



Image Analysis for Extracapsular Hip Fracture Surgery

A Thesis submitted to
Newcastle University
for the degree of
Doctor of Philosophy (PhD)
in the Faculty of Science Agriculture and Engineering (SAGe)

Akachukwu B. Okoli

School of Engineering
August 2018

Contents

Publications	ix
Acknowledgements	xi
Dedication	xiii
Abstract	xv
Declaration	xvii
Glossary	xix
1 Introduction	1
1.1 Motivation	3
1.2 Aim and Objectives	4
1.3 Contributions	5
1.4 Thesis Outline	7
2 Literature Review	9
2.1 Deformable Models	9
2.2 Constrained Local Models	13
2.2.1 PDM-based CLM	14
2.2.2 The Generalised CLM	15
2.2.3 Texture Models in Literature	17

Contents

3	Technical Background	19
3.1	Statistical Analysis of Shape	19
3.1.1	What is shape?	20
3.1.2	Point Distribution Models	21
3.1.3	Shape Alignment	23
3.1.4	Principal Component Analysis (PCA)	24
3.1.5	Probabilistic Principal Component Analysis	27
3.1.6	Automatic Landmark Identification	28
3.2	Machine Learning	30
3.2.1	Training and Learning	31
3.2.2	Types of Learning	32
3.2.3	Validation	34
3.3	Shape Matching	35
3.3.1	Feature Extraction and Detection	35
3.3.2	Constrained Local Models	37
3.3.3	Constructing the CLM	38
3.3.4	CLM Search	40
3.3.5	Initialising the CLM	41
4	Clinical Background	43
4.1	Extracapsular Hip Fracture Surgery	43
4.1.1	Intertrochanteric Fractures	45
4.1.2	The Tip–Apex Distance	46
4.1.3	Parker’s Ratio	48
4.2	Computer–Assisted Orthopaedic Surgery	50
4.3	Datasets	52
4.3.1	NHS Dataset	52
4.3.2	Online Digital Image Bank	53
4.4	“Interesting Parts” of the Femur	54
4.5	Project Rationale	55

5	Methodology	57
5.1	Overview	57
5.1.1	Automatic Guidewire Navigation Methods	58
5.2	Building the Shape Model	60
5.2.1	The Morphology of the Femur	60
5.2.2	Definition of Landmark Types	62
5.2.3	Morphometric Variables	65
5.2.4	The Shape Model	66
5.3	Building the Appearance Model	67
5.3.1	Correlation Filter-based Local Detector	67
5.4	Data	71
5.5	CLM Initialisation	71
5.5.1	Initialisation using Decision Trees	72
6	A Semi-Automatic Guidewire Insertion Planning Tool	75
6.1	Introduction	76
6.2	Method	77
6.2.1	Keystroke Level Modelling	80
6.3	Results	80
6.4	Analysis and Discussion	83
7	Implementation of the Fully Automatic Guidewire Navigation Tool	89
7.1	Introduction	90
7.1.1	Related Work	90
7.2	Clinical Background	91
7.3	Technical Background	92
7.4	Encoding the Shape of the Femur	93
7.5	The PCA-Shape Model	95
7.6	The PPCA-Shape Model	96
7.7	Constrained Local Models	97
7.7.1	A Point Distribution Model for the Femur	97
7.7.2	Local Texture Detectors	98

Contents

7.8 Correlation Filter Based Detection	99
7.8.1 MOSSE Filter-Based Local Detector	100
7.9 Experiments	102
7.9.1 Dataset	102
7.9.2 Applying the CLM Detector	102
7.9.3 CLM Search	103
7.9.4 CLM Initialisation	104
7.9.5 Automatic Landmark Detection	108
8 Conclusion and Further Work	115
Appendices	135
A.1 Tip-Apex Distance Illustrated	137
A.2 Typical implants used in Extracapsular Hip Fracture Surgery .	138
A.3 Centre coordinates and radii plots for the Semi-automatic Tool	139
A.4 Overview of Template-Based Object Matching Techniques . . .	140

List of Figures

2.1	Illustration of the ACM	10
2.2	Ideal contour location vs ASM or ACM failure	11
2.3	A proximate evolution of Deformable Template Models.	11
3.1	A Generic Framework of the CLM	39
4.1	Anatomy of the Hip Joint	44
4.2	Anatomy of the Femur	45
4.3	Two basic types of fractures.	46
4.4	Illustration of the Tip-Apex Distance	47
4.5	Illustration of TAD vs Cut-out Incidents	47
4.6	Parker's Ratio Illustrated	49
4.7	Implant postions based on Parker's Ratio calculation	50
4.8	Insertion of Guidewire shown	54
5.1	High-level overview of the Automatic Navigation Tool	59
5.2	Landmarks on the Femur	61
5.3	Illustration of the axis of ideal implant insertion	65
5.4	Workflow of MOSSE Filter-based Patch Model	68
5.5	Illustration of a decision tree	72
5.6	High-level overview of the CLM initialisation process	73
6.1	Steps to estimate the Femoral Head Centre	78
6.2	Figure showing how centre and radius of circle are calculated .	79
6.3	Components for image intensification during fracture surgery	80
6.4	Centre of Femoral Head Computed	82

List of Figures

6.5	Implant insertion vs Error Margin	83
6.6	Computed head centres plotted	85
6.7	Circle versus Ellipse Fitting of the Femoral Head	87
7.1	Plot of aligned shapes along with the mean shape	94
7.2	Successful instances of the localisation of femoral neck patches.	105
7.3	Successful neck localisation instances on the full femur.	106
7.4	Unsuccessful Localisation Attempts of the Femoral Neck Area	107
7.5	Landmarks detected by the CLM	108
7.6	CLM Localised Landmarks and their sample responses	109
7.7	Thresholded femur image with a poor contrast profile.	110
7.8	Computed values of Parker's Ratio based on landmarks detected	111
7.9	Detected landmrks and axes of insertion depicted	112
7.10	Bland-Altman Plot: PPCA vs Groundtruth	114
7.11	Bland-Altman Plot: PCA vs Groundtruth	114
A.1	Tip-Apex Distance Illustrated	137
A.2	Examples of implants	138
A.3	Implant viewed under x-ray	138
A.4	Plots of the computed femoral head centre for 6 images	139
A.5	Overview of Template-based Object Matching Methods	140

List of Tables

5.1	Landmark labels and descriptions	64
6.1	Femoral head centre coordinates and radii	81
6.2	Mean and Standard Deviation for head centre coordinates . . .	83
6.3	Comparison Between Circle and Ellipse-Fitting Approaches . .	87
7.1	Summary of Femoral Neck Localisation Results	107
7.2	Parker's Ratio P_r calculated for 10 images	111

Publications

Below is a list of publications resulting from this research:

- [1] **Okoli, A. B.**, Woo W.L., Dlay S.S. A Fully Automatic Guidewire Insertion Planning System for Extracapsular Hip Fracture Surgery. IEEE Transactions on Medical Imaging

- [2] **Okoli, A. B.**, Penny, J. B., Woo, W. L., Dlay, S. S. (2014, October). Guidewire insertion planning for extracapsular hip fracture surgery. In Biomedical Circuits and Systems Conference (BioCAS), 2014 IEEE (pp. 105-108). IEEE.

Acknowledgements

I would like to specially thank my supervisor, Professor Satnam Dlay for sharing his knowledge with me, his patience and understanding and for being ever present during the course of my research.

Special appreciation goes to my very good friends Nnamdi Maduka and Opeyemi Dele-Ajayi, who have been like brothers to me during the course of my research.

I appreciate my mother-in-law Mrs Yemisi Afenkhide, whose visit was timely in my last few months of research and for taking care of my girls while I spent long hours in the library.

I deeply thank my parents for all their support, prayers, love and encouragement.

And finally to my woman Bose, who has been in it with me all the way, I deeply appreciate her love, prayers and unflinching support in all areas. She has been like a rock and always given me reason to smile and be happy.

Dedication

This Thesis is dedicated to my parents Professors Bosa and Nkechi Okoli.

Abstract

During the implant insertion phase of extracapsular hip fracture surgery, a surgeon visually inspects digital radiographs to infer the best position for the implant. The inference is made by “eye-balling”. This clearly leaves room for trial and error which is not ideal for the patient.

This thesis presents an *image analysis* approach to estimating the ideal positioning for the implant using a variant of the *deformable templates model* known as the *Constrained Local Model (CLM)*. The Model is a synthesis of shape and local appearance models learned from a set of annotated landmarks and their corresponding local patches extracted from digital femur x-rays.

The CLM in this work highlights both Principal Component Analysis (PCA) and Probabilistic PCA as regularisation components; the PPCA variant being a novel adaptation of the CLM framework that accounts for landmark annotation error which the PCA version does not account for. Our CLM implementation is used to articulate 2 clinical metrics namely:

the *Tip-Apex Distance* and *Parker’s Ratio* (routinely used by clinicians to assess the positioning of the surgical implant during hip fracture surgery) within the image analysis framework. With our model, we were able to automatically localise significant landmarks on the femur, which were subsequently used to measure Parker’s Ratio directly from digital radiographs and determine an optimal placement for the surgical implant in 87% of the instances; thereby, achieving fully automatic measurement of Parker’s Ratio as opposed to manual measurements currently performed in the surgical theatre during hip fracture surgery.

Declaration

No portion of the work referred to in this thesis has been submitted in support of an application for another degree or qualification here at Newcastle University or any other university or other institute of learning.

Glossary

<i>jpeg</i>	Joint Photographic Experts Group
2D	Two-Dimensional
3D	Three-Dimensional
AAM	Active Appearance Model
ACM	Active Contour Model
AP	Aterior-Posterior: A directional reference in Anatomy indicating front-to-back views.
ASM	Active Shape Model
CalTAD	Calcar-referenced Tip-Apex Distance
CF	Correlation Filter
CLIC	Collection of Landmarks for Identification and Characterisation package: A software for extracting and managing landmarks.

Glossary

CLM	Constrained Local Model
DICOM	Digital Imaging and Communications in Medicine
DTM	Deformable Templates Model
FA	Factor Analysis
FFT	Fast Fourier Transform
GPA	General Procrustes Analysis
HOG	Histogram of Oriented Gradients
IT Fracture	Intertrochanteric Fracture: A fracture that occurs at the proximal part of the femur.
K-Wire	Kirschner Wire: sterilised metal pins used in binding bone fragments together in orthopaedic surgery.
KLM	Keystroke Level Modelling
LBP	Local Binary Patterns
LOA	Lower Limit of Agreement
MDL	Minimum Description Length
ML	Machine Learning

MOSSE	Minimum-Output-Sum-of-Squared-Error
MRI	Magnetic Resonance Imaging
OOI	Object of Interest
PACS	Picture Archiving and Communication Systems
PCA	Principal Component Analysis
PDM	Point Distribution Model
PPCA	Probabilistic Principal Component Analysis
RFRV	Random Forest Regression Voting
ROI	Region of Interest
SSA	Statistical Shape Analysis
SSM	Statistical Shape Model
SVM	Support Vector Machine
TAD	Tip Apex Distance
UOA	Upper Limit of Agreement
PR, P_r	Parker's Ratio

1 Introduction

In the orthopaedic management of hip fractures, a variety of techniques depending on the nature, site and severity of the fracture are utilised. The basis is usually to perform reduction, ensure proper healing and crucially regain or retain function of the affected area.

It is estimated that more than 700,000 mortalities due to hip fractures occur annually worldwide [1]; 24 per cent of which occur postoperatively within a year [2].

In UK, hip and thigh injuries have been identified as one of the leading causes of emergency admissions and readmissions [3]–[5].

The primary objective of the ideal surgical treatment of intertrochanteric fractures is to achieve stable fixation of the and early ambulation.

To this end, various intramedullary and extramedullary implants have been designed (See Figure A.2): *internal fixation with femoral nailing or sliding screws* [2], [6], [7].

These implants have advantages and disadvantages. Postoperative complications such as implant failure, iatrogenic injuries, etc. result in 20 percent of cases where these implant techniques are used [8]–[10].

The “debate” between which is a better option for treatment has raged for a long time in research circles. Some studies have noted that no closure has been reached concerning the ideal implant due to an incomplete understanding of the biological/biomechanical factors that govern the restoration of the affected area [11], [12]. However, one undesirable outcome common to both is a phenomenon known as “cutout” where the implant

1 Introduction

loses purchase to its attachment usually due to varus deviation or rotation of the femoral head which sometimes leads to fixation failure [13]–[15]. Several studies have considered factors that may be responsible for cutout. Most seem to lean towards the physical design of the implants [11], [16]–[18]; others site patient-related medical conditions like osteoporosis as a reason for this failure. Finite Element analysis has been used to assess the mechanical competence of the bone [19]–[22]. However, only few studies have considered placement of the implant as a predictor of fixation failure. A pioneering study, proposed the *Tip-Apex Distance* (TAD) (See Figure A.1) as a strong indicator of the probability of cutout [23]. In this study, the TAD is defined as the sum of the distance from the lag screw tip to the apex of the femoral head in millimetres(mm) on both the anterior-posterior (AP) and lateral radiographs adjusted for magnification [16], [23], [24]. In a later report, the results of a questionnaire completed by participating surgeons showed that although none of the surgeons directly quoted the use of TAD, they supported the theory of a central and deep placement [24]. It has been reported in several studies that a TAD value of greater than 25 mm is highly likely to result in cutouts [13], [16], [17], [25] and some studies have suggested that a value less than 20 mm should be the ideal [17], [26], [27]. So far, the TAD seems to be the “most accurate” predictor of cutout although opinions following previous research in the area are split concerning the ideal value of the TAD.

1.1 Motivation

While the TAD remains the most reliable way of predicting cutout, this metric lends itself to trial and error on the part of the surgeon as measurements are carried out manually. As some studies have rightly observed, expertise and experience are vital in the success of such procedures. Baumgaertner [24] observed that experienced practitioners favour a *central-deep* placement of the lag screw. While the technical competence of a surgeon should not be underestimated, their fallibility cannot be ruled out. The research challenge is: if it is possible to devise a means whereby surgeons do not have to rely on intuition during this process; such that a precise position can be predetermined before implant insertion commences. Consequently, it will be beneficial to devise a minimally invasive and more accurate means to obviate the intrinsic trial and error involved during the intraoperative period of the operation. A guidewire is usually employed prior to implant insertion and it is vital that the operator gets the insertion right in the minimum number of attempts ideally in ‘one-pass’ as a higher number of attempts is likely to be iatrogenic.

Thus the aim of this study is centred on the implementation of an *image analysis-based digital guidewire navigating tool* to enable surgeons plan the positioning to achieve an optimal location for the implant that satisfies both conditions of the TAD and central-deep placement.

1.2 Aim and Objectives

The aim of this project was the development of a digital-imaging surgical navigation tool that may be used to estimate the optimal position of implants on digital x-ray images during extra-capsular hip fracture surgery. The key objectives were:

- To determine a suitable framework for annotation of the femur and encoding of the shapes.
- To develop an automated system for identifying femur landmarks on digital x-ray images.
- Evaluation of the performance of the landmark identification system based on the clinical metrics utilised by surgeons for implant placement.
- Demonstrate how the system may be used in planning the insertion path of the surgical implant.

1.3 Contributions

This investigation carried out in this project culminated in the development of a new fully automatic digital system for guidewire insertion planning; based on the automatic detection of landmarks using PCA and PPCA-based shape models within the CLM framework.

The system has the potential of transforming existing surgical procedures in the management of intertrochanteric fracture surgery.

The following is a breakdown and brief description of the contributions.

1. **A novel landmark detection model based on using PPCA in the CLM framework** : The use of PPCA in shape analysis is not new but the use of PPCA within the CLM framework is a novel application which demonstrates advantages over the staple PCA variant.
2. **Identified Parker's Ratio and the Tip-Apex Distance metrics as actionable concepts**: Parker's ratio is metric that is used to determine an optimal placement for surgical implants in intertrochanteric fracture surgery. Existing estimates are made manually, however this project has demonstrated that this estimates can be computed automatically using landmark detection.
3. **A fully automatic means of determining the Parker's Ratio from radiographs**: These are measurements that are manually estimated under current surgical procedures. This research has been able to demonstrate that these measurements can be acquired through image analysis. The use of computer vision as a means of measuring Parker's Ratio is a new concept presented in this thesis.

1 Introduction

4. **A semi-automatic system for determining the centre of the femoral head:** The centre of the femoral head is a vital consideration in implant positioning. The semi-automatic approach establishes a means of locating the femoral head centre with minimal user input.

5. **A fully automatic digital guidewire insertion planner:** This is the big picture of the research carried out in the project and has high clinical relevance.

1.4 Thesis Outline

This thesis is made up of 6 chapters. The following are a summary of the work carried out in the thesis:

Chapter 2 presents a critical review of work related to this thesis.

It specifically treats the parametric variants of the deformable model, detailing the evolution from the primitive “Snakes” to the cutting-edge Constrained Local Model (CLM) used as the basis of this work. It also presents a clear distinction between 2 classes of CLMs not immediately apparent in literature and their predominant application areas.

Chapter 3 establishes the technical underpinning for the research carried out such as methods and approaches used in image analysis. It elaborates on techniques such as Principal Component Analysis (PCA), Probabilistic Principal Component Analysis (PPCA), Machine Learning applications for analysing the shape and texture of an object and how these techniques can be combined in the CLM framework to achieve automatic landmark identification.

Chapter 4 defines and introduces the key medical concepts that underlie the research project. It identifies gaps, challenges and areas of hip fracture surgery management that can be potentially improved by computer-assisted intervention.

Important concepts such as the Tip-Apex Distance and Parker’s ratio are introduced and reconciled with applications of Image analysis.

The dataset used for analysis along with the anatomy of the femur - the object of interest are also presented in this part.

1 Introduction

Chapter 5 describes the methodology used to implement the automatic guidewire navigation tool. It presents an overview of the components of the Constrained Local Model, namely: the shape model and appearance model. The construction of each model, its constituent parts and the entire process from image capture to localisation of landmark features is documented.

Chapter 6 contains preliminary investigations performed in the research which culminated in a conference paper.

A semi-automatic guidewire insertion planner that requires minimal input from an operator is introduced. The importance and potential benefits of the system are highlighted.

Chapter 7 details the development of a fully automated guidewire insertion planning system; how image analysis can be used as a tool to identify an optimal location for the implant. The various algorithms used to achieve automatic landmark identification are implemented in the CLM framework. The CLM is constructed and tested on a partitioned dataset of training and test data with the overall aim of automatically detecting landmarks of the femur. The results and analysis are also presented in this chapter.

Chapter 8 gives a conclusion and suggestion for improvement and future work. It also outlines the weaknesses of the designed system and recommends alternative courses of action.

2 Literature Review

This chapter presents a critical review of related studies that underpin the work carried out in this thesis.

It gives an overview of the *Deformable Templates Model* (DTM), which is an object localisation framework used to learn the attributes of the *object of interest* (OOI). The OOI in this instance, is the *proximal femur*. Particular instances of the model and a review of the *Constrained Local Model* (a variant of the DTM) which is at the crux of the Image Analysis framework utilised in this project is presented.

2.1 Deformable Models

“A deformable template model is a model, which deforms a shape to match a target object in a given image within implicit or explicit optimisation constraints.” – Fisker [28]

In recent times, utilising model-based approaches in image analysis has proven very popular within the computer vision community.

The *Deformable template model* is an abstraction of the *template-based object matching* technique (see [29] for a more detailed exposition) and has been successfully applied to image understanding problems across various domains such as facial recognition [30], [31], medical imaging [32]–[34],

2 Literature Review

general object recognition [35]–[37], etc. The attractiveness of these models lies in their ability to capture the properties of objects with large variability; and have naturally been the subject of intense research, evolving over the years into different forms and applications (see Figure 2.3) .

One of the earliest known deformable template models is the Active Contour Model (ACM) – also known as *Snakes* formulated by Kass [35].

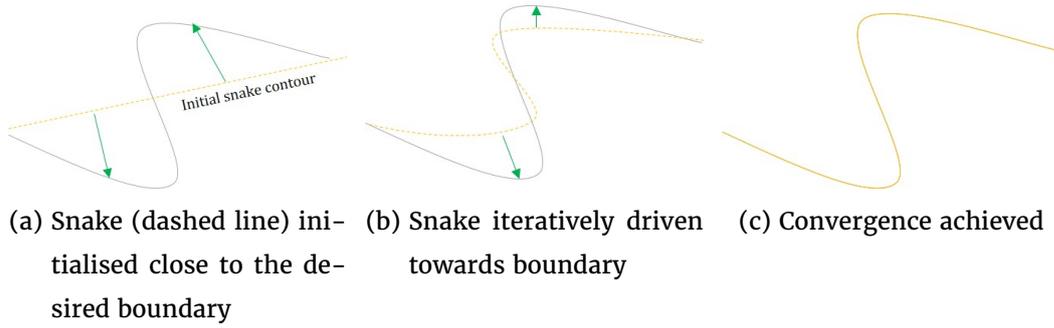


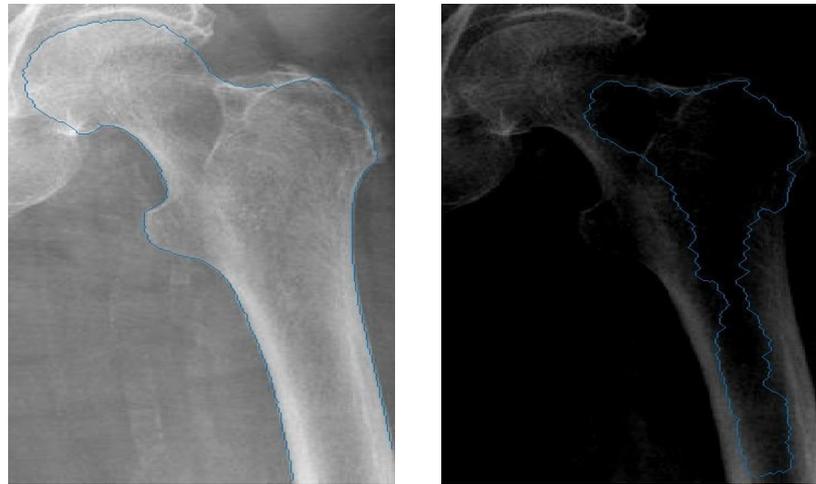
Figure 2.1: Illustration of the operation of the ACM: The blue arrows signify the migration of the Snake towards the desired boundary.

The ACM model models an object contour as a set of edge landmarks to which an energy-based function is applied to govern local shape variation. The Snake is initialised close to the desired boundary as shown in Figure 2.1 and driven towards the contour edges using an energy minimisation principle. The Snake assumed to be elastic, is defined as a set of points \mathbf{z}_i where $i = 1, 2, \dots, n$, along a contour $\mathbf{z}(s)$, with internal and external energy terms $E_{internal}$ and $E_{external}$ respectively. $E_{internal}$ controls the way the snake deforms and $E_{external}$ governs the fitting of the Snake to the desired object. The energy relation governing the Snake's movement may be written as:

$$E_{snake}^* = \int_0^1 E_{snake}(\mathbf{z}(s)) \quad (2.1.1)$$

$$\int_0^1 E_{snake}(\mathbf{z}(s)) = \int_0^1 (E_{internal}(\mathbf{z}(s)) + E_{image}(\mathbf{z}(s)) + E_{con}(\mathbf{z}(s))) ds \quad (2.1.2)$$

The weakness of the ACM is the lack of its specificity to the desired object, often being drawn to edges that are not part of the object to be identified; as shown in Figure 2.2, where the ACM fails to localise the desired contour.



(a) Ideal boundary to be located (blue contour)

(b) Failure to locate desired boundary typical of both ACMs and ASMs in challenging conditions.

Figure 2.2: Failure of the ASM to locate a given boundary. In (b) it can be seen that the model is attracted towards undesired edges.

To circumvent this challenge, the use of priori information has been explored to great effect [38].

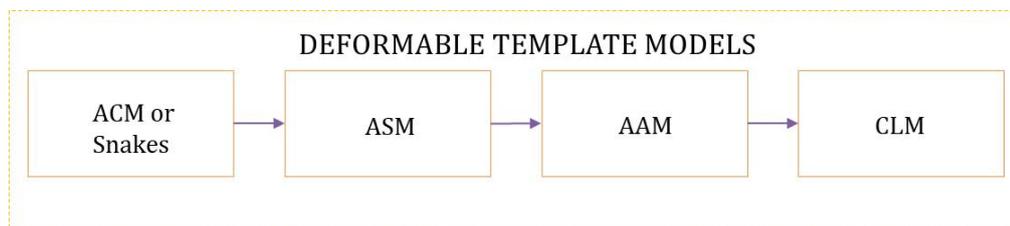


Figure 2.3: A proximate evolution of Deformable Template Models.

2 Literature Review

Cootes [39] formulated a ‘smarter’ version of snakes known as the *Active Shape Model* (ASM) which improved on specificity.

