estimation errors when we analysed DRRs where leaked cement was digitally removed.

An independent analysis of real patient data showed that cement leakage is prevalent in clinical practice [Malan et al., 2014]. Given this reality of cement leakage, HipRFX plays the important role of specifically estimating the injected cement fraction that reaches a target lesion. This fraction is presumed to be the one that directly contributes to the stability of a hip prosthesis [Andreykiv et al., 2012]. Cement leakage was exacerbated in our experiments by the cadaver legs having been sawed through, thereby providing an easy escape route for injected cement. Even in this challenging scenario the median cement volume estimate error was 5.2% (range 1.6%-14.9%). We are encouraged by this result, especially considering the limitations of our experimental set-up.

Residual image registration mismatches limited the accuracy with which our algorithm could deduce cement filling. From Fig. 8.17 we made the paradoxical observation that, for the fluoroscopy images, the estimation error was higher in computable than in interpolated regions. We explain this by random image noise being present in fluoroscopic images that cause per-pixel

fluctuations in filling estimation. By contrast, the interpolated regions are smooth and show a lower per-pixel absolute error. This explanation is supported by the observation that the errors in the computable regions decreased when we substituted the fluoroscopic images with DRRs, while the interpolated regions' error decreased by less. We observed that per-pixel estimation errors tended to cancel out when they were summed, resulting in a reduced estimation error for the total cement volume.

A limitation of HipRFX in its current form is its inability to directly record the position and projection parameters of the relevant fluoroscope. In fact, in our experiments these parameters were not numerically recorded at all. The simulated C-arm was visually and manually aligned to match the DRRs with the recorded fluoroscopic image. The inevitable discrepancies that resulted from such qualitative alignment required us to perform image registration, as schematically illustrated in Fig. 8.2. Iterative techniques that match radiographs with CT-derived DRRs [Guéziec et al., 1998, Russakoff et al., 2005] may in future improve registration accuracy. Calibrated systems using accurate angle encoders such the Philips XperGuide may in future be constructed for our application and make the system easier to use and more robust.

In clinical practice, needle placement is mainly performed under the supervision of single-plane fluoroscopy [Raaijmakers and Mulier, 2010, de Poorter et al., 2008a]. Intraoperative CT may also be used for more accurate 3D guidance. Similar to Philips XperGuide our software has the potential to be developed to guide a surgeon during needle placement. One possibility would be to simulate a correctly placed needle and overlay it in live fluoroscopy. We leave the exploration of this topic to future work.

Our system may in future be used to compute optimal C-arm orientations for performing an intervention. This can be done by maximizing the image difference between DRRs of the expected pre-operative and post-operative situation. Large image differences indicate that the changes caused by the planned procedure are clearly visible — this then represents an informative viewing angle. Differences may enable discernment between correct and incorrect needle placement, or between sufficient and insufficient cement filling. The parameter space that needs to be traversed in this case consists only of the C-arm's elevation and azimuth — i.e. a two dimensional optimization problem. When a C-arm fluoroscope needs to be positioned accurately to within 5 degrees, all possible orientations may be examined in fewer than 800 iterations. Given the speed with which DRRs and difference images can be computed, this brute force approach would be feasible for optimizing the view for each

desired step in the cement injection procedure, while being guaranteed to find the global optimum.

Another direction for future work would be to combine the estimates obtained from several distinctly oriented fluoroscopy images. A stereo fluoroscopy system [Baka et al., 2011] may be used for this purpose, or several single plane fluoroscopic images may be taken successively from different angles. The area occluded by a large metal prosthesis would be different for each image and the combined estimate may be considerably more accurate than either estimate on its own.

We are encouraged by the feedback received from the orthopaedic surgeon who used HipRFX in planning and executing the cadaver experiment. We foresee that, in future, the techniques described in this paper could provide surgeons with sufficient insight and feedback on percutaneous cement injection procedures to avoid the use of explicit 3D imaging tools in the operating room. We believe that the ability to glean quantitative estimates of the cement volume that is deposited in a desired target region, without having to resort to 3D imaging modalities, is novel. This approach may be applied to any application where radiopaque cement is injected, including vertebroplasty.

Acknowledgements

We wish to sincerely thank Fred van Immerseel, Dr Marc Nieuwenhuijse and Dr Huub van der Heide who assisted us in preparing the cadaver specimens. Also Noeska Smit, who operated the C-arm fluoroscope during the cadaver experiments.