ASMs capture the variability of the desired object from a *training set* – a set of similar instances of the object. The ASM learns the variability across the training set by applying

Principal Component Analysis (see Section 3.1.4) to the normalised training set. Normalisation is carried out after the shapes have been aligned using General Procrustes Analysis (Section 3.1.3).

The initial ASM algorithm made use of edge detection to determine the optimal landmark position relying on the profile normal to the boundary of the model contour. Hence the landmarks were required to be located on the strongest edges. In practice, this is not always feasible as the strongest edges might occur at positions unrelated to the contour of the desired shape. This effect is demonstrated in Figure 2.2(b) which is typical of ASMs when the model encounters spurious edges and/or is not initialised close enough to the desired contour. An iterative improvement incorporates statistics of the grey-level profile normal to the model boundary. While this modification does not require the strongest edges, ASMs struggle with initialisation of the feature positions i.e. the model has to be initialised as close as possible to the shape of interest. This makes ASMs unsuitable for medical imaging problems as images in this domain vary widely across a training set.

The often homogeneous texture around regions of interest also affects the distinctiveness of the desired features, rendering ASMs inefficient due to their tendency to be attracted to spurious edges or get stuck in local minima.

A further extension to ASMs known as the *Active Appearance Model* (AAM) was proposed by [40]–[42] which utilises both the shape and textural information extracted from triangulated patches across the entire object. The AAM utilises the entire object shape rather than the landmarks to drive

convergence. The model seeks to minimise the residual between the model texture and the image of interest. Due to its heavy reliance on textures, the AAM suffers from the variation between training set samples and the test image thus, fluctuations in illumination adversely affect the efficiency of AAMs; which is a typical occurrence in medical radiographs for example, the appearance of soft tissue and organs surrounding bones in x-ray images tend to give off higher illumination than the objects of interest.

A more robust model known as *Constrained Local Models* (CLM) that combines both textural and shape information was proposed by Cristinacce et al. [43], [44]. The following section gives a review of CLMs and their application.

2.2 Constrained Local Models

CLMs are similar to AAMs in the sense that they both utilise a template model to generate feature templates to be matched to a search image. However, while AAMs use triangulated patches over the entire image to encode textural information, CLMs capture only textures surrounding each landmark or feature point and uses a different convergence procedure. Thus the CLM is a joint model of shape (*finite configuration of landmarks*) and texture (*appearance around each landmark*) used to search for and locate the desired features in a test image. Local patches around each feature are readily learned from annotated training images for each feature point. The primitive CLM uses a statistical model to capture both shape and appearance variations across a training set composed of similar instantiations of the desired object. Developments in CLM methodology have utilised classifiers [45], regressors [46], [47], filters [31] and neural architecture [48], [49] to encode the local features of interest, achieving more robustness and accuracy in the process.

2 Literature Review

The most popular applications of the CLM in literature are in the facial tracking and alignment domain [31], [43]–[45], [47], [50]–[54]. Recently, CLMs have started becoming more popular in the medical imaging domain [46], [55]–[57].

In literature, CLM methodology has been shaped by two schools of thought. While both agree that local response maps can be generated using arbitrary detectors, both differ in the choice of shape constraints. One approach favours point distribution model (PDM) as the regularisation module of the CLM while the other considers the CLM as a generalised framework that combines an arbitrary constraint and local detectors – it therefore considers the former as an instantiation of a more general framework.

We shall refer to both approaches as *PDM-based CLMs* and *The Generalised CLM* respectively and now give an overview of these two areas in Sections 2.2.1 and 2.2.2.

2.2.1 PDM-based CLM

Due to the inherent pliability of CLMs, a number of variants to the original model have been proposed. The flexibility of the model lies in its modularity such that a number of regularisation techniques can be used to constrain the desired shape while a variety of local feature detectors can be chosen to learn local textures as noted in Section 2.2.

Research by [30], [43], [44], [46], [55]–[57] argues that the shape constraint must be a statistical shape model which is built using PCA, i.e. PCA is used to decorrelate the shapes in the training set after the annotated landmarks have been aligned using General Procrustes Analysis. This yields a compact model that is able to deform based on the modes learned from training. The resulting shape model is then used to regularise the local responses. While PDM-based CLMs have retained the PDM as the

core of the regularisation module, proponents of the approach have exploited the local detector component of the CLM. Wang [45] demonstrated this by using the Support Vector Machine (SVM) to learn local feature maps around each landmark subject to the constraints of a PDM. This approach, is also evident in excellent work carried out by Lindner [46], [55] in the fully automatic segmentation of the proximal femur on digital radiographs used to monitor osteoarthritis; where the local feature response maps were learned using *Random Forest Regression Voting* and regularised using the conventional PDM. Lindner et al [56] further demonstrated the versatility of the approach in the detection of facial features and the annotation of hand joints, achieving state-of-the-art performances in both instances. This approach was also successfully applied in the automatic localisation of vertebrae on x-ray images used in bone densitometry [57].

2.2.2 The Generalised CLM

Work carried out by [31], [53], [54] refers to the CLM a framework that utilises any form of constraint to fit the response surface of the local detectors; and this regards the PDM-based CLM as an instance of the Generalised CLM.

This realisation of CLMs was demonstrated by Saragih [53] who used *regularised landmark mean-shifts* rather than the canonical statistical shape model to fit the response surfaces generated by a classifier, citing the inability of the PCA shape model to perfectly estimate true landmark locations due to the truncation of eigenvalues.

Asthana [47] further elaborated on this technique by proposing a novel regression-based fitting method - *Discriminative Response Map Fitting* that outperformed Saragih's regularised landmark mean-shifts technique in facial feature fitting. Another variant of the Generalised CLM is found in Martin's works [31], [54] where the landmark alignment goal is posed as a Bayesian problem such that the observed shape is assumed to be the

2 Literature Review

likelihood term and the training set of shapes are encoded as the prior in Bayesian paradigm – the local response maps were extracted using *Minimum-Output-Sum-of-Squared-Error* (MOSSE) filters.

While both approaches have been used with a great degree of success, it is evident that the Generalised CLMs are more suited to facial feature fitting problems and the PDM-based CLMs, while also successful in face alignment, are predominantly suited to medical image analysis problems.

In this work we depart slightly from the canonical CLM by employing the use of the *Probabilistic PCA* (PPCA) technique as the regularisation model. This is to account for any errors in measurement during the landmark annotation process. The PPCA technique is a latent variable model formulated by Tipping and Bishop [58].

The idea of using PPCA to model shape variation is not a completely new idea as [59] employed its use as a hierarchical Bayesian prior to estimate 3D shape and motion. Mutsvangwa [34] also utilised a PPCA-based SSM to model the shape of the human scapula and humerus. However neither of these approaches utilised PPCA within the CLM framework.

In this thesis, a PPCA-based PDM is constructed from manually annotated training samples and used to fit local response maps extracted from surrounding textures at each feature point (see Sections 3.1.5 and 7.6 for further details).

2.2.3 Texture Models in Literature

The initial implementation of the CLM made use of *normalised correlation* templates, constructed by averaging patch samples and normalising them. However, such is the flexibility of the CLM, that a variety of methods could be used to capture the local texture appearance for each feature point in an image. For example Machine Learning methods have recently been used for this purpose. [45] utilised Support Vector Machines to learn the local texture features of the face while [46] used Random Forest Regression Voting (RFRV) to learn the local textures in the segmentation of the proximal femur. [57] also utilised RFRV in the localisation of vertebrae on DXA images. A recent formulation [31] makes use of correlation filters to learn an optimal filter representing the local response of local patches. A key advantage of using correlation filters is the minimal training required, particularly due to the fact that negative samples don't have to be explicitly assigned unlike in the cases of SVMs or RFRV. Correlation filters are popular in the object tracking domain [60], [61] but can be adapted for still images as is the case in our work. Further details on Correlation filters can be found in Section 5.3.1.

Chapter 3 presents an in-depth background/theory of PDM, PCA, PPCA and cognate techniques underpinning this work.

3 Technical Background

This chapter is a synopsis of the methods and techniques that underpin the development of the automated landmark identification system.

The following sections present the underlying concepts of statistical shape analysis used in this project (Section 3.1); an overview of Machine Learning as it applies to this work is presented in Section 3.2 and an introduction to the methods for automatic shape matching is presented in 3.3.

3.1 Statistical Analysis of Shape

Statistical Shape Analysis (SSA) encompasses the mathematical methods used to measure and represent the geometric properties of an OOI such that the variability and important attributes of the shape of the object within a population is captured and accounted for by a statistical model. SSA has broad applications in diverse fields such as biology, architecture, agriculture, etc. and can be broadly divided into two techniques. Namely: *Landmark-based* and *Deformation-based* techniques. In the Landmark-based approach, a given shape is characterised as a finite set of key points along the contours of the shape. A statistical model is used derive a compact representation of a set of similar instances of the given shape. An alternative to this approach is to regard the contours of the desired shape as a configuration of curves and/or surfaces. The Deformation-based approach considers the relationships between shapes as a mapping of one shape to

3 Technical Background

another based on a force applied to the input shape. Thus the difference between two given shapes is regarded as the smallest deformation between the two. For a more detailed exposition on statistical shape analysis, see Dryden's excellent work [62].

3.1.1 What is shape?

Generally, shape may be thought of as the physical form or appearance of an object. In image analysis, the shape of an OOI is identified by its *outline or boundary* and *texture* (analogous to form and appearance respectively). The geometric properties of the shape are captured in the outline while the colour may be represented by its texture or intensity.

In medical imaging, shape analysis may be contour-based or appearance-based or a combination of both. In medical images particularly radiographic or x-ray images, the homogeneity of the OOI varies across different instances due to the presence of artefacts as a result of surrounding organs or tissue, anatomy, disease, gender etc. A medical approach to shape analysis otherwise known as morphometrics usually utilises geometric/size measurements which do not account for the global shape of the object. In this work the local textures around the shape outline are also taken into account (see Section 3.3.1) in constructing a shape model. For the purpose of this project, the following definition [63] for shape is adopted:

“Shape is the geometric information retrieved after the location, scale and rotational effects have been filtered out of an object.”

With this in mind, to build a finite-dimensional representation of a shape, a “*connect-the-dots*” approach is adopted i.e. landmarks or key-points are placed along the contour of the object in a consistent manner.

Thus the shape is composed of a set of these landmarks. Landmarks should be consistent across a set of similar shapes where they can be located repeatedly and reliably.

Landmark-based representation facilitates the statistical analysis of shapes using the Point Distribution Model(PDM).

3.1.2 Point Distribution Models

The PDM also known as a Statistical Shape Model was first introduced circa 1992 [36].

It has since then been utilised extensively in numerous computer vision applications related to shape analysis [30], [33], [42]–[44], [46], [56], [64]–[68].

The PDM models a dataset of shapes according to their mean and a restricted number of eigenvectors. Each shape is represented by landmark points or coordinates located on their outlines. The PDM offers a compact statistical description of the dataset based on a set of training samples. The way it works is a dataset of shape vectors which represent the class of objects to be extracted, is decomposed and reduced to learn the variability patterns of the shapes in the dataset. The PDM captures the statistics of the displacements of the landmarks across the training set. PDMs capture the shape variation and eliminates redundancies across the dataset using eigenvectors computed from PCA.

The model may be viewed as a ‘flexible’ model that deforms according to constraints learned from the training set such that the model is able to generate shape instances to match a novel shape(i.e. a similar shape to those in the training set but not in the training set) with a high probability.

3 Technical Background

Mathematically, a shape \mathbf{x} may be computed as

$$\mathbf{x} = \bar{\mathbf{x}} + \psi \mathbf{b} \quad (3.1.1)$$

where $\mathbf{x} = (x_1, y_1, x_2, y_2, \dots, x_n, y_n)^T$ a $2n \times 1$ vector, represents a 2D shape instance of the training shapes consisting of n landmarks, $\bar{\mathbf{x}}$ is the mean of the set of shapes while ψ defined by equation 3.1.2 is the PCA projection matrix of t eigenvectors which corresponds to the t largest eigenvalues

$$\psi = \begin{bmatrix} \psi_1 \\ \psi_2 \\ \vdots \\ \psi_t \end{bmatrix} \quad (3.1.2)$$

$\mathbf{b} = (b_1, b_2, \dots, b_t)^T$ is the t -dimensional model coefficient vector given by:

$$\mathbf{b} = \psi^T (\mathbf{x} - \bar{\mathbf{x}}) \quad (3.1.3)$$

The first step in building PDM is obtaining a dataset of sample shapes whose landmarks have been annotated on the corresponding images. The dataset should be chosen such that the variations of interest are represented across the images.

Assuming there is a shape of the OOI to be modelled, usually, these objects are found on digital images. For example in this project, the dataset consists of digital x-ray images of the femur taken in the anterior-posterior(AP) view only.

The images are collated into a set on which the shapes contours are annotated as landmarks. The landmarks are usually 2 or 3 dimensional; although this project deals with the 2D variant. These points are used to establish a correspondence across the training set.

Due to rotation, scale and translation effects, the OOI tends to vary across the dataset thus it becomes paramount that the contours undergo some

form of standardisation which in this case is an alignment within a common coordinate frame. This ensures that all the shapes are compared on a standard baseline.

3.1.3 Shape Alignment

To extract only shape information in accordance with the definition of shape in [63], the members of the training set have to be aligned using General Procrustes Analysis (GPA). The aim is to eliminate translation, rotation and scaling effects leaving pure shape information by computing a transformation \mathcal{G}_ρ with parameters ρ that superimpose each training vector unto a common reference frame so that

$$\mathbf{x}^{ref} = \mathcal{G}_\rho(\mathbf{x}) \quad (3.1.4)$$

Prior to alignment, all shape vectors are whitened or centred i.e. the centre of gravity of each shape is aligned. GPA is then applied to calculate the transformation \mathcal{G}_ρ that best fits shapes \mathbf{x}_i to a reference mean shape $\bar{\mathbf{x}}$ - initialised with a random shape (usually chosen to be the first shape in the dataset) in the training set. The best fit is based on an iterative computation for the minimum of the sum of square distances between a moving shape and the reference mean shape i.e. based on equation 3.1.5

$$\arg \min_{\rho} \sum_i |\mathbf{x}_i - \mathcal{G}_\rho(\bar{\mathbf{x}})|^2 \quad (3.1.5)$$

until convergence is achieved. Principal Component Analysis is applied to the aligned dataset to build the PDM.

3.1.4 Principal Component Analysis (PCA)

PCA is a dimensionality reduction technique used to not only reduce the dimensions of a dataset but also capture the major variations across the dataset. That is what makes it a very useful and popular tool for modelling PDMs. In this case, PCA is used to encode the variation across the n shape vectors $\mathbf{x}^{(i)}$ in the dataset $\mathcal{D} = \{\mathbf{x}^{(i)}\}_{i=1}^n$ aligned using the technique specified in 3.1.3. The PCA algorithm to compute a k - dimensional approximation of the dataset \mathcal{D} with $\dim(\mathbf{x}^{(i)}) = d$ may be stated as follows:

1. Compute $\bar{\mathbf{x}}$ the $d \times 1$ sample mean vector and the $d \times d$ covariance matrix

$$\mathbf{m} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}^{(i)}, \quad \mathbf{C} = \frac{1}{n-1} \sum_{i=1}^n (\mathbf{x}^{(i)} - \bar{\mathbf{x}})(\mathbf{x}^{(i)} - \bar{\mathbf{x}})^T. \quad (3.1.6)$$

2. Find the eigenvectors $\mathbf{e}^{(1)}, \dots, \mathbf{e}^{(d)}$ of the covariance matrix \mathbf{C} , sorted such that $\mathbf{e}^{(i)}$ is greater than $\mathbf{e}^{(j)}$ for $i < j$.
Compose the matrix $\mathbf{E} = [\mathbf{e}^{(1)}, \dots, \mathbf{e}^{(k)}]$.

3. The lower dimensional representation of each data point \mathbf{x}^i is given as:

$$\mathbf{p}^i = \mathbf{E}^T (\mathbf{x}^{(i)} - \bar{\mathbf{x}}). \quad (3.1.7)$$

4. The approximate reconstruction of the original data-point

$$\mathbf{x}^{(i)} : \mathbf{x}^{(i)} \approx \bar{\mathbf{x}} + \mathbf{E}\mathbf{p}^{(i)} \quad (3.1.8)$$

5. The total squared error over the dataset due to the approximation ϵ is:

$$\epsilon = \sum_{n=1}^N (\mathbf{x}^{(i)} - \tilde{\mathbf{x}}^{(i)})^2 = (n-1) \sum_{j=k+1}^d \lambda_j \quad (3.1.9)$$

3.1 Statistical Analysis of Shape

where $\tilde{\mathbf{x}}^{(i)}$ is a reconstructed shape and $\lambda_{k+1} \dots \lambda_n$ are the truncated eigenvalues. ϵ in 3.1.10 indicates how close the PDM is to approximating the dataset. A small ϵ indicates a small deviation from the dataset and vice-versa.

A strong advantage of PCA is its innate ability to simultaneously reduce the dimensions of the data while retaining a strong variance across the training set. As seen in 3.1.8, it is adequate to store the k largest eigenvalues and corresponding eigenvectors such that $\sum_{i=1}^k \lambda_i$ stores a percentage of $\sum_{i=1}^{2n} \lambda_i$ e.g 90% – 95%.

After dimensionality reduction, an accurate representation of the shape vector \mathbf{x} is written as:

$$\mathbf{x} = \bar{\mathbf{x}} + \tilde{\mathbf{E}}\tilde{\mathbf{p}} + \epsilon \quad (3.1.10)$$

The basis matrix \mathbf{E} and vector \mathbf{p} are the parameters controlling the PDM such that $\mathbf{p}^{(i)}$ specifies the deviation from the mean and determines the plausibility of shapes across the distribution. Each eigenvalue λ_i specifies the variance of its corresponding eigenvector $\mathbf{e}^{(i)}$. The total sum of the eigenvalues gives the total variance across the dataset. The significant shape variations are represented by the larger eigenvalues and the smaller eigenvalues signal small and local variations.

A constraint is usually applied to the elements of $\mathbf{p}^{(i)}$ to ensure that only shapes with high similarity to the the shapes in the dataset are generated. The limits are applied with respect to the standard deviation from the mean i.e. $-3\sqrt{\lambda_i} \leq \mathbf{p}^{(i)} \leq 3\sqrt{\lambda_i}$ for $\sqrt{\lambda_i}, 1 \leq i \leq k$.

Assuming a Gaussian distribution across the elements of $\mathbf{p}^{(i)}$, a Mahalanobis distance constraint may be placed on the distribution. Given an aligned shape \mathbf{x}_D the square of the Mahalanobis distance d_M from the mean shape $\bar{\mathbf{x}}$ is given by:

$$d_M^2 = (\mathbf{x}_D - \bar{\mathbf{x}})\Sigma^{-1}(\mathbf{x}_D - \bar{\mathbf{x}}) \quad (3.1.11)$$

3 Technical Background

Interpreting the PCA result in terms of Mahalanobis distance in equation 3.1.11 gives:

$$d_M^2 = \sum_i^{2n} \frac{p_{D_i}^2}{\lambda_i} \quad (3.1.12)$$

The intuition behind equation (3.1.12)¹ is that the minimised squared distance of d_M implies a strong equivalence between \mathbf{x}_D and the shape variance distribution across the dataset \mathcal{D} .

After PCA is applied to the aligned shapes, the landmarks are projected onto a new coordinate system. A dataset usually contains correlated and redundant information since it consists of a set of different instances of the same object, PCA is used to eliminate this redundancy by finding the orthogonal axes along which the variance across the data is maximised. These axes are referred to as principal components \mathbf{e}_1 and \mathbf{e}_2 such that each vector $\mathbf{x} \in \mathcal{R}^2$ may be represented as

$$\mathbf{x} = \bar{\mathbf{x}} + p_1 \mathbf{e}_1 + p_2 \mathbf{e}_2 \quad (3.1.13)$$

To achieve dimensionality reduction, each vector may be written as

$$\mathbf{x} \approx \bar{\mathbf{x}} + \mathbf{p}_1 \mathbf{e}_1 \quad (3.1.14)$$

where each vector is approximated by its nearest neighbour on the first principal axis.

The dataset \mathcal{D} could then be represented a single parameter statistical model.

A transform \mathcal{G}_ρ is then used to transpose the data from the new axes to the image domain.

$$\mathbf{x} = \mathcal{G}_\rho(\bar{\mathbf{x}} + \tilde{\mathbf{E}}\tilde{\mathbf{p}} + \epsilon) \quad (3.1.15)$$

¹This equation is also equivalent to the χ^2 distance between two variables.

Note that $\tilde{\mathbf{E}}$ is the matrix containing the k eigenvectors corresponding to the k largest eigenvalues while $\tilde{\mathbf{p}}$ is the k dimensional vector of k shape modes.

3.1.5 Probabilistic Principal Component Analysis

Given the dataset $\mathcal{D} = \{\mathbf{x}^{(i)}\}_{i=1}^n$, PCA may be used to form lower dimensional representations of \mathcal{D} on the assumption that the elements of \mathcal{D} lie close to a linear subspace. However, no matter how close the approximation is to the original data, there is always some “discrepancy” due to the data not perfectly lying on the linear subspace. This is taken into account and modelled using the probabilistic formulation of PCA known as Probabilistic Principal Component Analysis (PPCA).

PPCA is a latent variable realisation of the classic PCA in section 3.1.4 which is deduced from the standard Factor Analysis (FA)² where the data \mathbf{x} is modelled as a linear aggregate of a latent or hidden variable \mathbf{z} i.e.

$$\mathbf{x} = \boldsymbol{\mu} + \mathbf{B}\mathbf{z} + \boldsymbol{\xi} \quad (3.1.16)$$

and \mathbf{z} is assumed to be Gaussian with zero mean unit isotropic variance, i.e. $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\boldsymbol{\mu}$ is the maximum likelihood estimate of the data mean and also sets the origin of the coordinate system, $\boldsymbol{\xi}$ is Gaussian distributed noise with zero mean and covariance $\boldsymbol{\Sigma}$ i.e. $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$. \mathbf{B} is an $n \times d$ factor loading matrix that is equivalent to the basis matrix \mathbf{E} in PCA.

However, where PPCA and FA differ is in the assignment of the covariance $\boldsymbol{\Sigma}$, where $\boldsymbol{\Sigma} = \sigma^2 \mathbf{I}$ for PPCA and $\boldsymbol{\Sigma} = \text{diag}(\lambda_1, \dots, \lambda_n)$ for FA.

Therefore, in PPCA,

$$\mathbf{B} = \mathbf{U}_k \text{diag}\{\lambda_1 - \sigma^2, \dots, \lambda_k - \sigma^2\} \mathbf{R} \quad (3.1.17)$$

²for in-depth treatment of FA, the reader is referred to [58], [69], [70]

3 Technical Background

Where \mathbf{U}_k is the $n \times k$ matrix whose column vectors span the k largest eigenvalues λ_i $i = 1, 2, \dots, k$ of the covariance matrix Σ of \mathbf{x} while \mathbf{R} is an arbitrary orthogonal matrix.

$$\sigma^2 = \frac{1}{n-k} \sum_{i=k+1}^n \lambda_i \quad (3.1.18)$$

$$\therefore \mathbf{B} = \mathbf{U}_k \left(\Sigma_k - \frac{\mathbf{I}_k}{n-k} \sum_{i=k+1}^n \lambda_i \right) \mathbf{R} \quad (3.1.19)$$

While PCA is effective in achieving dimensionality reduction it does not take into account noise in the dataset which could result from errors in the annotation of landmarks. PPCA accounts for this uncertainty by incorporating a noise parameter in the model.

3.1.6 Automatic Landmark Identification

Acquisition of high quality ground-truth is vital for the application of image identification systems. In several research domains, general-purpose datasets are utilised, however in real world applications such as medical imaging, such datasets are not readily available and have to often rely on manual annotation of images which is a subjective process, expensive, laborious and error-prone. Thus devising an alternate, objective and automatic means of annotation is a hot research area.

An automated annotation system for images would require that feature points along the contour of the image are located with a high degree of accuracy, close to or better than manual annotation, without human intervention. Each landmark therefore, would be a feature point of interest and a detection system would be required to predict the position of each landmark. While the feature detector operates locally, a shape model would be required to globally constrain the possible feature positions within the image.

Automatic image annotation is not a new technique and has been applied in many areas of image processing and computer vision such as satellite imaging, pedestrian detection, medical imaging, agriculture, image retrieval systems etc.

The application of automatic annotation in the medical domain is an important facet of patient management. Thus there have been applications tailored to specific challenges such as a semi automatic hybrid system developed to annotate anatomical structures in brain MRI images [71], a system that automatically labels landmarks in mammograms [72], the automatic detection of retinal lesions [73], “Snakules” is an active contour-based algorithm for annotating spicules on mammography [32], Lindner *et al* [46] devised an excellent automatic segmentation system for the femur on x-ray images. The basis for the automation was a random forest voting algorithm which is a machine learning (ML) technique used in making predictions usually without human interference. One of the main objectives of this work is to automate the landmark identification process and to achieve this, ML techniques will be a key component of the system.

3.2 Machine Learning

*A computer(machine) is said to **learn** from experience E with respect to a task T and performance measure P , if its performance at task T , measured by P improves with experience E – Mitchell, 1997.*

Machine Learning³(ML) is a branch of AI that enables computers learn from data and improve from experience without explicit programming by a human. Its application varies across several fields including medicine, finance, engineering, etc. The following subsections present key concepts of ML, and the specific techniques used in this work. Mitchell’s work along with several others contain comprehensive and in-depth expositions on the subject [69], [76]–[82]. If $\mathbf{x}^{(i)}$ is the input vector representing the i th n -dimensional training sample and $\mathbf{y}^{(i)}$ the desired target of the i th sample. Given a training set in the form $(x_1, y_1), \dots, (x_n, y_n)$ the aim of training the model is to learn a function g that maps the input X to the output space Y , i.e. $g : X \rightarrow Y$ such that g can correctly infer the desired output when presented with valid test data. The overall goal of training a model is to achieve good *generalisation*.

To accomplish this, the following usually precede model training :

- *Data acquisition and preprocessing*: This refers to identifying and obtaining suitable training samples and extracting training data that is typical of the problem to solved.

³The term “Machine Learning” was first coined by Arthur Samuel in 1959 to describe the transfer of human intelligence to machines [74]. A machine in this sense is an abstract term that describes any autonomous system that may be implemented as software and should not necessarily be perceived as a physical entity [75].

- *Feature extraction and selection:* The input vector usually represents those set of features, which are the most likely to influence the accuracy of g .
- *Identifying the suitable learning algorithm(s) and selecting the corresponding parameters.*

The algorithms are applied to the training dataset and the performance of the model is evaluated using test data.

3.2.1 Training and Learning

The learning paradigm in machine learning is often accompanied by training, as both depend on the other. A system is trained to learn, and the level of learning determines when training should stop. The way an ML system learns is analogous to the way humans learn. Humans are able to learn how to complete a task after being presented with few examples. The same principles apply to ML where a computer is presented with examples with the aim of completing a task. Its performance is evaluated based on how well it learns from the examples and performs the given task. The following paraphrased definition by [70] captures the concept of training and learning in ML:

“A set of methods that may be used to train a computer to automatically learn patterns in data and use the learned patterns to predict novel data or make decisions under uncertainty.”

3 Technical Background

In ML, a model is (usually) a concise representation of the information learned from the underlying relationship between variables in a training dataset.

Thus learning may be referred to as model learning or training - where a model is taught to deduce some function from training data.

Usually, training data is a set of observations, where each is an input training vector, x_i and target y_i .

Just like humans have different modes of learning, ML systems have different learning techniques.

3.2.2 Types of Learning

Learning in ML consists of 3 broad classes namely: Supervised Learning, Unsupervised Learning and Reinforcement Learning.

- *Supervised Learning*: This refers to a mode of learning where the system is trained with labelled data from which a model is learned and subsequently used to make a decision on novel or unlabelled data. *Classification* which is the grouping of data into homogeneous classes and *Regression* which is the prediction of numeric data are 2 examples of Supervised Learning.
- *Unsupervised Learning*: refers to training a system with unlabelled data such that the system is able to learn a model that best describes the pattern within the data. Examples of this kind of learning are *Clustering*: where the algorithm discovers intrinsic groups or clusters within the unlabelled data and *Dimensionality Reduction*: where the aim is to achieve compression of the data i.e. the data is represented with fewer components than the original form.

- *Reinforcement Learning*: This kind of learning is almost like a hybrid between supervised and unsupervised learning. The system is not presented with labelled examples but is rewarded for making the right decision and penalized in some applications for making the wrong decision. The system learns based on the rewards received. The system learns via “*trial-and-error interaction*” with data. It seeks actions that are the most likely to maximise rewards within the learning space.

No particular method is deemed superior to the other, rather the type of data and nature of the task to be accomplished tend to influence the method applied. Sometimes two methods are combined to achieve a complement such that each gives leverage to the other. For example unsupervised learning may be used to discover a hidden pattern in a dataset with no known structure and then the discovered pattern is used as a template for successive datasets.

A vital stage of any learning task, is getting the right balance between good generalisation⁴ and fitting in estimating a model from the dataset. *Over-fitting* occurs when the model “over-learns” the training data with the noise and other irrelevant information to an extent that it becomes sensitive to minor perturbations in the dataset.

Consequently, it does not generalise well to novel data within the same distribution. Conversely, if the model is unable to learn the underlying relationships in the data such that it neither fits nor generalises well to the training data and novel data respectively, then *under-fitting* is said to have occurred.

⁴Generalization refers to how well an estimator or model is able to respond to samples similar to but not contained in the training set. This is a widely accepted benchmark for assessing the performance of ML models or algorithms.

3.2.3 Validation

In general, ML is not only about learning from data but also about utilising the learned information to make decisions when presented with similar data. ML algorithms would not be very useful without some means of evaluating the performance of a system. This is achieved via model validation where a trained model performance is evaluated with test data – a subset of the data from which the training data is acquired. The test set is used to ascertain how well the trained model can generalise – i.e. make good decisions when presented with novel data.

Ideally, the goal of the model should be to strike a good balance between under-fitting and over-fitting. To achieve this, the performance of the learning algorithm is monitored over time as it learns training data. As learning progresses, the errors on the training and test data diminish naturally. However, over-training would lead to the model learning exact information and noise from the training set while the error on the test set increases and the generalisation capacity decreases.

To avoid or limit over-fitting, a resampling procedure may be used such as k-fold cross validation where the model is trained and tested k-times on separate subsets of the training data. The model is then evaluated based on its performance across the subsets. Another means of reducing over-fitting is by reserving a *validation dataset* – a subset of the training data that may be used to evaluate the performance of the model after the training and testing procedure and indicate how the model may perform on unseen samples. Overall, the concept of generalisation is key in building object recognition models which is essentially what automatic shape annotation is.

3.3 Shape Matching

To achieve robust automatic annotation, a good model should be able to generalise from specific examples within the training set to generic samples without the set. Exploiting prior information about the shape under question is an intuitive means of guiding the model and improving the generalisation ability. The prior could be the global shape configuration and local textures surrounding the boundary of the desired object. These elements may be captured in a *deformable model*⁵ whose modes of deformation are governed by statistical parameters inferred from the distribution of the shapes in the dataset. The deformable object framework facilitates the representation and analysis of objects as a collection of salient features⁶ extracted from the dataset.

3.3.1 Feature Extraction and Detection

Feature detection is a very vast area of research; and many times, what constitutes a feature is not readily discernible due to various levels of abstraction, which are usually proprietary to an application. As mentioned in Section 3.3, a shape or object may be represented by a set of features.

In context of this project, the OOI is the femur and the features to be extracted are the landmark locations along the boundary of the femur. Automated annotation of the femur may be achieved then, by casting the problem as a feature detection one. There are several features that may be extracted from a shape [83]–[85] depending on the application and a slew of features that may be acquired from the texture of an image [86].

⁵Deformable models refer to a set of model-based techniques in Computer Vision that are used to capture the variability of a homogeneous class of objects [37], [38].

⁶A feature in plain terms, is the attribute of an object acquired through some form of measurement.

3 Technical Background

In broad terms, shape features may be *global* or *local* while textural features may be *structural* or *statistical* [87]. In the context of this project, the features of interest are:

- The global shape of the femur encoded as a feature vector whose coordinates are separate landmarks along the contour (more like a sparse representation than dense) [88]; and
- The raw local textures surrounding each landmark, where each landmark is the centre of a reference window within which the enclosed texture (patch) is extracted and concatenated in a vector of similar patches.

The rationale behind using low-level shape and texture representations is two-fold:

1. The encoding of the shape as vector of landmark coordinates, facilitates visualisation and also presents the shapes in a format readily accessible for statistical manipulation a la the PDM discussed in Section 3.1.2.
2. The acquisition of raw textures, enables an ML algorithm to learn directly from the primitive textures and construct higher level representations of the local patches.

Incorporating the global shape of the femur in detecting features ensures that only feature points within the range of plausible shapes are explored during detection while the textural component ensures that the landmarks are matched to their corresponding local patches. There are many methods in literature that have been used to explore the shape and texture attributes of an object, such as snakes/active contours which is an edge-dependent energy-based minimisation model that migrates a contour of points to-

wards the desired feature [35] they are sensitive to noisy edges and are not shape specific; the Active Shape Model(ASM) incorporates both shape and texture attributes [36], [64]–[66] however it suffers from spurious edges, gets stuck in local minima and depends on a good initialisation very close to the contour of the desired object; it also limits its search range to only textures around the landmark and doesn't exploit the full textural information [89]. The Active Appearance Models(AAM) [42], [67], [89] has some advantages over the ASM, but it is sensitive to texture variation and illumination [31], [67], [89]. In an attempt to overcome the weaknesses of ASMs and AAMs, Mitchell [90] combined the ASM and AAM to segment cardiac MR Images.

A relatively recent model combining both shape and texture information is the Constrained Local Model(CLM) [43]–[45], [53] which circumvents the disadvantages of ASMs and AAMs.

3.3.2 Constrained Local Models

The Constrained Local Model(CLM) is a framework or composite model that integrates a shape model and local pattern model for matching a statistical model shape to a novel image. The parameters of the model are deduced from the training dataset of landmark annotated images.

The variances in shape between training images are modelled using PCA or PPCA as described in Sections 3.1.4 and 3.1.5.

A function which could be a classifier, discriminant, regressor or filter is trained to encode the local texture or appearance of a patch around each landmark. The stack of the learned patches forms the patch model. Subsequently, the patch model and shape model are synthesized to form the CLMs.

3 Technical Background

In the testing phase, the CLMs are matched against the test images such that divergence between the model and each test sample is minimized.

CLMs are similar to AAMs but the difference is that the CLM uses local texture patch information surrounding each landmark and employs a non-linearised search protocol which is more efficient and reliable [44], [53], [55] than that of the AAM. Moreover, the construction of the CLM is flexible and modular such that a variety of shape models and classifiers could be integrated in its framework. For example [46] utilised the random forest classifier to learn the appearance of the local textures, [45] used SVMs to achieve the same purpose, [53] used the logistic regressor. A recent formulation uses Minimum-Output-Sum-of-Squared-Error(MOSSE) filters to capture local texture responses within a Bayesian regularisation framework [31] in tackling the facial alignment problem. Another method that employs Discriminative Response Maps to fit local textures in a 3D facial fitting problem, is yet another variation of the CLM [47].

3.3.3 Constructing the CLM

A composite shape and texture model is constructed from a training dataset of manually labelled images as described in Section 3.1.2. Texture patches around each landmark are extracted within a defined reference frame. Each patch is whitened and a vector of textures representing an image instance is formed by concatenating the whitened patches.

The combo of the vectorised patches and the normalised shape vectors underpin the construction and operation of the CLM. Mathematically, the model may be represented thus:

$$\mathbf{x} = \bar{\mathbf{x}} + \mathbf{E}_\omega \mathbf{p}_\omega \quad \mathbf{g} = \bar{\mathbf{g}} + \mathbf{E}_\tau \mathbf{p}_\tau \quad (3.3.1)$$

where \bar{x} and \bar{g} are the mean shape and texture respectively; E_ω and E_τ are the shape and texture modes respectively; p_ω and p_τ represent the shape and texture parameters respectively.

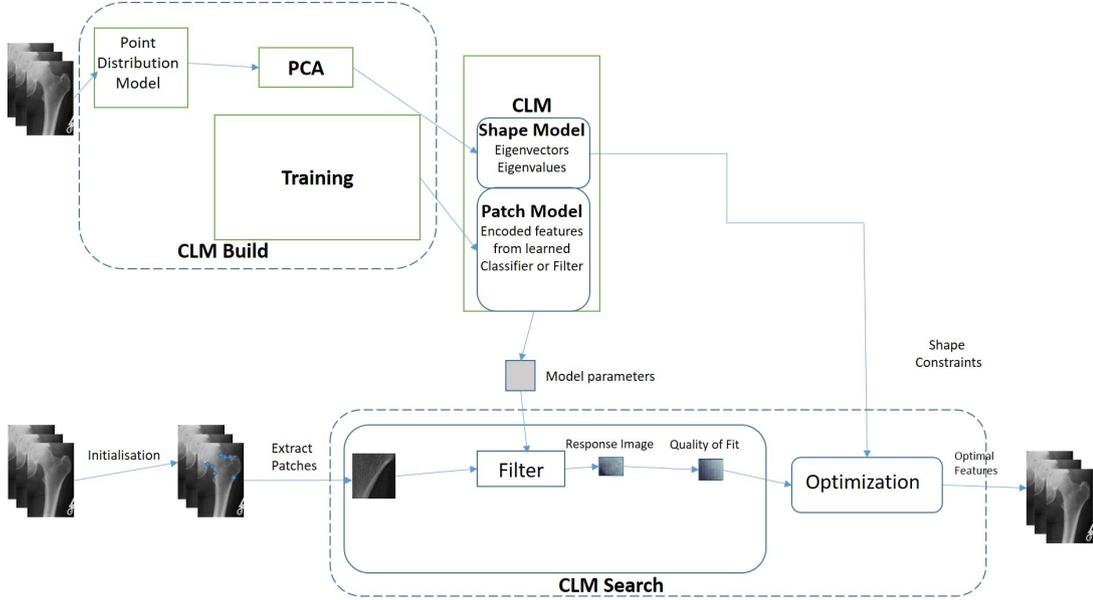


Figure 3.1: A Generic Framework of the CLM

As mentioned earlier, the CLM integrates both shape and texture models into a single model. This may be represented as

$$\mathbf{p} = \mathbf{B}\boldsymbol{\rho} = \begin{pmatrix} \mathbf{M}_\omega \mathbf{p}_\omega \\ \mathbf{p}_\tau \end{pmatrix} \quad (3.3.2)$$

\mathbf{M}_ω is the matrix of diagonal weights to compensate for the disparity in units between the shape and texture models. PCA is applied to the combined model to yield:

$$\mathbf{p} = \begin{pmatrix} \mathbf{B}_\omega \\ \mathbf{B}_\tau \end{pmatrix} \boldsymbol{\rho} \quad (3.3.3)$$

\mathbf{M} represents the combination of shape and texture modes while $\boldsymbol{\rho}$ is the vector of coordinates encoding the variation of the composite model.

3.3.4 CLM Search

After the composite model is constructed, it is used to search for the OOI in a test image. The following procedure is used to search:

- 1: Make an initial guess of the feature point positions based on the transformation \mathcal{G}_ω , landmark coordinates and shape parameters \mathbf{p}_ω .
- 2: At each estimated feature point, within the reference frame, estimate \mathbf{p}_τ and ρ by fitting equation (3.3.3)
- 3: Compute a response image $\mathbf{R}^{(i)}$ per landmark by extracting surrounding patches at each landmark and matching it to the local texture template. A matching score is assigned to each feature point where a high score implies a high response.
- 4: Update the landmark coordinates in accordance with a local high response regularised within the global shape variation constraint i.e. subject to the fitting of the shape model to maximise the match score by optimising the objective function:

$$M(\omega) = \sum_{i=1}^n \mathbf{R}^{(i)}(x_i, y_i) - \eta \sum_{i=1}^k \frac{p_i^2}{\lambda_i} \quad (3.3.4)$$

p^i are entries in \mathbf{p}_ω and λ_i are the corresponding eigenvalues of the PDM.

$M(\omega)$ implies searching for a pair (x_i, y_i) that yields a high local response \mathbf{R} while maintaining the global shape variation represented by the second term – the squared Mahalanobis distance. η is a weight that determines the extent to which a high Mahalanobis distance is penalised. η may be empirically determined from the manually labelled training dataset by computing the ratio of the second term to the first term in equation 3.3.4. $M(\cdot)$ may be solved by any non-linear optimisation algorithm although the

Nelder–Mead Simplex method is recommended [44], [55]. The optimiser drives the update of the parameters and feature points until a satisfactory minimum is achieved. The choice of similarity metric used to determine the match between the local templates and patches in the reference frame can also be varied. For example, normalised correlation [44], Mahalanobis distance [36] and Haar-like based feature classification [43].

3.3.5 Initialising the CLM

The general object detection problem requires some form of initialisation. In the medical domain, with several computer-assisted procedures, initialisation is usually performed where an operator manually sets up cues to aid a computer program in performing the desired task [91]–[94].

Deformable models also generally require an initial estimate of the feature point positions prior to searching and optimisation. Recent trends however, strive to achieve total automation of the whole detection process. In automatic landmark detection, a good initialisation is critical to the performance of the detector. To this end, a typical approach is to employ a 2-stage detection routine by incorporating a global detector prior to the final detection stage. The purpose of initialisation is to confine the search space of the algorithm and reduce the probability of getting stuck in local minima. Traditional image processing methods employ low-level routines like thresholding and edge detection to create a binary masks that assign the object's pixels or region of interest (ROI) to the foreground and other pixels to the background. While this might be effective for images with good homogeneous intensity variation, it produces unsatisfactory results for images with inconsistent gray-level distribution typical of medical images that usually have artefacts due to interference from surrounding organs and tissues which tend to obscure the desired object.

3 *Technical Background*

Template-based approaches offer some improvement by raster scanning the whole image where each pixel scanned is compared to the template and classified based on a matching score. However, changes in size, scale or orientation yield undesirable results; thus, these methods perform better when coupled with a heuristic, statistical descriptor or an ML algorithm as is the case now with many modern object detection techniques which employ sophisticated learners in a sliding window with template matching to achieve state-of-the-art object detection [46], [95], [96]. Computational advantages can also be gained by scanning the image based on a grid of locations or coordinates imposed onto the image, so that there is no need to visit and evaluate every pixel.

4 Clinical Background

This chapter covers the clinical background of the project and contains the information on the dataset and the anatomy of the hip and fractures and how computer-assisted surgery can improve the prospects of patient recovery and surgeon performance.

4.1 Extracapsular Hip Fracture Surgery

Hip fractures¹, are usually the subject of intensive orthopaedic treatment. It is estimated that 86 000 of such casualties occur yearly in the UK. Globally, the numbers have been projected to increase from 1.3 million (in the 1990s) to 7 - 21 million by year 2050 [101], [102]. Mortality rates due to a hip fracture are 5 - 10% past one month, 33% after a year while those that survive, have been observed to become disabled or experience pains [103].

Hip fractures are managed under perioperative care which comprises the *preoperative*, *intraoperative* and *postoperative* phases [104].

Computer-Assisted mechanisms may be utilised in any of these phases.

¹Also known as proximal femoral fractures: refer to fractures of the femur that occur anywhere between the articular cartilage of the hip joint to 5cm below the distal region of the lesser trochanter [97]. For more information and the different classes of fractures, see [98]–[100].

4 Clinical Background

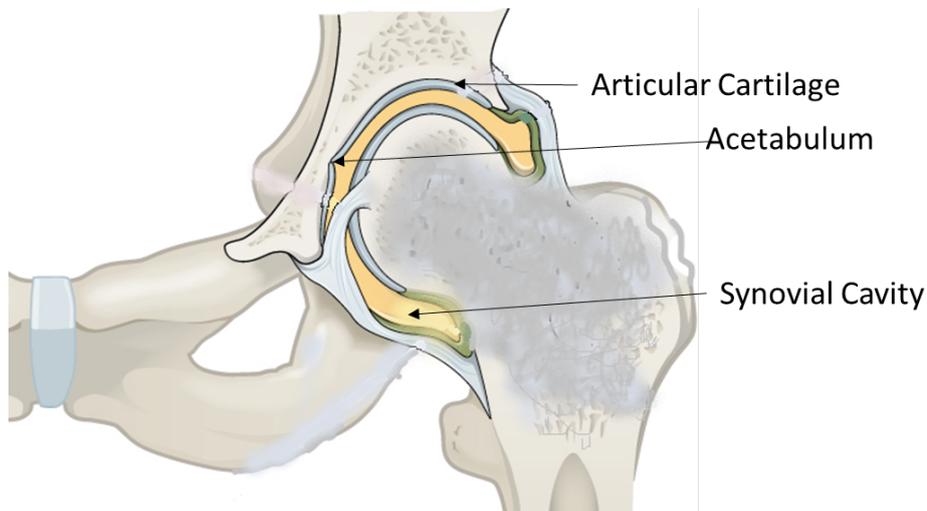


Figure 4.1: Anatomy of the Hip Joint: The hip joint is often described as a ball-and-socket joint, with the head-acetabulum structure. The synovial cavity contains fluid that lubricates the joint and allows for rotation of the head within the socket.

The *preoperative stage* is where the patient is prepared physically and psychologically for surgery. Tests are conducted to determine the best course of action. Preoperative planning usually includes computing assistance for visualisation and simulation of the actual operation.

The *intraoperative stage* is when the actual surgery is carried out in the operating ward and is actually the most critical phase, as care is taken to ensure no iatrogenic harm comes to the patient, ensuring the patient is stable and not in any form of discomfort. Autotransfusion is also performed in most cases to limit blood loss.

Postoperative care involves monitoring the patient's recovery and ensuring there is no relapse or adverse side effect after surgery.

Computer-assisted devices and software are used to perform procedures that are otherwise laborious and error-prone when performed manually.

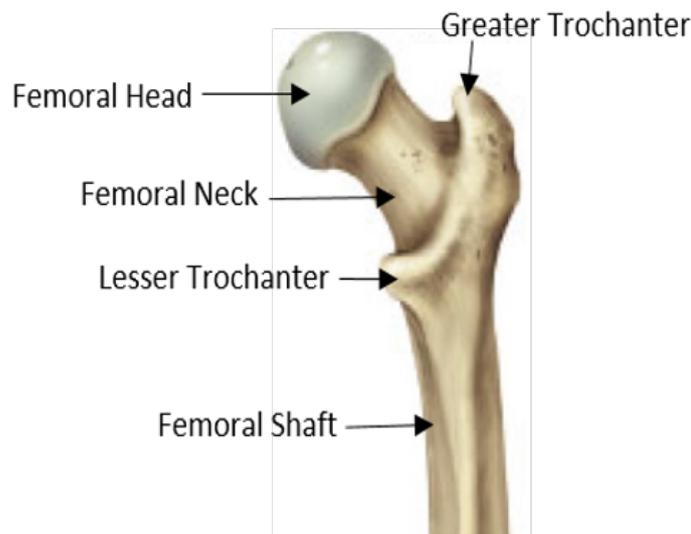


Figure 4.2: Anatomy of the Femur

4.1.1 Intertrochanteric Fractures

An *intertrochanteric* (IT) *fracture* is one that occurs within the proximal or upper region of the femur. The femoral head, neck and trochanteric areas make up the proximal femur (see Figure 4.2).

An IT fracture occurs between the greater trochanter and the lesser trochanter. The classification of fractures is important in determining suitable treatment and management for the patient [105]. For example, the type of implant required is specific to the kind of fracture.

Intertrochanteric fractures are fixed using metallic surgical implants² or nails to stabilise and secure the affected components in their original position. i.e. after any displaced fragments or bones have been bound back in place. This minimises the risk of postoperative complications.

²Sliding hip screws are the standard implant of choice when managing intertrochanteric fractures [27], [106]

4 Clinical Background

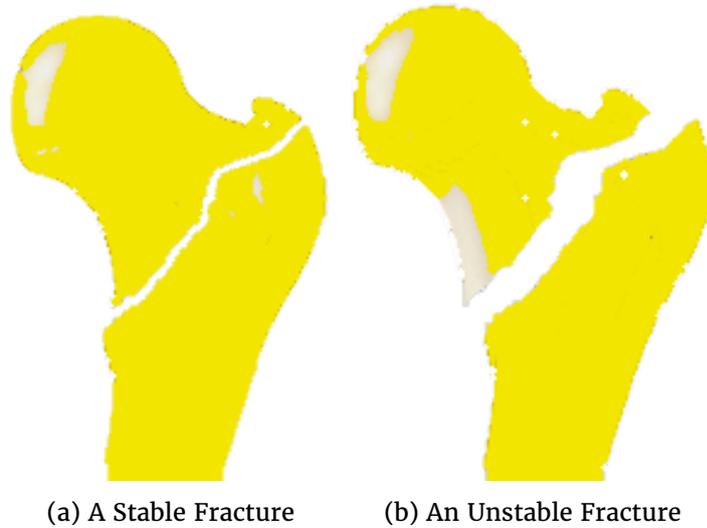


Figure 4.3: Two basic types of fractures.

4.1.2 The Tip–Apex Distance

A complication that is usually associated with IT fractures is *cutout* – a phenomenon where the inserted implant loses purchase of the femoral head and is detached from the bone. Studies have attributed this occurrence to the influence of several factors.