Appendix A: DRR generation by ray casting

Assuming a bundle of mono-energetic X-ray photons with mean energy E and no scattering, the X-ray flux X that reaches a detector after passing along path s is:

$$X = X_0 e^{-\int_0^D \mu(s,E) ds}, \quad (8.6)$$

where X_0 is the source's emitted X-ray flux, D is the length of the path that the X-ray photons travel to reach the detector, and $\mu(s, E)$ is the X-ray absorption coefficients of the matter that the X-rays

pass through [Kalender, 2005]. It is important to note that the superposition principle holds for X-ray absorption coefficients; if an X-ray bundle successively passes through two material regions with X-ray absorption coefficients $\mu_a(s)$ and $\mu_b(s)$ respectively, then Eq. 8.6 yields $X = (X_0 e^{-\int_0^a \mu_1(s) ds}) \cdot e^{-\int_a^b \mu_2(s) ds} = X_0 e^{-(\int_0^a \mu(s) ds + \int_a^b \mu_2(s) ds)}$, showing that the equivalent total absorption coefficient is the sum of the two individual absorption coefficients. Consequently the equivalent X-ray absorption coefficient for an entire ray, $\mu_{ray} = \int_0^D \mu(s, E) ds$, varies linearly with the distance the ray passes through a uniform substance. If we know the values of μ_{ray} for two X-ray paths passing through different thicknesses of cement, we would be able to relate these thicknesses to one another.

Because X-rays are photons, it comes as little surprise that Eq. 8.6 is identical to the light absorption term used in direct volume visualization, a well established 3D image rendering technique that simulates light propagation through physical objects [Preim and Botha, 2013a]. This allows for the X-ray imaging process to be approximated with efficient GPU-optimized ray tracing algorithms that have become an essential part of modern medical visualization software [Preim and Botha, 2013b].

The required X-ray absorption μ that needs to be plugged into Eq. 8.6 can be directly read from a CT volume. The constituent volume elements ("voxels") in a CT volume are quantified in Hounsfield Units (HU). The Hounsfield Unit is dimensionless and is defined by the physical X-ray absorption coefficient of the material that was imaged, as well as that of water and air:

$$CT = \frac{1000(\mu - \mu_{water})}{\mu_{water} - \mu_{air}} \approx \left(\frac{1000}{\mu_{water}} \right) \mu - 1000$$

This relationship between CT number and X-ray absorption coefficient is linear, allowing Eq. 8.6 to simulate X-ray absorption when the absorption coefficient is set to $\mu = \frac{\mu_{\text{water}}}{1000}(CT + 1000)$ for every point along the ray's path.

To understand how one may estimate cement volumes from numerically comparing fluoroscopy images it is also necessary to understand the relationship between X-ray flux and fluoroscope image brightness. The transmissivity of a developed X-ray film is defined as $T = I/I_0$ at every point, where I can be interpreted as the intensity of light emanating from a position on the observed image when it is back-lit by a light-source with intensity I_0 . The blackening of the X-ray film or digital detector's output image as it is exposed to X-rays is called its optical density, and is defined as the base-ten logarithm of the reciprocal of its transmissivity, i.e. $D = \log_{10}(I_0/I)$.




Figure 8.18 – Optical film's characteristic curve also known as the H&D curve

Over a certain working range, an X-ray detector's optical density has a linear relationship to the logarithm of received photon flux. This is illustrated in the detector's characteristic H&D curve as illustrated in Fig. 8.18. The slope of this linear relationship is called the detector's Gamma (Γ).

Medical fluoroscopic images are displayed as the “negative” image of traditional X-ray radiographs, with the background being white and radiopaque structures such as bone being dark. This is also the default display mode we used in generating DRRs, in which case the characteristic curve is the vertically mirrored version of the H&D curve in Fig. 8.18, i.e. with a linear region that has a slope of negative Γ , or

$$\begin{aligned}
\log_{10} \frac{I_0}{I} &= -\Gamma \log_{10} X + c \\
\therefore \ln \frac{I_0}{I} &= -k \ln X + p \\
&= -k \ln(X_0 e^{-\int_0^D \mu(s,E) ds}) + p \\
&= -k \cdot (\ln X_0 - \int_0^D \mu(s,E) ds) + p \\
&= k \cdot \int_0^D \mu(s,E) ds + (p - k \cdot \ln X_0) \\
&= k \cdot \mu_{ray} + q
\end{aligned}$$

, where c , k , p and q are placeholders for constant values and μ_{ray} represents the integral of the X-ray absorption coefficients along the whole path which the X-ray beam passed from the source to the detector. In a digital fluoroscopy image, the image brightness and contrast may be adjusted arbitrarily. To be accurate, our cement-filling algorithm requires that the image intensities are scaled so that I_0 has a value of 1.0 and q is zero. This can be achieved by automatically computing the appropriate image brightness and contrast. Under the aforementioned conditions this implies that:

$$\ln I = -k \cdot \mu_{ray} \tag{8.7}$$

where the image I is interpreted as having an intensity range of $[0, 1]$.