However, the most influential indicator according to research has been the positioning of the implant [107], [108]. To this end, a metric known as the *Tip–Apex Distance* (TAD) first formulated by [23], [24] is used to minimise the risk of cutout. TAD is defined as the sum of the distances between the tip of the lag screw and the apex(A) of the femoral head on both the anteroposterior and lateral radiographs after calibration for magnification. Fig. 4.4 illustrates the parameters and measurements used in calculating the TAD. The TAD is calculated using the following formula:

$$TAD = \left(z_{ap} \times \frac{\varnothing_{ap}}{\varnothing_{true}} \right) + \left(z_{lat} \times \frac{\varnothing_{lat}}{\varnothing_{true}} \right) \quad (4.1.1)$$

4.1 Extracapsular Hip Fracture Surgery

Where \varnothing_{ap} , \varnothing_{lat} are the diameters of the implant, z_{ap} , z_{lat} are the distances between the tip of the implant and the apex, A, of the femoral head measured on the radiograph in the anteroposterior(ap) and lateral(lat) views respectively. \varnothing_{true} is the actual diameter of the implant. It should be noted that the TAD metric is still manually estimated by orthopaedic surgeons.

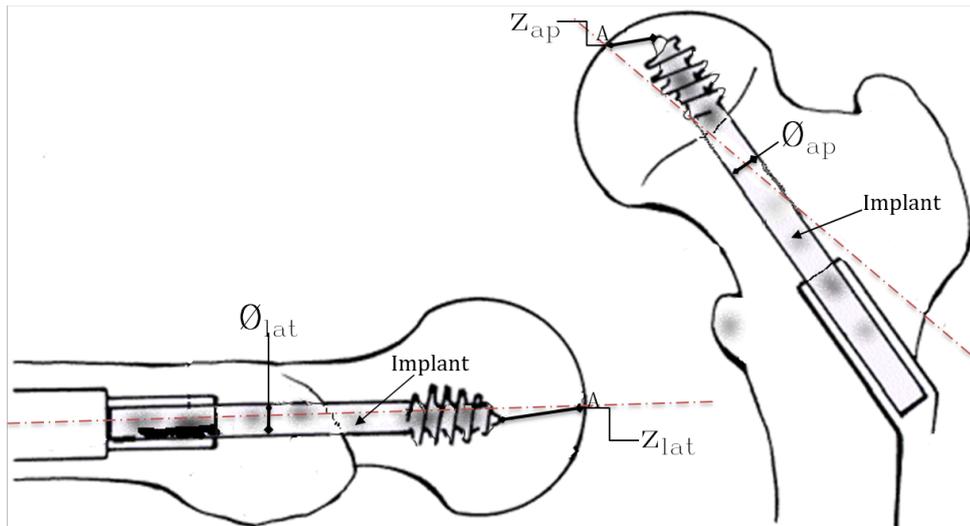


Figure 4.4: Parameters of the Tip-Apex Distance illustrated.

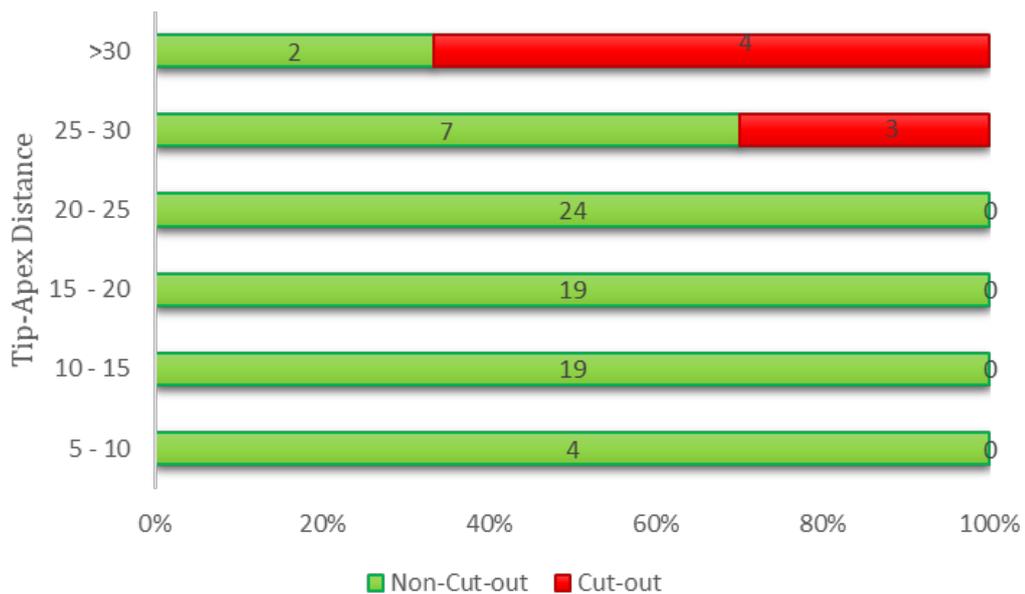


Figure 4.5: Chart showing cutout incidents vs TAD.

4 Clinical Background

The common consensus is that the implant should be placed centrally and deep [109] which lends credence to both metrics. However, the positioning of the implant still depends on manual estimation which is subjective and at best performed by the more experienced surgeons but would still leave room for error due to human fallibility. Hence the need for a more reliable and objective means of estimating an optimal position of the implant arises. [110] has already demonstrated that, the TAD can be estimated digitally from a PAC archiving system without the need for continuous reference to the live fluoroscopes.

Nowadays, computer-aided surgery is on the rise as it boosts the surgeon's productivity and enables manually difficult and laborious repetitive tasks to be performed with ease.

4.1.3 Parker's Ratio

Parker's work [109] similar to [23], [24] in that it also investigated the relationship between the positioning of the implant and the cutout phenomenon, fundamentally established a statistical basis for the priorly subjective conjecture pioneered by [111], [112] which recommends a deep and central placement of the implant within the femoral head.

From Figure 4.6, Parker's ratio, P_r , is given by:

$$P_r = \frac{AB}{AC} \times 100 \quad (4.1.2)$$

This evidently yields $0 \leq P_r \leq 100$.

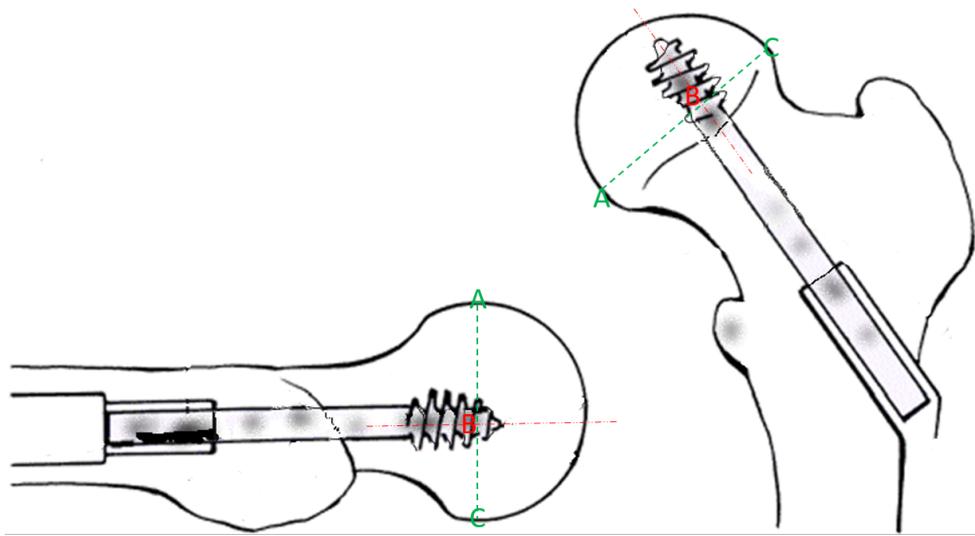


Figure 4.6: Parker's Ratio Illustrated

In Figure 4.7, the "circles" (red) represent the cutout incidents while the "crosses" (green) signify the non-cutout incidents. A dense set of crosses can be observed around the middle of the chart, which represents the ideal location of the implant.

Note that the posterior (x -axis) and inferior (y -axis) cover the range : 0 – 50 while the anterior (x -axis) and superior (y -axis) cover the range : 50 – 100.

P_r , like TAD, takes into account both the AP and lateral views on the radiograph and demonstrated that for average values of $P_r = 45$, non-eventful union occurred; cutout occurred in cases with average values of $P_r = 58$ corresponding to superior or posterior placements in the AP view and $P_r = 36$ corresponding to non-central placements in the lateral view. [26] also reported that a value of $P_r = 66$ or more corresponds to a superior/anterior position of the implant on AP/lateral radiographs which is likely to result in cutout. Overall, P_r underpins the principle of a central-deep placement of the implant.

4 Clinical Background

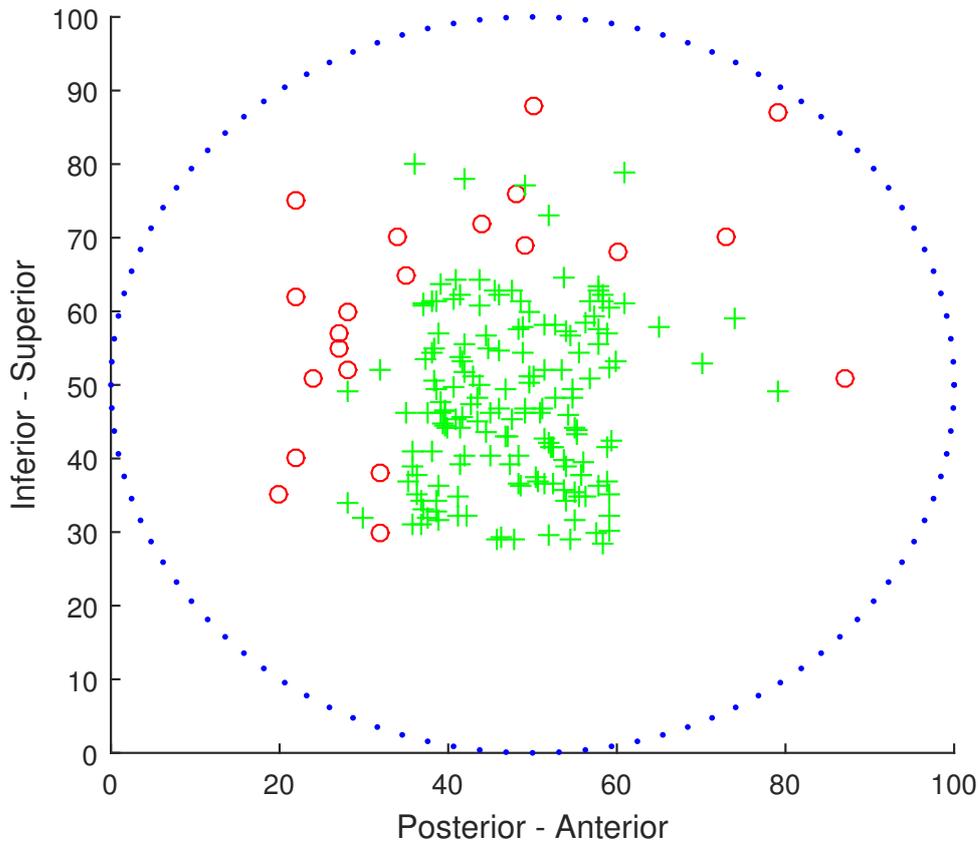


Figure 4.7: Interpretation of implant positions based on Parker's Ratio.

4.2 Computer-Assisted Orthopaedic Surgery

Computer-assisted surgery involves the use of intelligent equipment to navigate difficult procedures during intraoperative surgery. They play a vital role where precision, accuracy, repeatability and reliability are required. More often than not such procedures are risky for the patient involved and demand a high level of expertise to execute.

Accurate positioning, implant alignment, precise drilling and construction of the bone cavity to match the implant shape [113] is an instance of the use of computer guided systems to manoeuvre surgical procedures. The aim is to achieve minimally invasive treatment, ease of use, reduction in

4.2 Computer-Assisted Orthopaedic Surgery

moving mechanical parts and even preclude human intervention in some aspects as in the use of a virtual targeting device for K-Wire insertion [114] which showed improvements over freehand methods. A relatively new advancement is the integration of medical imaging modalities and spatial positioning systems. Such provisions afford a surgeon the luxury of a 'mock surgery' prior to the live operation while improving ergonomics and boosting safety levels [115]. X-ray imaging is the visual mode of choice for orthopaedics when it comes to the management of orthopaedic trauma. However, its disadvantages include prolonged exposure to radiation of both patient, surgeon and other associated personnel, scattered radiation, repeated repositioning and refocusing of the c-arm to capture x-rays [116]–[118]. The possibility of virtual fluoroscopes in lieu of live fluoroscopes have been explored to reduce operating time [119], limit the exposure to beams and improving the accuracy of positioning [120].

In the management of extracapsular hip fractures, x-rays are used to examine the nature of the fracture and establish the best form of treatment. The aim is to achieve a stable reduction and fixation of the affected limb. Ideally, a conservative approach to treatment should be adopted to avoid complications during surgery. However, for hip trauma, surgical intervention is highly recommended as conservative treatment often leads to deformity and usually delays ambulation [97], [121]. Studies have shown that the use of computer-based auxiliaries in the surgery theatre is beneficial for example [122] carried out an evaluation of a 2D fluoroscopy-based navigation system for guidewire positioning, it recorded a reduction in the number radiographs and drilling attempts required. Similar results were also recorded by [123] further highlighting the merits of computer-assisted surgery. Preliminary work by [124], demonstrates how an implant may be positioned with minimal input from an operator.

4 Clinical Background

For Computer Assisted Surgery to be successful, there must be some form of data to work with. Any computing-based system would be useless without data to analyse. The following sub-section 4.3 gives a brief on the data specific to this project.

“Data is the evidence of things sought for” – Okoli, 2017

4.3 Datasets

The thrust of this project is the image analysis of the femur on digital radiographs or x-rays. To this end, 2 datasets of images were acquired. The initial set: a collection of intraoperative digital radiographs obtained from the local NHS, RVI and the second set sourced from an online digital image bank. All images acquired were AP and Lateral views for the first set and strictly AP views for the second set.

4.3.1 NHS Dataset

The images from NHS archive were retrieved by an orthopaedic surgeon. The images were of varying quality and resolutions most of which were plain radiograph films scanned to digital format. Due to the randomness of the collation, some images were missing interesting parts – like the femoral head; such images along with degraded samples were filtered out. All images were in *.jpeg* format, free of knowledge tags and metadata, thus there was no confidentiality breach.

4.3.2 Online Digital Image Bank

The second set of images were sourced from an online stock images database³. The images were deposited without any standard protocol thus they were of varying sizes and quality, most of which were snapshots of the full pelvic region along with the hips. These images were then preprocessed to acquire the useful parts for example the hip joint/femur areas were cropped out and resized, preserving aspect ratio to avoid aliasing. All images were saved in *.jpeg* format.

It should be noted that the standard format of images utilised in the medical domain are usually stored as *Digital Imaging and Communications in Medicine* (DICOM) format which can readily be converted to other image formats such as the *.jpeg* used throughout this project. DICOM is a standard protocol universally adopted by clinics and medical practitioners. It is used to store or transmit medical images and facilitates the integration of a variety of medical imaging devices ranging from PCs, workstations, servers, scanners, printers, networking hardware to *Picture Archiving and Communication Systems* (PACS) irrespective of the vendor or manufacturer.

³www.dreamstime.com

4.4 “Interesting Parts” of the Femur

The femoral head, neck and lateral cortex are referred to as the “interesting” parts of the femur in this work. These are the main features considered when inserting an implant to fix an IT fracture. After a position for the implant has been estimated, a path is drilled through the cortex, via the neck into the head of the femur, then a guidewire is inserted through this path over which the surgical implant is then inserted.

The guidewire position is observed on the x-ray to ensure an optimal position before final insertion of the implant. See Figure 4.8(a) for an illustration of guidewire insertion. Figure 4.8(b) illustrates the ideal positioning for the guidewire i.e. insertion along the shaft axis $A - B$ via the neck axis B and through the head centre $C - Apex$ axis.

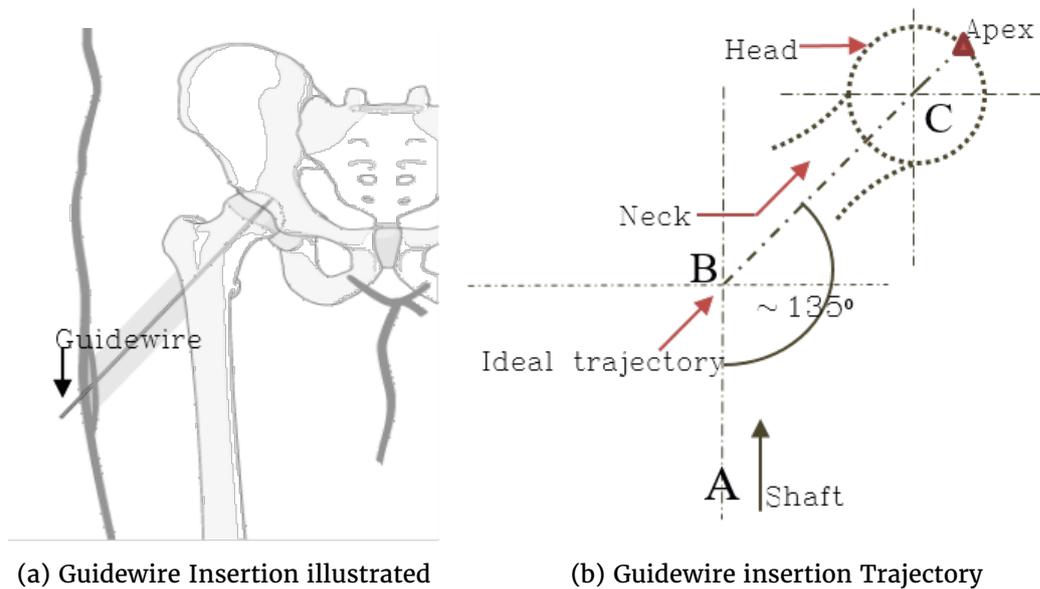


Figure 4.8: Guidewire ideal insertion positions and direction approximately 135° to the shaft (A) along $B - C$ via the neck and head centres respectively.

4.5 Project Rationale

It has been previously established in the preceding sections 4.1.2 and 4.1.3 that the positioning of the implant largely relies on the skill and/or experience of the surgeon performing the procedure and that cutout occurs due to malpositioning of the implant. Therefore one of the key objectives is to identify a means of aiding the surgeon; by precluding the manual aspect of estimating the optimal position of the implant; thereby, improving patient outcomes.

The studies carried out in this respect are geared towards the intersection between orthopaedic surgery and computer vision (CV) or medical imaging and investigates how CV may be applied in alleviating some of the problems present in current intertrochanteric fracture surgery procedures.

Using the TAD and Parker Ratio metrics, we have also been able to establish that the ideal positioning for the implant would be through a point along the shaft axis, through the middle of the neck to the centre of the femoral head. Chapter 6 presents the preliminary work carried out in this respect.

5 Methodology

This chapter presents an overview and description of the methods and techniques underpinning the implementation of the *Semi-Automatic Guidewire Insertion Planning Tool* and the *Fully Automatic Guidewire Navigation Tool* in Chapters 6 and 7 respectively.

5.1 Overview

The research gap this thesis addresses, is the challenge of replacing the existing manual, subjective positioning of implants currently employed by orthopaedic surgeons with a more objective and automatic means. The approach adopted is image analysis-based as radiographs are required to assess the position of the implant. Due to the nature of the work and the data available, we adopt a *proof-of-concept* approach, i.e. all concepts and implementations are simulated to demonstrate that the proposed methods are applicable in the real world. The actual application of the methods in clinical trials is beyond the scope of this work.

As with any image-analysis pipeline, the first stage is the acquisition of images, where each image must contain the OOI; in this project, the *proximal femur* is the OOI (see Figure 4.2). For the semi-automatic digital tool, six intraoperative x-ray images were obtained and 100 images were acquired for the implementation of the automatic digital navigation tool.

5 Methodology

After the images have been captured, subsequent image analysis techniques are applied. The assumptions taken into account are:

- The images consist of instances of the femur without fractures or implants this is important because the presence of an implant in the image will impede the performance of the navigation algorithm in the case of the automatic digital tool since the implant will ‘compete’ with the regions of interest for signal detection.
- Magnification is not taken into account in our implementations as it is only required in a live clinical demonstration.

For more details of the methodology used in the implementation of the semi-automatic tool, refer to Section 6.2.

We now delve into the details of the methods used in implementing the automatic guidewire navigation tool.

5.1.1 Automatic Guidewire Navigation Methods

The deformable models background presented in Chapter 2 is the basis for the automatic tool, specifically the CLM variation. Figure 5.1 illustrates the high-level procedure for implementing the CLM and can be coarsely defined as an integration of two modules: a *Shape Model* and an *Appearance Model*.

The aim of the developing the proposed digital tool, is to automatically compute the optimal placement for an implant subject to the constraint of the clinical metrics described in Sections 4.1.2 and 4.1.3.

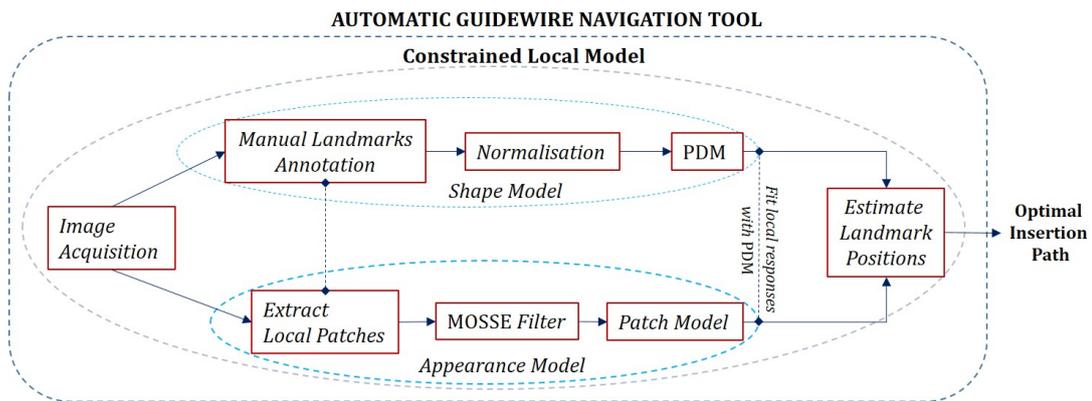


Figure 5.1: A high-level Overview showing the steps for implementing the CLM technique used for the Automatic Navigation Tool.

The challenge therefore, is to articulate these metrics within the CLM framework. To achieve this, landmarks were strategically located along the ROIs on the femur to enable measurement of these metrics based on the annotations. The manually annotated landmarks and their corresponding local patches are the input to the CLM. The expected output is an automatic annotation of a test image with the optimal path of implant insertion displayed and the corresponding P_r score computed.

5.2 Building the Shape Model

After the images have been captured containing the ROIs, a shape model is constructed using landmarks extracted from the ROIs. The landmarks are manually annotated according to predefined criteria. This section describes the concepts and processes involved in constructing the shape model.

5.2.1 The Morphology of the Femur

The term *morphology* in this project, refers to the study and analysis of the shape of the femur. This includes choice and identification of relevant parts and the annotation of the outline of the femur. For this sub-study, the second dataset of images acquired from the online database as described in 4.3.2 is utilised.

As mentioned in subsection 4.4 there are specific parts of the femur which are of interest during surgery. These are the regions that are annotated in this project. [46] carried out a ‘complete’ annotation of the proximal femur using 65 landmarks to cover the entire outline of the femur - as a precursor to a fully automatic segmentation system used to analyse radiographs as a means of diagnosing osteoarthritis of the hip joint in patients. The application in this work is towards a different purpose and will therefore, follow a different annotation scheme.

In Figure 5.2, 27 landmarks are identified whose locations are concentrated around the interesting parts of the femur, namely: The lateral cortex, neck and head. The annotation was manually carried out using the *Collection of Landmarks for Identification and Characterization* (CLIC) package [125], [126]. Figure 5.2 shows a typical instance of the annotated femur.

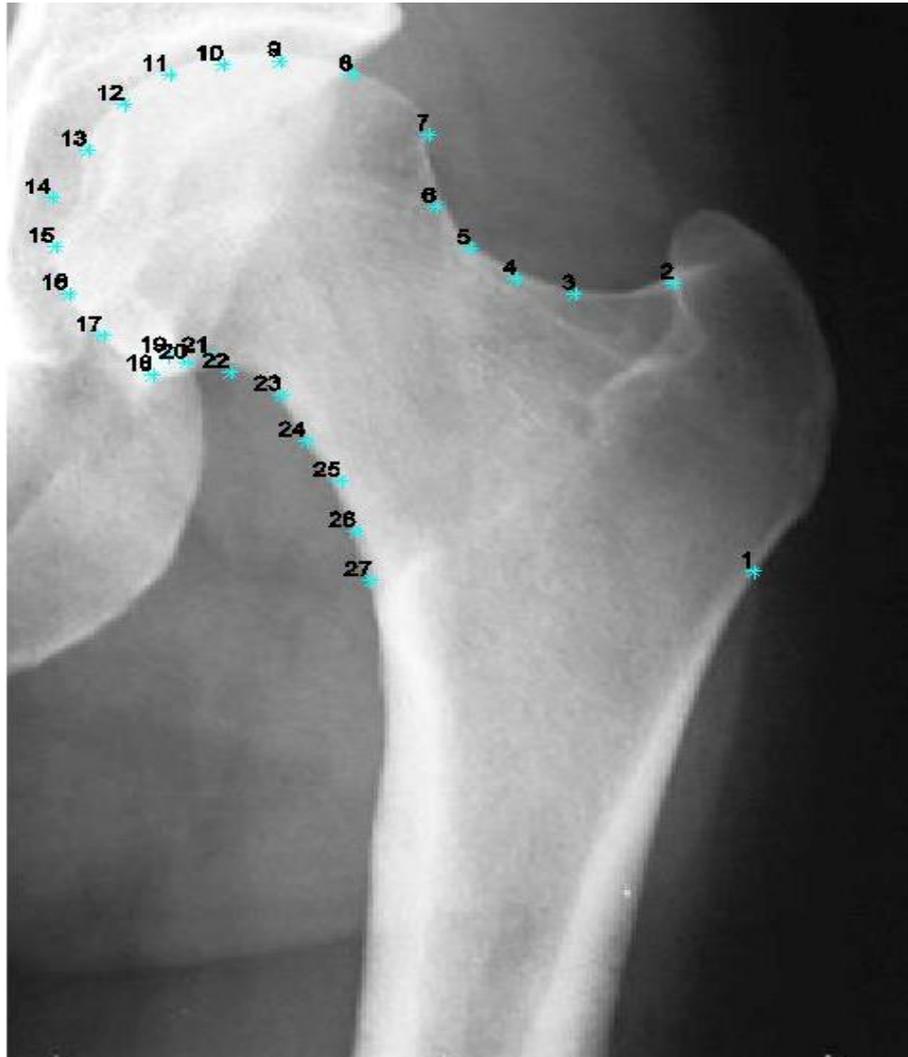


Figure 5.2: Typical Annotation of the Femur showing the landmarks labelled according to the context of this project.

5 Methodology

The selection of landmarks is an important phase in the shape modelling process as it directly influences the performance of the PDM. Hence the landmarks are placed in a consistent manner on each image to achieve a good correspondence across the dataset.

A radiograph is essentially a $2D$ representation of a $3D$ object, thus, the perceived shape may not always be “perfect” and the annotation of only distinct landmarks is likely to yield a sparse representation of the shape. To manage this, a combination of pure landmarks and pseudo-landmarks equally spaced between the pure landmarks are inserted.

5.2.2 Definition of Landmark Types

For the purpose of this work, a landmark is defined as a *context*-significant point on an object. By *context*, we refer to attributes that are specific to particular applications e.g. anatomical, biological, location, mathematical etc. The landmarks used in the annotation scheme are categorised as *application-specific*, *anatomical* and *geometric*. We briefly describe the various categories:

- *Application-specific*: Landmarks in this category refer to those that are not distinct in appearance but the location is within a ROI and makes it relevant to the task to be performed.
- *Anatomical*: Anatomical landmarks are those that have a distinct appearance by virtue of their location at a corner or protrusion on an object.
- *Geometric*: These landmarks do not have a distinct appearance or location but may be situated using mathematics or geometry. e.g. along the curved regions. They are also referred to as *pseudo-landmarks*.

Since a clinical interpretation is required, anatomical landmarks in combination with some inconspicuous landmarks take on added significance. For example, *landmark 1* in Figure 5.2 is highly significant in the context of IT fracture surgery, since this is the entry point of the implant, but it is unlikely to have a high significance in an alternative application, e.g. the full segmentation of the femur. The other significant landmarks are *landmarks 5* and *23* which represent the shortest distance between the 2 ‘concave arcs’ of the femoral neck; and *landmark 12* that represents the central axis through the neck and head centre. Due to the variability and sometimes ambiguity of the local features of the femur, particularly around the head–acetabulum region, a geometric approach was incorporated to annotate the head region by dividing the local circular contour between *landmarks 8* and *18* into equally spaced segments to achieve consistency and repeatability in the annotation. The 27 annotated landmarks and their labels are summarised in Table 5.1.

Table 5.1: Landmark labels and descriptions

Landmark label(s)	Description	Type
1	lateral cortex	application specific
2	end of greater trochanter	anatomical
3 – 6	pseudo-landmarks between 2 and 7	geometric
7	end of superior neck	anatomical
8	beginning of head	anatomical
9 – 12	pseudo landmarks between 8 and 13	geometric
13	mid-point along head contour	anatomical
14 – 17	pseudo-landmarks between 13 and 18	geometric
18	ligament of femoral head	anatomical
19 – 21	pseudo-landmarks between 18 and 22	geometric
22	beginning of interior neck	anatomical
23 – 26	pseudo-landmarks along interior neck	geometric
27	beginning of lower trochanter	anatomical

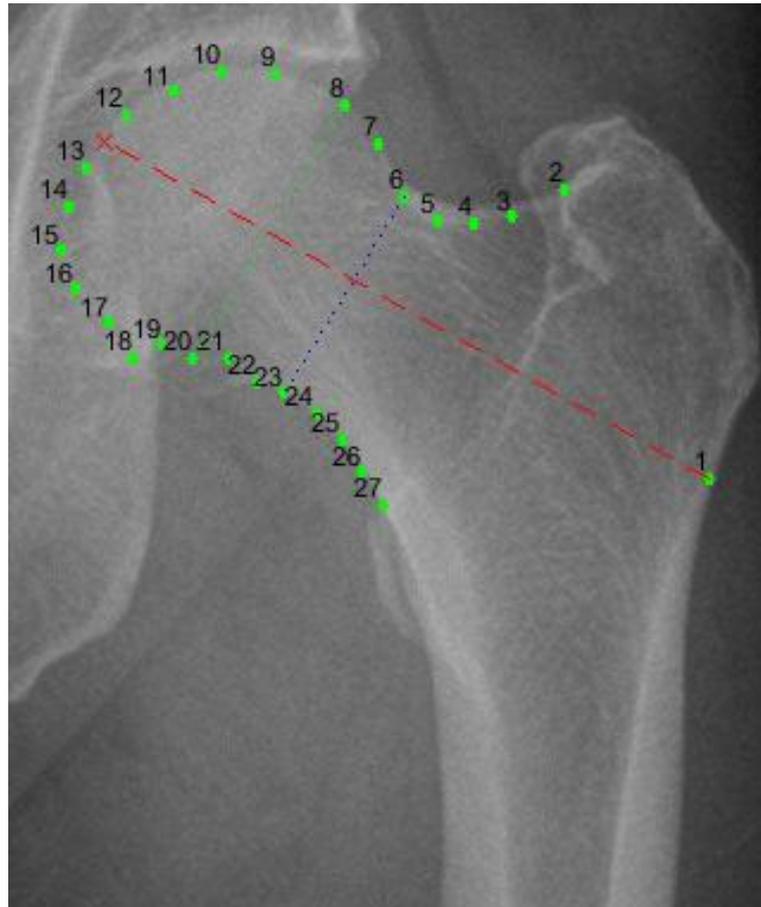


Figure 5.3: A figure showing the annotated landmarks and the axis for optimal insertion.

5.2.3 Morphometric Variables

The parameters used in this study were selected based on clinically relevant estimators. i.e. landmarks were strategically chosen to underpin the significance of both the TAD and PR which is a central-deep axis via the lateral cortex, through the centre of the neck and head. Figure 5.3 shows the axis of optimal insertion via *landmark 1*, the midpoint between *landmarks 6 and 23* through to the apex (red 'x') midpoint between *landmarks 12 and 13* of the femoral head which minimises the TAD and yields an acceptable PR score. Together these variables are used to calculate or measure PR on a given image.

5.2.4 The Shape Model

It is obvious that the annotation carried out in Section 5.2.1 and illustrated in Figure 5.3, yields a sparse outline of the OOI. Due to the variations in scale, rotation and translation across the dataset, normalisation is applied to the extracted shapes. A 2-stage normalisation procedure is applied.

Firstly, GPA is used to align the shapes in the dataset w.r.t. scale, rotation and translation. The following process describes GPA :

1. Select an arbitrary reference shape from the training dataset;
2. Superimpose the rest of the shapes in the set unto the reference shape;
3. Compute the mean shape of the superimposed set of shapes;
4. Compute the Procrustes distance between the mean shape and the reference shape, go back to step 2 until convergence is reached i.e. until there is negligible difference between the reference and the other shapes in the set.

Mathematically, The Procrustes distance is computed by minimising the closed form of equation 5.2.1.

$$|\mathcal{T}(\mathbf{x}_1) - \mathbf{x}_2|^2 \quad (5.2.1)$$

Where \mathcal{T} is given by equation 5.2.2.

$$\mathcal{T} \begin{pmatrix} a \\ b \end{pmatrix} = \begin{bmatrix} a & -b \\ b & a \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} t_x \\ t_y \end{bmatrix} \quad (5.2.2)$$

equation 5.2.1 is then differentiated w.r.t. (a, b, t_x, t_y) .

Secondly, the aligned shapes are decorrelated by setting their *mean* to zero and *variance* to one. After the shapes have been suitably normalised, a statistical model known as the PDM is used to encode a compact and constrained description of the shapes. This is usually achieved using PCA as described in Section 3.1.4.

5.3 Building the Appearance Model

An appearance model is simply a description of the local texture surrounding a landmark. The appearance model is constructed from the local patches of specified sizes extracted from each landmark. After the patches have been extracted, they are processed by a detector and a local response is computed. The aim is to learn a discrimination between a feature of interest (foreground) and its background. We previously mentioned that a number of detectors may be used to achieve this. However, the method used in this project is correlation filter-based technique, described in Section 5.3.1.

5.3.1 Correlation Filter-based Local Detector

Correlation Filters (CFs) were first introduced by [127] as a means of learning templates in the frequency domain as opposed to the spatial domain. The filter operates by convolving a square patch at its feature points, where a high response or correlation peak corresponds to a positive feature point. The canonical CF requires extensive training which makes it less attractive especially for real-time applications. However, the advent of the MOSSE filter based on an adaptive training process, has proven to be a significant improvement on the traditional CF.

5 Methodology

One of the strengths of the MOSSE filter is its tendency to cope with rotation, scaling, lighting effects and partially occluded features.

These filters are also computationally efficient.

Another important attribute of MOSSE filters is the ease with which patches may be implicitly assigned as *positive* or *negative* samples¹.

Figure 5.4 illustrates the workflow for using MOSSE filters to extract local response patterns for building the patch model.

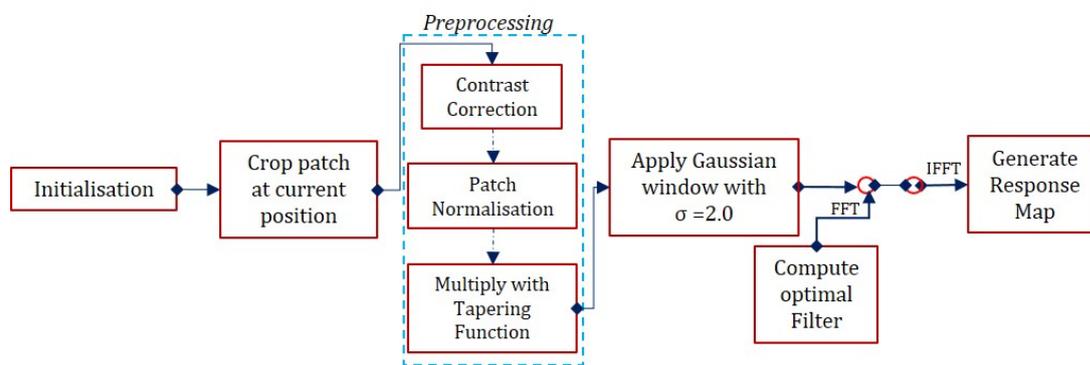


Figure 5.4: Workflow of MOSSE filter-based patch model for extracting local response patterns.

The various steps from Figure 5.4 are described next.

- *Initialisation*: The landmarks extracted from the object serve as the centre of the patches to be extracted. A typical initialisation routine is to use the mean shape coordinates.

Given that deformable models are sensitive to initialisation, more specialised routines could be employed to initialise the models as close to the optimal shape as possible.

- *Extract patch*: After the model has been suitably initialised, surrounding patches at each landmark are cropped to a specified size, the size should be large enough to capture meaningful local attributes for each

¹A positive patch is simply a local window containing the foreground features, while a negative patch is a window, not necessarily local, containing the background features.

feature point but not too large, such that local context is lost.

- *Preprocessing*: The local patches are preprocessed following extraction by firstly undergoing *contrast correction*, usually achieved by using *log transformation* given by equation 5.3.1.

$$S(i, j) = c \log(1 + |P(i, j)|) \quad (5.3.1)$$

Where $S(\cdot)$ and $P(\cdot)$ are the transformed and input intensity values respectively at coordinates i, j ;

c is a scaling constant given by: $c = 255 / \log(1 + |r|)$ and r represents the maximum pixel value in a patch. After the contrast of the patch has been transformed, the patch is *normalised* to have *zero-mean* and *unit* standard deviation; the normalised patches are then multiplied by a *tapering function*², typically a *2D cosine window* to minimise the artefacts that occur due to frequency effects at the borders during FFT operations by gradually driving the values away from the centre towards 0.

- *Apply Gaussian Window*: The preprocessed patches are convolved by a Gaussian window with $\sigma = 2.0$, to further attenuate the background pixels; such that values that fall within and without the Gaussian window, are assigned as positive and negative samples respectively. The output of the Gaussian is the desired output.
- *Compute Optimal Filter*: Let \check{x} be the patch under consideration and \check{h} be the correlation filter. According to the Convolution theorem,

$$\mathbf{x} \otimes \mathbf{h} = \mathcal{F}^{-1}(\check{\mathbf{x}} \odot \check{\mathbf{h}}^*) \quad (5.3.2)$$

Where $\mathcal{F}^{-1}(\cdot)$ signifies the inverse Fourier transform and $\check{\mathbf{x}}, \check{\mathbf{h}}$ are the transformed vectors; ‘ \otimes ’ is the convolution operator, ‘ \odot ’ indicates

²A tapering function, also referred to as a *window* or *apodization function*, is one used to *smooth* a sample signal such that the edges of the sampled region tend towards zero.

5 Methodology

element-wise multiplication and ‘*’ represents the conjugate of a vector.

A correlation output \mathbf{g} , for an arbitrary positive sample \mathbf{x}_k , is given by:

$$\mathbf{g} = \mathcal{F}^{-1}(\check{\mathbf{x}}_k \odot \check{\mathbf{h}}^*) \quad (5.3.3)$$

Hence,

$$\check{\mathbf{h}}^* = \frac{\check{\mathbf{g}}}{\check{\mathbf{x}}_k} \quad (5.3.4)$$

Note that the division in equation 5.3.4 is element-wise. Performing the operations in the frequency domain using FFT yields significant computational savings for an $n \times n$ patch i.e. $\mathcal{O}(n^2 \log n)$ using FFT as opposed to $\mathcal{O}(n^4)$ in the spatial domain. The MOSSE filter is an adaptation of equation 5.3.4 to multiple target features. MOSSE is essentially an adaptive training method that is used to compute the optimal filter \mathbf{h} from a set of training patches. This is achieved by minimising the sum of squared differences between the desired and actual correlation outputs. Mathematically, for N sample patches, the optimal filter may be written as:

$$\min_{\check{\mathbf{h}}^*} \sum_{n=1}^N \|\check{\mathbf{x}}_k \odot \check{\mathbf{h}}^* - \check{\mathbf{g}}_k\|_2^2 \quad (5.3.5)$$

The solution to equation 5.3.5 is found by solving for $\check{\mathbf{h}}$. It can be shown that [61]:

$$\check{\mathbf{h}}^* = \frac{\sum_{k=1}^N \check{\mathbf{g}}_k \odot \check{\mathbf{x}}_k^*}{\sum_{k=1}^N \check{\mathbf{x}}_k \odot \check{\mathbf{x}}_k^*} \quad (5.3.6)$$

After the optimal filter has been learned in the frequency domain, the inverse FFT is computed to yield a local spatial response, which is fitted using the constraints posed by the PDM as illustrated in Figure 5.1 and described in Section 5.2. A composite equation 3.3.4, representing the fitting process is described in Section 3.3.4. The output of the whole process is an optimal set of features, representing the landmark coordinates used to guide an optimal insertion of the implant.

5.4 Data

The dataset used in this thesis comprises digital x-ray images, the annotated and extracted landmarks and their corresponding square patches. The specific details of the images and their specifications are described in Section 4.3.

5.5 CLM Initialisation

In Sections 2.1 and 3.3.5 it was highlighted that the CLM requires initialisation as is the case with every deformable template model. A naive technique of initialisation is superimposing the mean shape onto the given image and starting the search around the coordinates of the mean image. However, this method of initialisation is only effective when mean shape parameters are very close to the optimum of the target shape. In cases where the target shape is far from the mean, the speed of convergence will be slow, or the model could get stuck in local minima, or return a sub-optimal shape. To circumvent these likely challenges, more streamlined initialisation routines are usually employed. The use of Haar features is a popular technique for most object detection problems [44], [46], [128], the use of Hough forests is also an option [129]. The output of the initialisation phase is typically a bounding box that specifies the detected object after which the CLM is used to search for the optimal configuration of landmarks within the bounding box parameters. In our implementation, we employ a machine learning technique known as boosting. Boosting is a method that combines the output of simple classifiers known as weak learners to predict or estimate a target feature. Weak learners are essentially classifiers marginally better than random guessing. Examples of weak learners are naive Bayes classifier, logistic regression, decision trees etc.

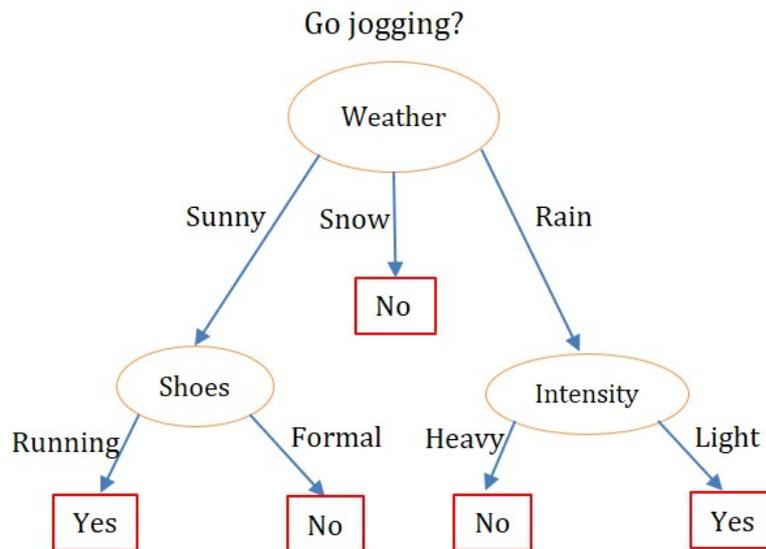


Figure 5.5: Illustration of a decision tree

5.5.1 Initialisation using Decision Trees

Decision trees are a *directed graph-like* representation of *queries* and *responses* designed in a manner to arrive at a desired outcome. The queries are usually directed towards features of interest which typically require *yes* or *no* responses. Figure 5.5 illustrates a trivial example of the tree representation consisting of *nodes* and *edges*, where the *internal nodes* represent queries and *terminal* or *leaf nodes* represent responses. In image analysis, the nodes typically represent features of interest which could be raw pixel values or extracted values from a feature descriptor such as Local Binary Patterns (LBP) [130], Histogram of Oriented Gradients (HOG) [131], Haar features, etc.

We utilise *1-level decision trees* also known *decision stumps* or ‘1R’, as the base learners in a *boosted classifier* framework³. The initialisation process is illustrated in Figure 5.6. The decision stumps are decision trees with only *one split*, i.e. just one level of queries and responses.

³Boosted classifiers are an ensemble of weak learners whose output has strong prediction capability, due to the aggregation of the outputs from each weak learner.

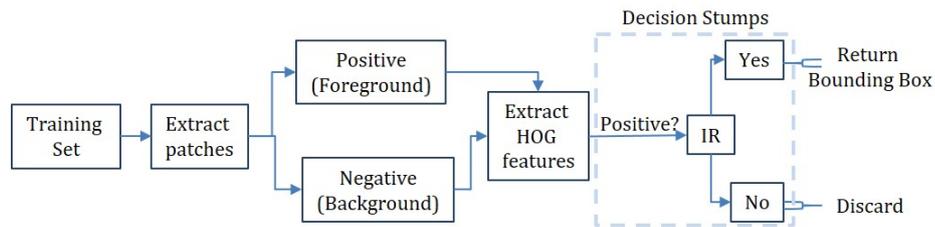


Figure 5.6: High-level overview of the CLM initialisation process

The input to the tree consists of HOG features extracted from patches within the ROIs. HOG features transform pixel-based values into gradient-based representations by accumulating the total number of gradient orientation instances within a localised window or frame such as a patch⁴. For samples corresponding to the foreground, the classifier returns a bounding box location within which the shape model in Section 5.2 is initialised. The implementation of this method can be found in Section 7.9.4.

⁴Full details of the HOG descriptor are beyond the scope of this work, but, may be found in [131]

6 A Semi-Automatic Guidewire Insertion Planning Tool

This chapter details the preliminary work carried out by the author and has been published in [124]. In Chapter 4 the significance of a deep-central placement was established. To achieve this, the implant must be inserted along an axis corresponding to the centre of the femoral head. Thus we propose a new method that computes the centre of the femoral head by fitting a circle along the boundary of the femoral head where the centre of the circle corresponds to the femoral head centre. This enables the surgeon or operator to objectively locate the femoral head centre rather than manually estimating the position.

6.1 Introduction

In the orthopaedic management of extracapsular hip fracture surgery, surgeons usually rely on visual assessments of radiographs to estimate the position of implants and often relying on intuition and experience. This is the case in the treatment of intertrochanteric fractures, where the surgeon estimates the position of the femoral head centre prior to the insertion of the guidewire over which the implant is placed. This is a key phase of the intraoperative procedure and strongly influences the postoperative outcome of the surgery.

There are already several existing digital software such as OrthoView™, MediCAD™ and TraumaCAD™ for carrying out preoperative planning [132]; while Surgix™[133] is one of the few that is utilised during the intraoperative phase of surgery.

The work carried out in this chapter lays the groundwork for the design of a digital guide which will assist the specialist in the positioning of the guidewire during intraoperative surgery.

Due to the natural shape of the femoral head which is almost circular, the aim is to use the properties of a circle to locate the centre of the femoral head such that the circle centre corresponds to the centre of the head. The significance of the centre of the femoral head is underpinned by years of research and the experience of practitioners who strongly recommend that the implant should be positioned centrally and deep [24]. This inherently supports the TAD concept described in Section 4.1.2.

It can be inferred that a path along the medial axis of the femoral head will minimise the TAD. We aim to achieve this by exploiting the shared properties of the circle and femoral head.

6.2 Method

Owing to the sensitivity and ethical requirements of the project, a simulation-based approach was adopted. The experiments carried out during the development of the digital tool required digital x-ray images; to this end, six typical intraoperative digital x-ray images - three in the AP plane and three in the lateral plane with dimensions 728 by 1036 were retrieved by an orthopaedic surgeon on the research team, from the local trauma centre - Royal Victoria Infirmary, Newcastle upon Tyne.

These were used to simulate the experiments. Any ethical obligations were fulfilled by the surgeon.

A semi-automated approach is adopted to determine the centre of the femoral head by exploiting the geometric properties of the femur. A conventional approach to the problem would entail image enhancement, segmentation and feature extraction. However, given the dimensions and quality of the images, such an approach would not only be computationally intensive but also unsuitable for real time applications. Thus a simple method employing the use of circle geometry to locate the centre of the femoral head is proposed. Figure 6.1 shows a summary of the adopted approach.

It is easy to see that the contours of the femoral head closely model that of a circle or partial circle. Circle geometry is applied to compute the centre-location of the femoral head.

6 A Semi-Automatic Guidewire Insertion Planning Tool

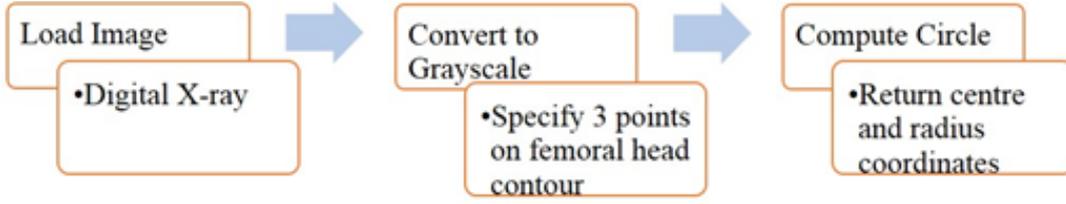


Figure 6.1: Steps to estimate the Femoral Head Centre

If we assume that \mathbf{p}_1 , \mathbf{p}_2 and \mathbf{p}_3 are 3 arbitrary non-collinear points lying on the circular contour of the femoral head in \mathcal{R}^2 , as shown in figure 6.2, setting:

$$\mathbf{u} = \mathbf{p}_2 - \mathbf{p}_1 \quad (6.2.1)$$

$$\mathbf{v} = \mathbf{p}_3 - \mathbf{p}_1 \quad (6.2.2)$$

then $\mathbf{u} = (u_1, u_2)$; $\mathbf{v} = (v_1, v_2)$; and $\mathbf{c} = (c_1, c_2)$.

Where $u_k, v_k, c_k | k \in \mathcal{R}^2$ represent the coordinates of chords $O\mathbf{u}$, $O\mathbf{v}$ of the circumscribing circle and its centre \mathbf{c} respectively. Through vector algebra and circle geometry, the centre \mathbf{c} and radius r of the circle bounding the femoral head may be computed from equations (6.2.3) and (6.2.4) respectively. L_1, L_2 in figure 6.2 are perpendicular bisectors of \mathbf{u} and \mathbf{v} respectively. Their intersection corresponds to the centre \mathbf{c} of the circle.

$$c_1 = 0.5 \frac{|\mathbf{u}|^2 v_2 - |\mathbf{v}|^2 u_2}{u_1 v_1 - v_1 u_2}, c_2 = 0.5 \frac{|\mathbf{v}|^2 u_2 - |\mathbf{u}|^2 v_2}{u_1 v_1 - v_1 u_2} \quad (6.2.3)$$

$$r = 0.5 \frac{|\mathbf{u}| |\mathbf{v}| |\mathbf{u} - \mathbf{v}|}{|\mathbf{u} \times \mathbf{v}|} \quad (6.2.4)$$

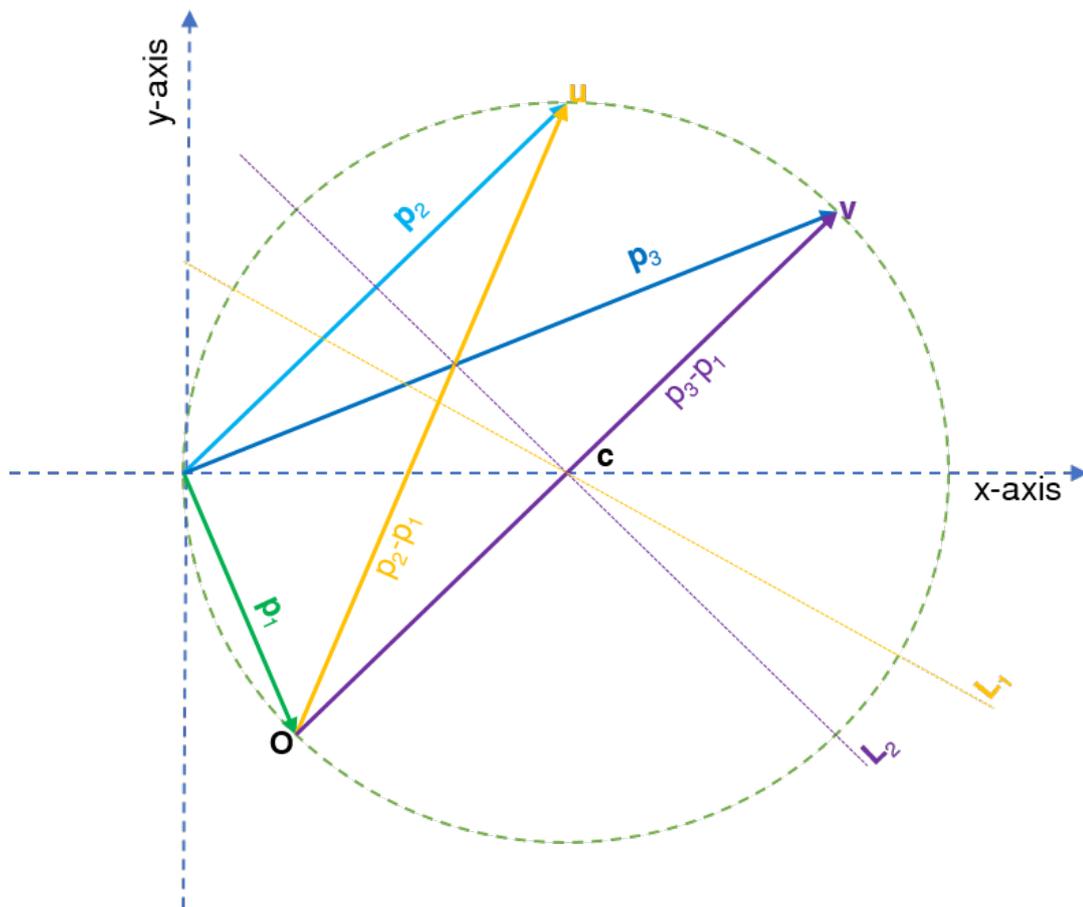


Figure 6.2: Illustration of planes and points used to find the femoral head centre.

6.2.1 Keystroke Level Modelling

Keystroke Level Modelling (KLM) is a process used in interactive computing systems to predict the time required by an expert user to accomplish routine tasks without encountering errors. The evaluation of the proposed digital tool takes into account the time required to generate a mouse click for the purpose of selecting points along the contour of the femoral head, as equivalent to the time taken to click a link or button in [134], given as 3.73 seconds. This is because the clicks are not random clicks across the image but must occur along the contours of the femoral head. Thus KLM allows us take into account both the time taken to point to a target on the display and the time taken to press the mouse.

6.3 Results

During procedures in theatre, surgeons have access to a radiographer to operate the image intensifier which consists of a C-Arm and display module tethered together by a large communications/power cable (Figure 6.3).

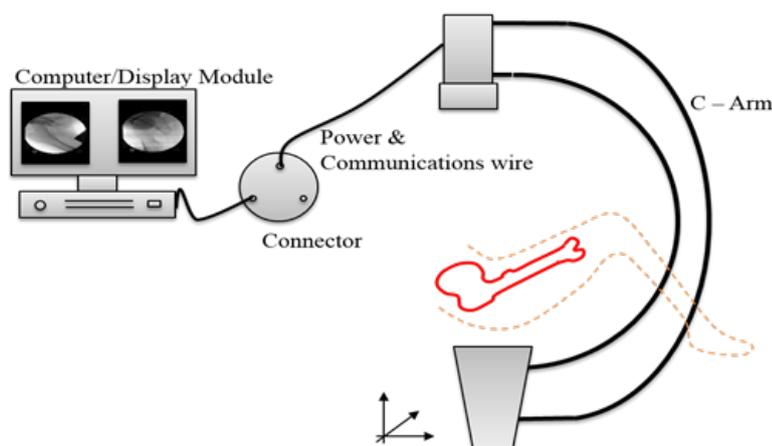


Figure 6.3: Components for image intensification during fracture surgery

Table 6.1: Typical Values Computed for Centre Coordinates and Radii, with Mean and standard deviation(SD).

	Centre (c_1, c_2)	Radii (r)
	509.24, 405.22	107.06
	506.31, 407.92	108.91
	506.33, 406.86	107.18
	505.15, 406.03	106.76
	507.72, 407.24	108.25
	511.03, 405.14	107.23
	504.32, 406.89	108.54
	510.18, 405.75	104.74
	508.67, 406.94	107.66
	505.86, 406.39	107.31
Mean	507.48, 406.44	107.36
SD	2.13, 0.85	1.09

It is envisaged that three arbitrary points will be interactively chosen along the arc of the displayed femoral head. The program then determines the femoral head centre point from the returned coordinates along this arc. The surgical procedure continues by placing an angled guide sleeve oriented towards a predefined femur-shaft angle (typically 125° – 135°) prior to guidewire insertion.

The projected circle and its centre are used a visual reference to aid the surgeon in aligning the guidewire within the centre of the femoral head. The circles circumscribing the femoral head were computed 10 times for each image, the centre coordinates and radii were recorded. The average position of the coordinates and their standard deviations were also computed.

6 A Semi-Automatic Guidewire Insertion Planning Tool

Table 6.1 shows typical values generated for the femoral head centre and radius for 10 iterations. Plots of the positions and radii and a summary of the results is shown in Figure 6.4 and Table 6.2. As deduced from the data, the determination of the femoral head centre using the proposed method shows that it is reliable with the individual values not differing significantly from the average position.

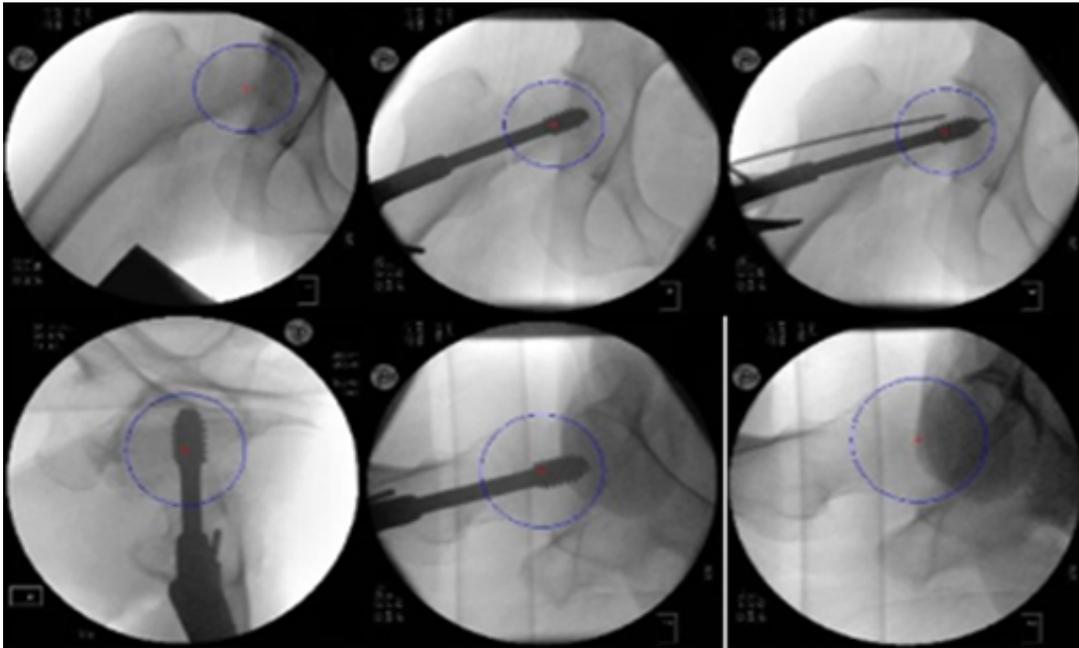


Figure 6.4: Computed centre positions of the femoral head in the antero-posterior and lateral views of the intraoperative images.

The proposed method offers the surgeon a quick visual guide to locate the centre of the femoral head which will facilitate implant placement with better accuracy.

Table 6.2: Mean and Standard Deviation (SD) of Centre Coordinates/Radii Computed for 10 Iterations

Images	Centre	Radii (r)
	c_1, c_2 Mean(SD)	Mean(SD)
IM1	507.48(2.13) 406.44(0.85)	107.36(1.10)
IM2	200.66(1.79) 392.55(1.95)	101.35(1.75)
IM3	273.73(2.02) 502.43(1.50)	140.85(1.40)
IM4	301.79(1.65) 441.26(3.56)	98.03(1.74)
IM5	281.67(1.25) 502.99(1.69)	102.78(1.60)
IM6	338.56(3.06) 519.33(3.58)	126.59(2.95)

6.4 Analysis and Discussion

The computed values of the mean and standard deviation demonstrate that the technique is highly accurate especially when the diameter of the implant is considered. It should be noted that the computed values are in pixels. The implants in the images used for this study are approximately 33 pixels in diameter which means the surgeon has approximately ± 16 pixel-“margin for error” about the medial axis (Figure 6.5).

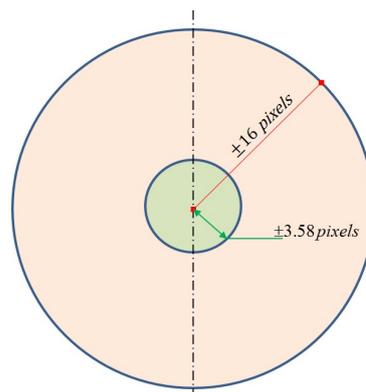


Figure 6.5: Illustration of standard deviation and its relationship to the diameter of the surgical implant.

6 A Semi-Automatic Guidewire Insertion Planning Tool

Taking into account the generated data, the largest deviation from the mean position is ± 3.58 pixels. This bodes well for accuracy purposes.

Figure 6.6 shows that the desired measurements can be measured reliably as 90% of the computed head centre coordinates fall within 2 standard deviations of the mean. Figure A.4 shows plots of the individual coordinates of the head centre and the corresponding radii of the head. The “straight lines” denote regularity in measurements of the head centre. The slight fluctuations along the lines are due to variances in the manual selection of boundary points. It should also be noted that the largest standard deviation computed for the radii is ± 2.95 pixels. This is also acceptable since the contours of the femoral head in radiographic images are usually more than a pixel thick.

During the implant insertion phase of the surgery, if the guidewire is wrongly positioned, it has to be reinserted until a satisfactory position is achieved. The position of the guidewire is observed under the image intensifier and once the implant is inserted, subsequent retractions are impossible as this could destabilise the fracture. Thus, it is vital that the guidewire is inserted correctly with the fewest number of attempts to ease the pressure on the affected region.

Overall, the potential benefits of the tool to the surgeon include less manual estimation, improved accuracy, fewer guidewire passes (ideally one pass), less wear on guidewire and surgical instruments, potentially with a lower probability of iatrogenic harm for the patient.

The digital planner, has the potential of: improving the surgeon’s operative performance during insertion of the guidewire, aiding in the training of junior clinicians and reducing surgical and anaesthetic time. The proposed technique is both intuitive and simple, thus, with minimal training, personnel would be able to grasp the operation of the tool. One of the typical challenges encountered in image analysis of radiographs also ac-

6.4 Analysis and Discussion

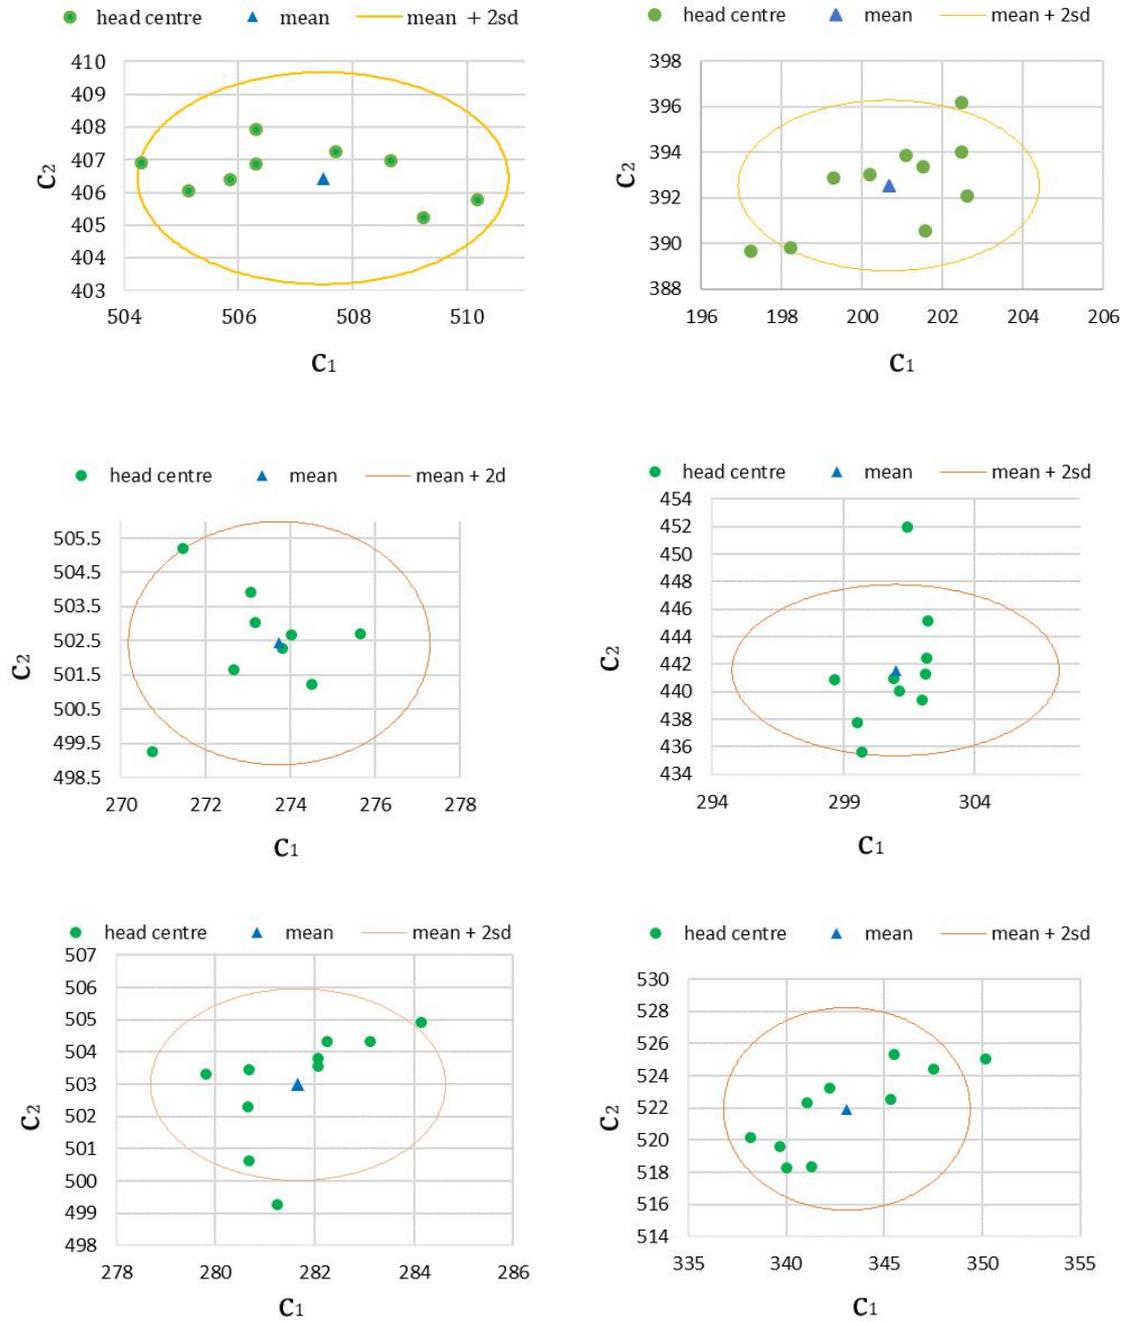


Figure 6.6: Plots of the computed femoral head centre for 6 images shown along with the mean and 2 standard deviations.

knowledge in [135], is the determination of the femoral head centre which is very difficult due to the poor contrast profile around the femoral head-synovial cavity interface and the interference by the acetabulum. However, the new proposed method circumvents these challenges even on poorly delineated edges, since it only requires three “visible” paths along the boundary.

An experiment carried out by [135] suggests that the participation of an operator in the manual selection of boundary points is laborious and reports that five points are required for determining the femoral head centre. Conversely, our proposed method requires only three points to achieve the same purpose (Table 6.3).

In [135], the author posed the problem of locating the centre of the femoral head as an ellipse-fitting problem our approach however, is posed as a circle-fitting problem (See Figure 6.7 for comparisons). Although it might be argued that more points chosen along the boundary will aid in the accuracy of locating the head centre; while using fewer points (3 in this instance) could skew the results with the presence of any outliers, it is noteworthy that the outline of the circumscribing circle would clearly indicate any abnormalities; thus, the operator could repeat the procedure to achieve a better outcome. While [135] considered only AP view radiographs of the femur in their experiments, it is worth mentioning that our work takes into consideration both the AP and lateral view radiographs of the femur.

Generally, using digital planning auxiliaries for orthopaedic surgery has several merits. One of such advantages is the cost savings as observed in a previous study on digital planning; where the financial outlay and the effects of radiation were evaluated during the application of Template-Directed Instrumentation in Total Knee Arthroplasty [136]. Although the

Table 6.3: Comparison Between Circle and Ellipse-Fitting Approaches

Femoral Head Parameters	Circle	Ellipse
Number of points selected	3	5
Actual Time taken for manual selection	<5s	– [†]
Time taken using keystroke-level analysis [†]	< 11.19s	18.65s
Time taken to process selected points	< 3ms	1s

[†] [134] predicts at least 3.73s per selection for a fast and skilled operator.

[†] data unavailable.

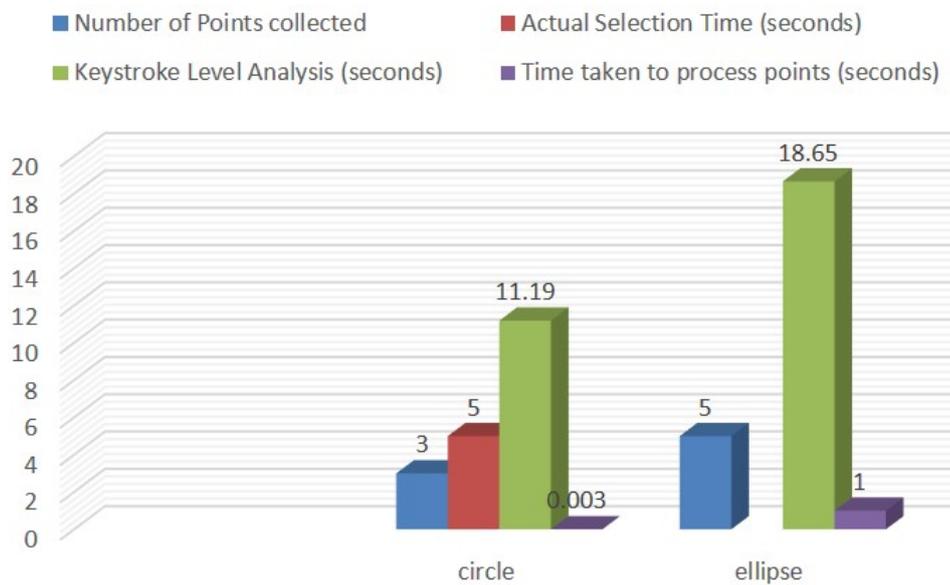


Figure 6.7: Comparison between Circle and Ellipse-Fitting Approaches

proposed method was successfully applied over a small dataset, there is an intention to build on the existing application to account for a large dataset. A potential benefit of this, is the possibility of using the tool as a high-level image technique to facilitate other low-level image analysis tasks such as: object detection, segmentation, feature learning and extraction to aid in automatic inference when presented with novel data.

The main weakness of the proposed method could stem from human fallibility and variability in measurements. But the analysis demonstrates

6 *A Semi-Automatic Guidewire Insertion Planning Tool*

that the error is negligible. Another weakness is that the speed of the process is likely to differ from operator to operator since mouse operation will vary amongst individuals. However, as noted previously, this work is the foundation for a digital tool-kit that will be expanded into a more robust application.

Re-operation or readmission due to misalignment of implants or post-operative implant migration/failure incidents occurs in 2% to 11% of fracture fixation cases [137]. Thus, the use of digital guidance intraoperatively to improve accuracy and minimise adverse outcomes is a welcome development.

The work carried out in this chapter forms the basis for a digital tool-kit to be utilised by surgeons during extracapsular hip fracture surgery. It prescribes a new and reliable method for computing and visualising the femoral head centre as a guide to positively influence the outcome of hip fracture surgery, requiring guidewire insertion during fixation. The ease of application and repeatability of the proposed method makes it a beneficial adjuvant to orthopaedic surgery. It is promising that the method is applicable even with poor quality radiographs occasionally seen during intraoperative hip surgery and still produces reliable results despite the degradation in quality.

Further research in this area will include implementation of the technique in a simulation environment with sawbone¹ models² to evaluate the influence of the technique on intraoperative parameters such as overall surgery time, radiation exposure time, the number of fluoroscopic images used and most importantly, the number of attempts and time taken for guidewire insertion.

¹www.sawbones.com

²Sawbones models refer to prototypes of real bones specially developed for use in tasks where real anatomical models are required for "hands-on" practice of surgical procedures. They are usually built out of foam, plastic or cortical shell.

7 Implementation of the Fully Automatic Guidewire Navigation Tool

The work in this chapter has been prepared as a paper for publication in IEEE Transactions on Medical Imaging under the following title: A Fully Automatic Guidewire Insertion Planning Tool for Extracapsular Hip Fracture Surgery.

In previous Sections 4.2 and 6.4 we noted that the use of automation to leverage existing manual procedures that are error-prone and laborious is one of the key advantages of computer-assisted surgical systems. This chapter describes the implementation and evaluation of a fully automated system for estimating an optimal position for the insertion of the guidewire and invariably the surgical implant.

7.1 Introduction

In the management of intertrochanteric fracture treatments, surgical intervention rather than conservative approaches is the recommended form of fixation, whereby, the fracture is openly reduced and fixed with a surgical implant. However, a recurring bottleneck is the process of locating an optimal position to avoid or minimise the likelihood of cutout. Existing procedures require the surgeon to manually estimate the optimal position for the implant. However, this is usually a trial and error process which could cause further injury to an already traumatised patient.

Two metrics known as the Tip–Apex Distance(TAD) [23], [24] and Parker’s Ratio (PR) have been shown to be significant indicators used to establish a strong correlation between the positioning of the implant and the probability of cutout. However these metrics are still manually estimated.

Computer–assisted surgery has been reported to improve implant positioning and a reduction in the number of guidewire passes required prior to implant placement [123], [133].

The purpose of this chapter is a demonstration of image analysis as a tool for estimating an optimal position for implant placement in the management of IT fracture surgery.

7.1.1 Related Work

The most similar work to this project was carried out by [135]. Although in a slightly different application regarding the femur neck fracture. It utilised a knowledge–based approach – a combination of geometry and low–level image processing techniques in the localisation of parts of the femur

on radiographs as a cue for the positioning of a surgical implant. However while this yielded the desired results, it still required an operator's interference.

Surgix™[133] is a proprietary surgical guide system IT fracture surgery. Although it is an image analysis based system, the details of the underlying algorithms are not disclosed. However, clinical trials conducted using Surgix demonstrated an improvement in the positioning of the implant and reduced radioactive exposure.

7.2 Clinical Background

During surgery, the surgeon relies on visual assessments of radiographs to guide the insertion of the implant. The TAD and PR previously described in Chapter 4 are means by which surgeons manually evaluate the positioning of an implant. TAD is defined as the sum of the distances between the tip of the lag screw¹ and the apex of the femoral head on both anteroposterior and lateral radiograph views, adjusted for magnification. It is recommended that TAD should be no more than 25mm to avoid cutout.

PR is also another metric that reconciles the positioning of the implant with the occurrence of cutout. Both metrics though different in terms of measurement, indicate a central-deep placement as the ideal position for an implant. Figures 4.4 and 4.6 in Chapter 4 illustrate the concepts of TAD and PR respectively.

While it is evident that experienced practitioners tend to have more favourable outcomes in terms of implant failure, this is not the case for more junior practitioners [24]. Hence, the need to develop a tool that enables an

¹The lag screw is the specific part of the implant that lodges in the femoral head as a means of fixing the fracture. We refer to it as an implant in this work.

objective evaluation of the implant position irrespective of the personnel performing the surgical procedure.

7.3 Technical Background

Although the advantages of computer-assisted interventions in orthopaedic surgery are evident, such systems are not readily available in IT fracture surgery. The presiding surgeon relies on a visual inspection of radiographs or x-ray images of the affected region to estimate an optimal position for the implant prior to the insertion of a guidewire.

A key component of a computer-guided system is a visualisation auxiliary. We demonstrate how computer vision or image analysis could be incorporated into the surgical workflow.

Feature detection in medical imaging is a prerequisite stage in many computer vision and medical imaging tasks. In the context of this work, the features to be detected are the landmarks. Feature detection in medical imaging is a notoriously challenging task due to inconsistencies in intensity variation caused by sensor noise, occlusions from surrounding organs, variation in patients' anatomy, etc.

However, most of these challenges can be circumvented by modelling a prior of the object(s) of interest. A popular prior is the a shape model which acts a global restraint that effectively allows only plausible shape instances or features to be generated. The shape prior may be constructed using various techniques, e.g. the minimum description length (MDL) [138], [139], Gaussian Processes [140] and Bayesian regularisation [31], [54]. In our implementation, we utilise the Statistical Shape Model (SSM) also known as a PDM which is a staple in the medical imaging community.

7.4 Encoding the Shape of the Femur

100 instances of femur were captured as x-rays, collated into a dataset which was in turn partitioned into training and testing sets each consisting of 70 and 30 instances respectively. Each femur in the training set was annotated with landmarks along the regions or boundaries relevant to the application. i.e. the head, neck and lateral cortex of the femur. Some landmarks were identified as having special significance according to the TAD and PR. These landmarks represent axis of the femur which would exploit the significance of TAD and PR. i.e. a central and deep position for the implant. The problem then arises – how to automatically detect these landmarks and then trace the optimal path for the positioning of the implant.

After the landmarks on the femur have been annotated and extracted, the next phase requires computing a representation or model that captures the variance across the training dataset. The function of the shape model is to constrain the global configuration of the landmarks by optimising the parameters of the chosen shape model. The parameters are learned from the training dataset. These parameters are usually second order statistics namely: the mean and covariance matrix. These are often sufficient to encode intra-class variability within and global variation across the training dataset. PCA earlier described in Section 3.1.4 is the conventional technique used to build the shape model, however, other Bayesian regularisation techniques exist for this purpose [31], [53], [141].

7 Implementation of the Fully Automatic Guidewire Navigation Tool

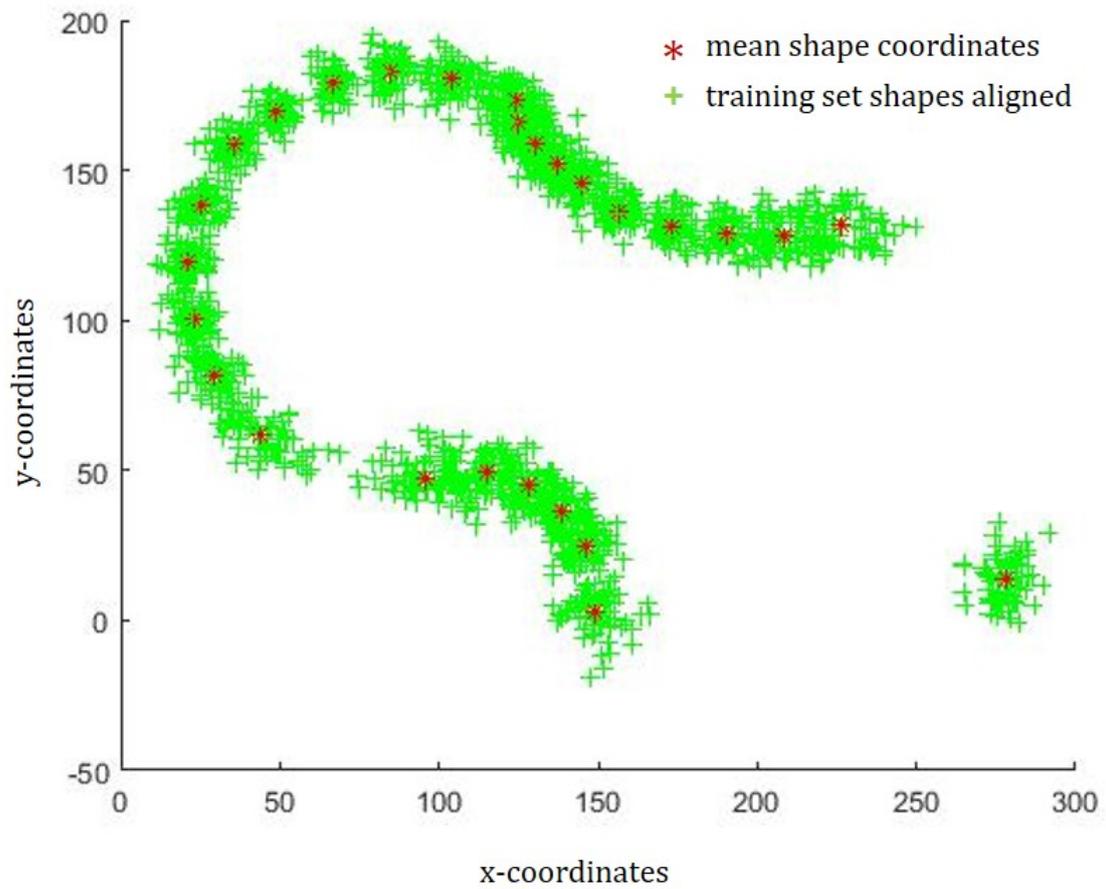


Figure 7.1: A plot of the aligned shape landmarks from the training set with the red points representing the coordinates of the mean shape.

7.5 The PCA-Shape Model

The training data comprises the shapes of different instances of the femur aligned using Generalised Procrustes Analysis after which PCA is performed on the aligned shapes to both eliminate redundancy and capture a lower dimensional representation of the data. PCA yields the parameters of the PDM used to represent the training shapes in a succinct form. PCA has proven to be a reliable model for shape representation and dimensionality reduction, but its strictly orthogonal bases means that even plausible shapes outside this subspace will not be accounted for. This inherent 'rigidity' in PCA would negatively affect its performance in estimating the more noisy samples. The probabilistic version of PCA, is more flexible in operation and takes into account the noise that is likely to be present in the data; the noise in this context is measurement error due to the manual labelling or annotation of landmarks. PPCA can be used in any instance where PCA is used.

The following Section 7.6 elaborates on the use of PPCA in shape modelling a domain where PCA has dominated for long. Recently, however, probabilistic shape representations have become more popular [31], [34], [53], [142].

7.6 The PPCA-Shape Model

The PPCA shape model is not an entirely new idea; however it has not been used anywhere in literature within the CLM framework. PPCA is regarded as a generative model that estimates a signal based on an unobserved variable; i.e. rather than directly learning a probability distribution from the data, it learns the distribution of a hidden variable that is most likely to generate the data itself.

$$\mathbf{x} = \bar{\mathbf{x}} + \sum_{i=1}^n \mathbf{B}_i \mathbf{z}_i + \boldsymbol{\xi} \quad (7.6.1)$$

$\bar{\mathbf{x}}$ is the sample mean otherwise referred to as the maximum likelihood estimate, with $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \Sigma)$ the noise or error term where $\Sigma = \sigma^2 \mathbf{I}$.

PPCA is able to model a richer probability distribution of the given data based on a linear relationship with a latent variable \mathbf{z} written as:

$$\mathbf{x}|\mathbf{z} \sim \mathcal{N}(\mathbf{B}\mathbf{z} + \boldsymbol{\mu}, \sigma^2 \mathbf{I}) \quad (7.6.2)$$

$$\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma_x) \quad (7.6.3)$$

$$\Sigma_x = \mathbf{B}\mathbf{B}^T + \sigma^2 \mathbf{I}; \quad (7.6.4)$$

Σ_x is the covariance matrix of the data \mathbf{x} . Note that as $\sigma^2 \rightarrow 0$ results in the classic PCA.

Although \mathbf{B}_i and \mathbf{z}_i in equation 7.6.1 play the same role as in PCA, (i.e. the matrix of basis eigenvectors and vector of eigenvalues); they are computed differently. See Section 3.1.5. Thus a new shape may be generated after these parameters have been estimated.

7.7 Constrained Local Models

CLMs refer to a framework that integrates a PDM and a local appearance model into a composite model. The PDM is realised by applying PCA to the set of aligned² landmarks extracted along the outlines of the objects in the training set comprising 70 instances of the femur on digital radiographs.

7.7.1 A Point Distribution Model for the Femur

Let $\mathcal{D} = \{\mathbf{x}^{(i)}\}_{i=1}^n$ represent the training set of finite landmark coordinates in \mathcal{R}^2 ; represented by n column vectors. The shape variation across \mathcal{D} may be approximated by the following linear relation, known as a PDM:

$$\mathbf{x}(\tau) = s\mathbf{R}_\tau(\bar{\mathbf{x}} + \mathbf{E}\mathbf{p}) + \mathbf{t}_\tau \quad (7.7.1)$$

where $\mathbf{x}(\cdot)$ represents the location of a landmark in the PDM with parameters $\tau = \{s, \mathbf{R}, \mathbf{t}, \mathbf{p}\}$; s represents the global scale, \mathbf{R} and \mathbf{t} are rotation and translation point-wise transformations respectively; \mathbf{p} denotes the deformation parameters. $\bar{\mathbf{x}}$ is the mean position of a landmark in the reference frame, \mathbf{E} is the matrix of eigenvectors estimated from the covariance matrix of \mathcal{D} .

Note that \mathbf{p} is defined by equation:

$$\mathbf{p} = \mathbf{E}^T(\mathbf{x} - \bar{\mathbf{x}}) \quad (7.7.2)$$

which is derived from the projection of \mathbf{x} onto the subspace spanned by \mathbf{E} . As described in Section 7.3 the PDM is used to impose a global constraint on the configuration of feature points. An exhaustive exposition on PDMs may be found in [67]. For each landmark position, c , in the PDM within a reference frame, we may rewrite equation 7.7.1 as:

²Alignment is usually carried out using General Procrustes Analysis

$$\mathbf{x}_c(\tau) = \mathcal{T}_\tau(\bar{\mathbf{x}}_c + \mathbf{E}_c \mathbf{p} + \epsilon) \quad (7.7.3)$$

Where ϵ represents the residual that determines how close the PDM approximation fits the dataset and \mathcal{T}_τ is a global transform with parameters of τ .

7.7.2 Local Texture Detectors

Without the appearance of neighbourhood surrounding the feature points, the local context of the model will be lost. Thus, an appearance model is built, using an ensemble of n independent local detectors which encode the pattern of local textures around each landmark feature point [33], [68], [143]. The PDM in 7.7.1 regularises the positions of these local detectors.

At each landmark location \mathbf{x}_i , in an image \mathbf{I} , a correlation function d_i for each detector \mathcal{C}_i may be computed as:

$$\mathcal{C}_i(\mathbf{I}(\mathbf{x}_i)) = \mathbf{d}_i^T \mathbf{I}(\Omega_{\mathbf{x}}) \quad (7.7.4)$$

where $\Omega_{\mathbf{x}} \in \mathbf{I}$ connotes the local neighbourhood in a reference frame centred on a landmark. The local detectors are used to perform an exhaustive search within $\Omega_{\mathbf{x}}$ from which texture exemplars are learned.

A number of detectors may be used to learn the local exemplars³ including haar-like features [46], [128], Gaussian likelihood [39], linear logistic regressor [31], [45], [54]. The outputs may be probabilistic as in [31], [45], [53] or make use of similarity metrics [39], [43], [46].

A recent formulation exploited the correlation filter specifically the Minimum-Output-Sum-of-Squared-Error (MOSSE) filter [31] in extracting response

³These local exemplars are in effect local response images or maps, characteristic of the surrounding landmark textures.

signals from local texture patches. We favour this approach in our work due to the simplicity and ease with which filtering operations can be applied to images using convolution.

7.8 Correlation Filter Based Detection

From an object detection perspective, correlation filters are essentially 2D kernels or templates that may be used to convolve an image at each of its pixels. As is expected of filters, the idea is to capture a high response in the regions of an image that exhibit a high degree similarity to the kernel. However, the template has to be constructed in such a manner that it captures the optimal characteristics of the area under consideration to achieve good discriminatory performance.

Correlation based filters also tend to exhibit excellent performance even in the presence of partial occlusions, rotation, scaling and lighting effects. Another key advantage of correlation filters is their significant computational efficiency where it takes $\mathcal{O}(ND \log D)$ in the frequency domain and $\mathcal{O}(D^3 + ND^2)$ in the spatial domain to perform computations.

While conventional classifiers require an explicit assignments of positive and negative samples, correlation filters do not require explicit associations between positive and negative samples. Rather, these filters assign the peak of the output signal as the positive sample and the surrounding areas as negative samples. The task then, is to estimate the parameters that minimise the difference between the ideal response of the training image and a test image.

There are several variants of correlation filters employed in literature for mostly tracking objects such as the Polytypic Sum of Squared Errors (POSSE)

filter [144] which is an ensemble of vector-valued filters within the MOSSE filter framework; [145] proposed another variation that significantly reduces ‘boundary effects’ in video tracking and a multi-channel spatial derivation of the MOSSE filter has also been formulated by [146] which casts the optimal filter learning task as a ridge regression problem in the spatial domain.

We delve into the specifics of the MOSSE filter in the following section as this forms the basis for the local patch models in this work.

7.8.1 MOSSE Filter-Based Local Detector

The MOSSE filter first proposed by [61] in a visual object tracking framework, has shown impressive performance as a standalone detector and in tandem with other detectors [31].

MOSSE computes an optimal filter that minimises the Sum of Squared Differences (SSD) error between the ideal correlation output and a test image.

A convenience afforded by this approach is that the correlation may be performed in the frequency domain as a component-wise multiplication⁴ between the 2D Fast Fourier Transform \mathcal{F} of an input image I and another filter, H . The relation may be expressed mathematically as:

$$G = \mathcal{F}\{I\} \circ H^* \quad (7.8.1)$$

where ‘ \circ ’ denotes a Hadamard product and “ $*$ ” a complex conjugate. The correlation output is then expressed in the spatial domain by applying the inverse of the Fourier transform, i.e. $\mathcal{F}^{-1}\{G\}$. In effect, MOSSE learns an optimal map, H between N training images and the ideal output response G .

⁴According to the Convolution Theorem.

Formally this may be written as:

$$\underset{H^*}{\operatorname{argmin}} \sum_{n=1}^N \|\mathcal{F}\{\mathbf{I}_n\} \circ \mathbf{H}^* - \mathbf{G}_n\|_2^2 \quad (7.8.2)$$

where $\mathbf{G} \in \mathbb{R}^{k_1 \times k_2}$ is usually deduced as a uniformly sampled 2D Gaussian with μ at the true landmark location and a sufficiently small value of σ^2 .

\mathbf{H}^* is found by solving:

$$\mathbf{H}^* = \frac{\sum_{n=1}^N \mathbf{G}_n \circ \mathcal{F}\{\mathbf{I}_n\}^*}{\sum_{n=1}^N \mathcal{F}\{\mathbf{I}_n\} \circ \mathcal{F}\{\mathbf{I}_n\}^* + e} \quad (7.8.3)$$

where e regularises the division operation to preserve the uniqueness of the division operation or avoid a *division-by-zero* error.

To circumvent the effects of artefacts due to the FFT operation, each sample \mathbf{I}_n is preprocessed prior to filtering. The preprocessing steps as described in Section 5.3.1 are:

1. Contrast correction by a log transformation;
2. Normalisation of pixel values so that \mathbf{I}_n has a *mean* 0 and *norm* 1;
3. Multiplication of the image with a tapering function which is usually a cosine window that gradually phases values close to the edge towards zero; and has the added advantage of highlighting values at the centre of the target object.

The MOSSE filter \mathcal{C}^M , is also a linear based detector and may be written as:

$$\mathcal{C}_i^M(\mathbf{I}(\Omega_{\mathbf{x}_i})) = \mathcal{F}^{-1}\{\mathcal{F}\{\mathbf{I}(\Omega_{\mathbf{x}_i})\} \circ \mathbf{H}_i^*\} \quad (7.8.4)$$

where \mathbf{H}^* is the same as the MOSSE filter computed from equation 7.8.3.

7.9 Experiments

This section describes the methods and results of the application of the CLM to the dataset.

7.9.1 Dataset

The dataset used in the experiment consists of 100 digital x-rays of the femur in the AP view acquired from an online database as described in 4.3.2.

The dataset was randomly divided into 2 subsets: a training set of 70 images and a test set of 30 images. The test set were used as ground-truth to evaluate the performance of the automatic landmark detection system. Each one was annotated with known landmarks corresponding to the annotation scheme in Section 5.2.1. 27 landmarks per image were extracted and stored in a 54×70 matrix used to build the PDM while the corresponding local patches were concatenated in a 1024×27 matrix of local intensity values per image.

7.9.2 Applying the CLM Detector

The goal of the CLM is to align the local response of each filter to its optimal location. In this section, 2 techniques are formulated to achieve optimal global alignment and fit the local feature responses.

2 Statistical models previously described in Sections 3.1.4 and 3.1.5 are the basis for the alignment of the global shape and are combined with the MOSSE filters to perform feature detection. We coin these as PCA-MOSSE

and PPCA-MOSSE to reflect the use PCA-based PDM and PPCA-based PDM as global constraints. The PCA based PDM is the standard variant applied in several literature and already covered in 3.1.2. The PPCA variant is developed in this section.

7.9.3 CLM Search

After the CLM has been constructed, the next task is to locate novel instances of the learned features from images within the test set. The search process per landmark, may be summarised in the following few lines:

1. Make an initial estimate of the feature position;
2. For each feature point, apply a local detector to search the neighbourhood of the landmark, from which a local response image is extracted;
3. Apply a fitting function to each response image;
4. Optimise the fitting function and shape/pose parameters, τ to locate the best feature point;
5. Do steps 2 to 4 until the convergence criterion is met.

The search procedure is initiated within a standard window centred around each feature point; every pixel response within this frame is evaluated by the local detector, according to the quality of fit. The computed responses within each window are aggregated to form a response image corresponding to a particular landmark.

Note that the highest response is not necessarily the optimal feature point since it is only based on the local search. The PDM is used to constrain the

7 Implementation of the Fully Automatic Guidewire Navigation Tool

response locations within the space of plausible shapes. This is achieved by iterative optimisation of the following equation:

$$\mathbf{x}(\boldsymbol{\tau}) = \sum_{i=1}^n \mathcal{C}_i^{\mathcal{M}}[\mathcal{T}_\tau(\bar{\mathbf{x}}_c + \mathbf{E}_i \mathbf{p} + \boldsymbol{\epsilon})] \quad (7.9.1)$$

The search is conducted within a specified radius r at each position then equation 7.7.3 is used to solve for τ with $\epsilon = 0$ and update each landmark position.

In our experiments, we used a search window of 128×128 in the training phase but reduced to a radius of 32×32 during the search in the testing phase⁵ which is somewhat analogous to the relaxation employed by [46] in the segmentation of the proximal femur.

7.9.4 CLM Initialisation

As mentioned earlier in Section 3.3.5, the CLM requires initialisation like every other part-based or deformable model; to reduce the computational burden and also the possibility of getting stuck in local extrema.

Since the upper femur is the main ROI and lies in the upper half of the image, any initialisation is performed within this region. Ideally, the femoral shaft region should be the best candidate for initialisation, but our PDM doesn't take into account the shaft region. Thus, we found the neck-greater trochanter region to be the most suitable due to the relatively strong contrast between the foreground and the background unlike the head region where the acetabulum response competes with that of the actual contour of the femoral head.

⁵[31] used a similar approach though it was a 40×40 window. We use 32×32 since x-ray medical images are more noisy and completely different from facial images.

The pose of the femur across the dataset guarantees that the exterior neck-greater trochanter region would be ‘north-eastwards’ in the image. Note that since the the femur pose is fairly consistent (i.e. upright in the AP view), a coarse initialisation is sufficient as opposed to [46], where a rigorous global initialisation routine was employed to account for the large variance in pose across the dataset.

A boosted classifier⁶ was trained on Histogram of Oriented Gradients (HOG) features to return a bounding box of the femoral neck area. The base learners consist of Decision Stumps. The CLM was then initialised within the bounding box locations.

Table 7.1 shows the detection performance of the classifier which yielded a True Positive Rate (TPR) of 90%, a False Positive Rate of 83.3% and a True Negative Rate of 91.7%.

Figure 7.2 shows examples of successfully detected femoral neck region with associated bounding boxes while 7.4 shows instances the failed localisation attempts.

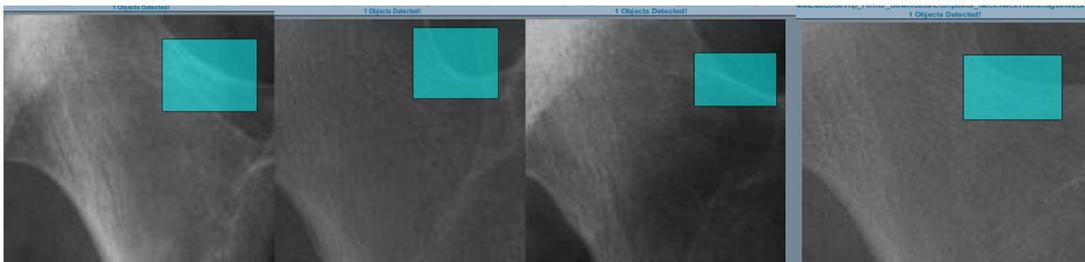


Figure 7.2: Successful instances of the localisation of femoral neck patches.

⁶A standard implementation of the boosted classifier is readily available in MATLAB’s Computer Vision System Toolbox.

7 Implementation of the Fully Automatic Guidewire Navigation Tool

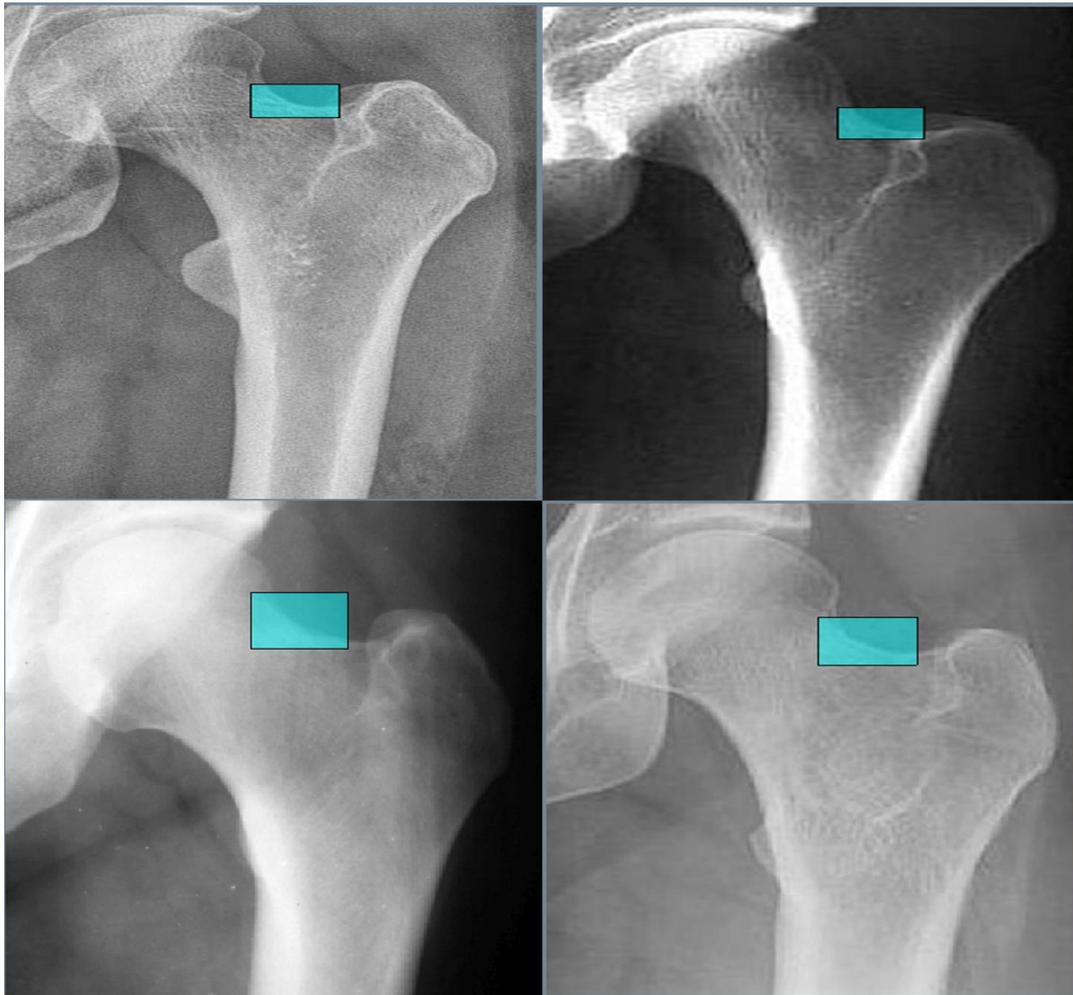


Figure 7.3: Successful neck localisation instances on the full femur.

Table 7.1: Summary of results showing the True Positives (TP), False Positives (FP) and True Negatives (TN) from the detection of the femoral neck region using Decision Stumps .

TP	FP	TN	Positive Samples	Negative Samples
9	1	11	10	12

On inspection, it was observed that the failure was due to the texture pattern in one of the images which resembles artificial textures not usually associated with such medical images, while the other instance is an example of an image with a poor contrast profile such that the foreground and background are almost indiscernible.



Figure 7.4: Unsuccessful Detection of the Femoral Neck Region

7.9.5 Automatic Landmark Detection

After the initialisation, the search for optimal features were carried out using the PCA-MOSSE and PPCA-MOSSE detectors, the optimal landmark locations were returned and plotted on the corresponding test image.

Since the 2 detectors use the same patch model, the response images at each landmark are essentially the same.

Figure 7.5 shows a sample image and the detected landmarks while 7.6 illustrates the detections with a subset of its landmarks and corresponding responses.

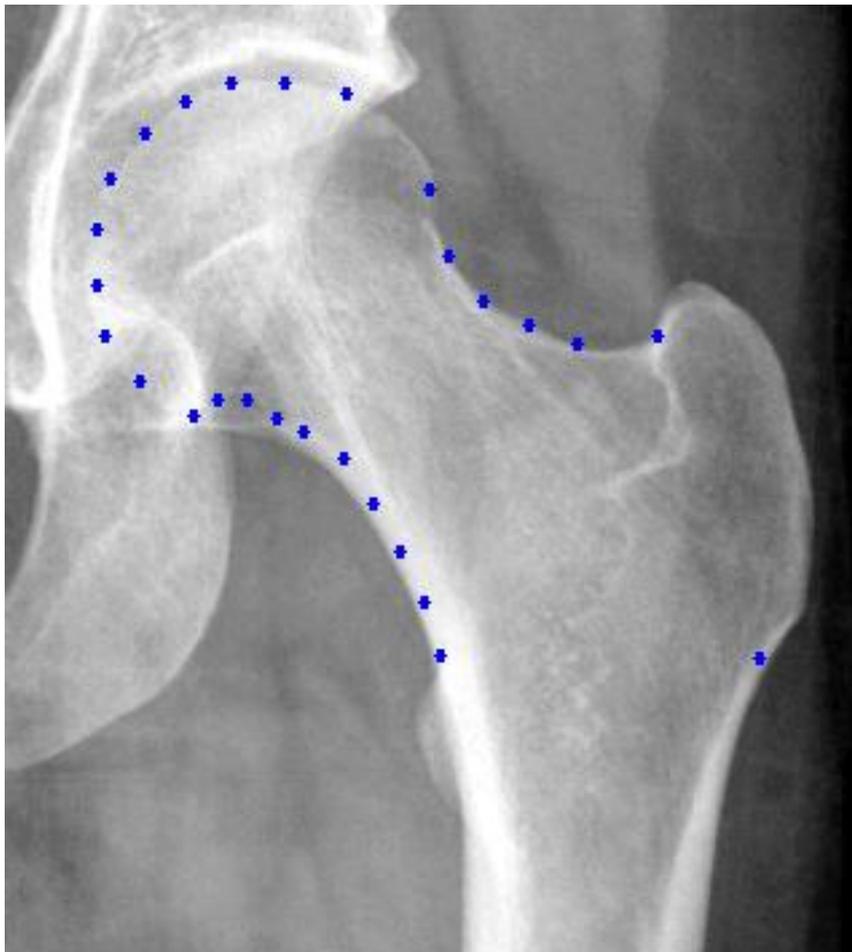


Figure 7.5: The Landmarks detected by the CLM

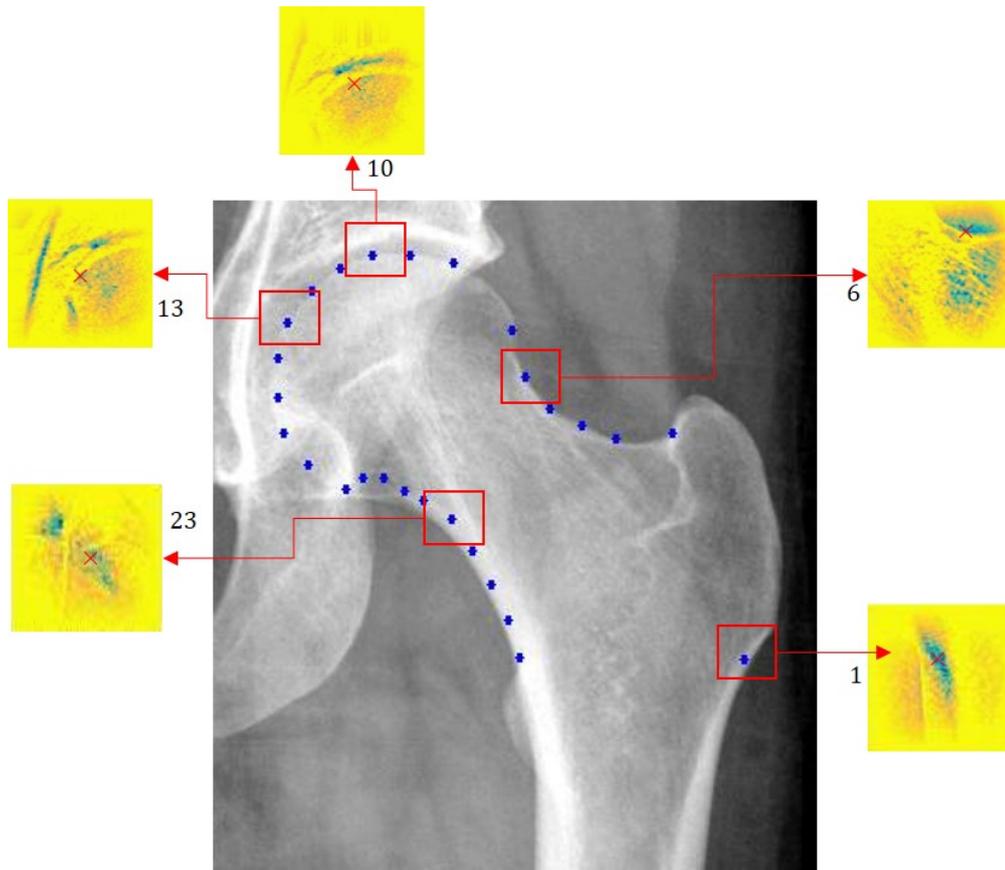


Figure 7.6: Illustration of responses from *landmarks* 1, 6, 10, 13, and 23

As this is a project that requires significant clinical interpretation, the detectors were assessed in terms of the clinical connotation. Parker's Ratio metric P_r was used to evaluate the positioning of the detected axis of entry for the surgical implant.

P_r was calculated for 26 of the images, which yielded an average of $P_r = 53.64$; well below the threshold indicated by [26]. 4 of the images failed to yield any detections due to the degraded nature of the images.

7 Implementation of the Fully Automatic Guidewire Navigation Tool

A well known challenge encountered in the processing of medical images particularly in x-rays is the inconsistency of the contrast profile of the images.



Figure 7.7: Thresholded image of the femur showing a merge of the head and acetabulum due to the poor contrast profile around the region.

This happened to be a problem around the femoral head region of the acetabulum and the ball joint. There is always a high tendency for the acetabulum to be detected rather than the femoral head itself. Figure 7.7 illustrates the problem, where a simple thresholding operation causes the head-acetabulum region to merge and lose its distinct structure. This is a major reason why ‘naive’ low-level imaging processing tasks perform poorly on x-ray images.

In Figure 7.6 this effect can be seen around landmarks 10 and 13, where the acetabulum response is the highest in the local region. The use of a PDM as a prior helps in circumventing the challenge by guiding the convergence away from the acetabulum towards the head. This is possible because the locations of the acetabulum are not encoded in the PDM.

Table 7.2: Parker's Ratio P_r^{apex} , P_r^{12} and P_r^{13} computed for 10 images at the apex, *landmarks* 12 and 13 respectively.

	P_r^{apex}	P_r^{12}	P_r^{13}	P_ϕ
	53.21	54.62	60.93	2.25
	52.19	53.38	59.48	2.26
	52.40	54.31	59.35	2.25
	52.72	52.60	60.47	2.26
	52.47	54.48	61.83	2.25
	52.74	52.11	60.00	2.27
	55.23	52.97	60.72	2.26
	54.73	52.17	59.66	2.27
	52.79	52.34	61.21	2.26
	54.26	54.90	59.75	2.25
Mean	53.27	53.69	60.42	2.26
SD	1.07	1.09	0.82	0.0073

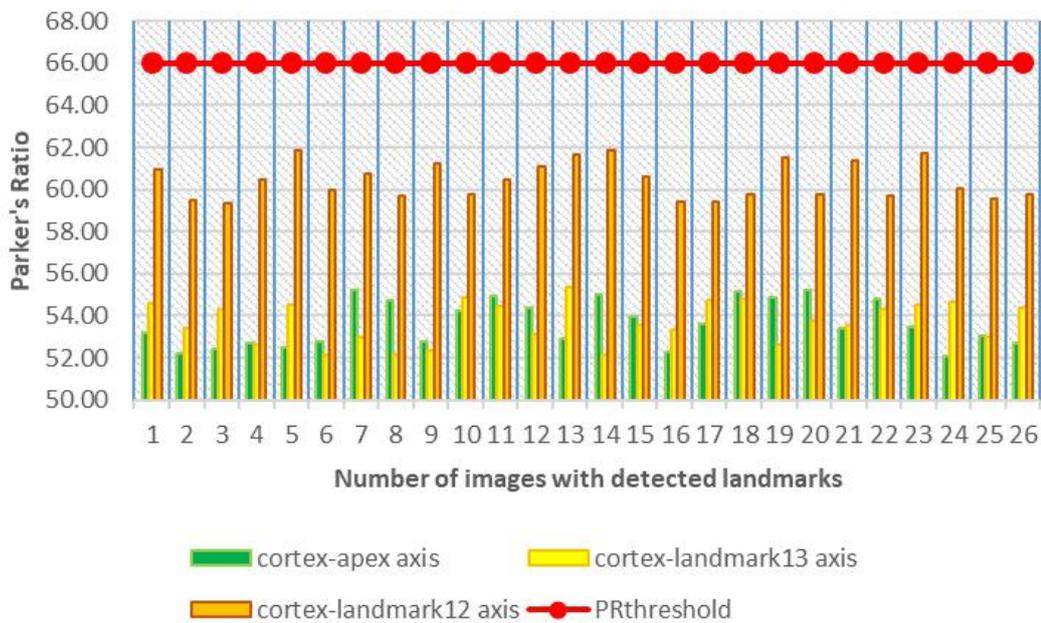


Figure 7.8: Chart showing the values computed for Parker's Ratio based on the landmarks detected. It shows how the estimated values are well below the threshold indicated by Pervez [26].

7 Implementation of the Fully Automatic Guidewire Navigation Tool

Table 7.2 shows P_r calculated for the apex, *landmarks* 12 and 13 with average values of 53.27, 53.69, 60.42, and 2.26° respectively; Note that the P_r averages of the ‘displaced axes’ are still below the threshold. P_ϕ is the average angle between the axes formed by the apex and *landmarks* 12 and 13;



Figure 7.9: The depiction of the axes of insertion based on the *landmarks* 1, 12 and 13

The ‘displaced axes’ refer to the dotted lines from *landmark* 1 to *landmarks* 12 and 13 (See Figure 7.9).

The angles between these axes and the central axis were calculated to be approximately 2°. Intuitively, an increase in this angle either side of the central axis will increase both P_r and the TAD and likely result in the occurrence of cutout, similar to the Calcar-referenced Tip-Apex Distance

(CalTAD) for intertrochanteric fractures as reported in [147] and [148].

Although there has been no significant research or literature that draws a correlation between P_ϕ and the cutout phenomenon, our hypothesis is that the axis of insertion into the femoral head should not be more than 2° from the cortex–apex axis to forestall incidents of cutout.

Figure 7.10 shows the results of using the PPCA shape model to estimate ground truth and Figure 7.11 shows the results of using the PCA shape model to estimate ground truth. On observation, the PPCA model exhibits narrow limits (*upper and lower*, i.e. UOA and LOA) of agreement: ± 2.44 and ± 1.34 in Figures 7.10(a) and (b) respectively, implying a small bias and therefore, a strong equivalence to the ground–truth. The PCA model tends to have wider limits of agreement particularly evident in 7.11(b), ± 24.06 . This implies that the PPCA model estimates are closer to the ground–truth values than that of the PCA.

The PPCA model shows an even spread of the landmarks while the PCA model tends to be more ‘compact’ and ‘cluster’ in certain spaces. This tendency is particularly evident in Figure 7.11(b) and is likely due to the inherent nature of both models; where the PPCA model is more flexible and can learn non–linear deformations and cope better than its PCA counterpart on noisy data. The PCA model, due to its rigid bases and the absence of a noise term is restricted in its ability to model non–linear shape deformations.

7 Implementation of the Fully Automatic Guidewire Navigation Tool

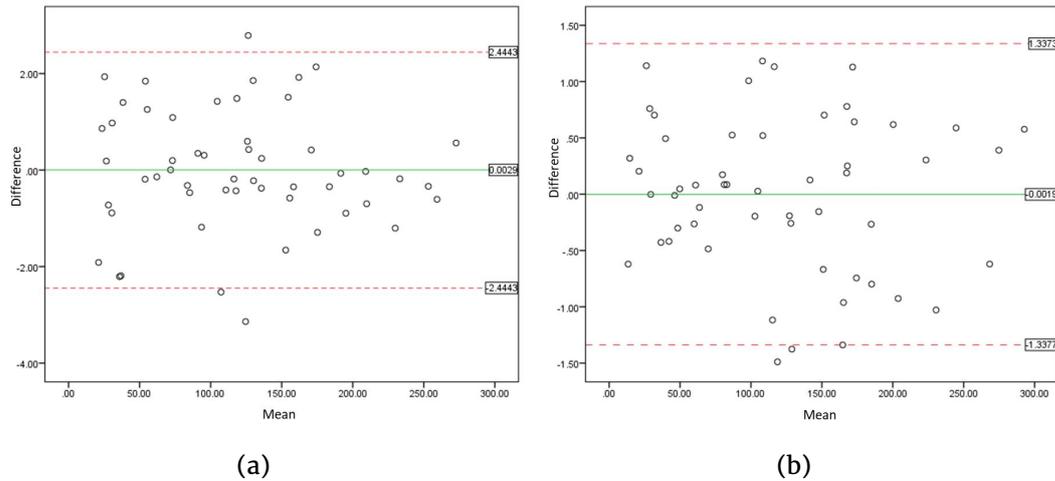


Figure 7.10: Bland-Altman Plots showing PPCA Estimate versus Ground truth Annotations

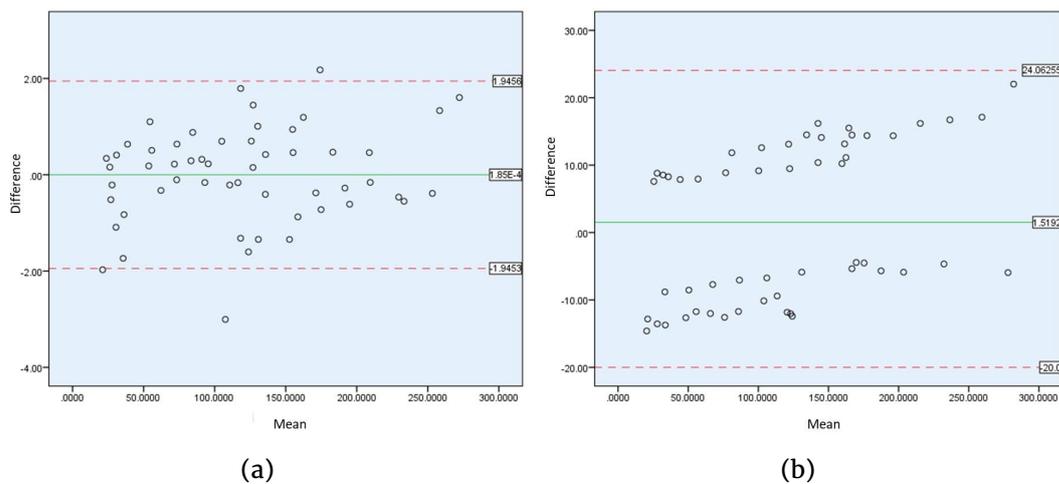


Figure 7.11: Bland-Altman Plots showing PCA Estimate versus Ground truth Annotations

A grey area in the implementation of this system is the uncertainty of the location of the line spanning the femoral head used to calculate P_r a potential improvement in this regard would be the provision of expert annotated images as this will further heighten clinical relevance.

8 Conclusion and Further Work

The main objective of the project was the development of a fully automatic guidewire insertion planning system.

There are already some existing computer-assisted interventions used for intraoperative surgery, but intertrochanteric fracture surgery management is yet to benefit from such interventions and the most pivotal part of the surgery which is the proper placement of the femoral head is still carried out manually.

There have been metrics in literature namely the Tip-Apex Distance and Parker's Ratio that have been known to predict the probability of cutout.

The investigations carried out in this work established a potential means for estimating an optimal position of the surgical implant through image analysis of radiographs. While the results look promising, the radiographs used in the analysis were not intraoperative radiographs (see Chapter 6 for samples) but those commonly used in the preoperative stage for diagnosis or surgical planning and sometimes used in the postoperative phase for monitoring surgical outcomes.

However, we argue that the methods developed can still be applied on intraoperative radiographs with little adjustment to the methods and techniques presented in this project. Getting access to such radiographs would require the support of a medical researcher or a clinician.

8 Conclusion and Further Work

An aspect of this research that is still open for further investigation is clinical testing of the developed system as this will ultimately determine how successful or beneficial the system is. To trial such systems in the first instance, a mock surgery session is staged where sawbone prototypes of the femur or cadaver are used in testing. A bottleneck in the research carried out is in the morphometric analysis of the femur as manual annotations of the femur were carried out by the researcher. It is expected that the annotations would be more consistent and accurate were the annotation task to be performed by an expert. An alternative option is to perform multiple instances of the annotation by different individuals based on a gold standard specified by the expert; after which inter-observer and intra-observer variability analysis are used to evaluate the quality of manual annotations.

An area for further research that might be worth investigating further is the minimisation of the manual annotation process or the elimination entirely particularly with regards to deformable object modelling or parts-based object analysis. Incorporating other imaging modalities such as MRI and CT scans might be useful in this aspect as such images are usually less noisy than x-rays.

This project has the potential to transform the orthopaedic surgery workflow by minimising surgery time, undue x-ray exposure and placing both experienced and junior surgeons on a 'level playing field'; whereby the surgical outcome is not necessarily dependent on the experience of the surgeon. The system developed could also be a potential training and practising tool for practitioners to hone their skills further.

References

- [1] O. Johnell and J. Kanis, 'An estimate of the worldwide prevalence, mortality and disability associated with hip fracture', *Osteoporosis International*, vol. 15, no. 11, pp. 897–902, 2004.
- [2] R. Smektala, C. Ohmann, S. Paech, E. Neuhaus, M. Rieger, W. Schwabe, P. Debold, A. Deimling, M. Jonas, K. Hupe *et al.*, 'On the prognosis of hip fractures. assessment of mortality after hip fractures by analyzing overlapping segments of longitudinal data', *Der Unfallchirurg*, vol. 108, no. 11, pp. 927–8, 2005.
- [3] J. Luthi, B. Burnand, W. McClellan, S. Pitts and W. Flanders, 'Is readmission to hospital an indicator of poor process of care for patients with heart failure?', *BMJ Quality & Safety*, vol. 13, no. 1, pp. 46–51, 2004.
- [4] M. Roland, M. Dusheiko, H. Gravelle and S. Parker, 'Follow up of people aged 65 and over with a history of emergency admissions: Analysis of routine admission data', *Bmj*, vol. 330, no. 7486, pp. 289–292, 2005.
- [5] E. Demir and T. Chausalet, 'A systematic approach in defining readmission', in *Computer-Based Medical Systems, 2009. CBMS 2009. 22nd IEEE International Symposium on*, IEEE, 2009, pp. 1–7.
- [6] A. Lešić, M. Jarebinski, T. Pekmezović, M. Bumbaširević, D. Spasovski and H. D. Atkinson, 'Epidemiology of hip fractures in belgrade, serbia montenegro, 1990–2000', *Archives of orthopaedic and trauma surgery*, vol. 127, no. 3, pp. 179–183, 2007.

References

- [7] M. Szpalski and R. Gunzburg, 'Prevention of hip lag screw cut-out in osteoporotic patients', *Bulletin Hospital for Joint Diseases Volume*, vol. 60, no. 2, pp. 2001–2002, 2001.
- [8] A. Lenich, E. Mayr and A. Rüter, 'Hip replacement after failed internal fixation in patients with proximal femur fracture—a simple procedure?', *Zentralblatt fur Chirurgie*, vol. 127, no. 6, pp. 503–506, 2002.
- [9] W. Werner-Tutschku, G. Lajtai, G. Schmiedhuber, T. Lang, C. Pirkl and E. Orthner, 'Intra- and perioperative complications in the stabilization of per- and subtrochanteric femoral fractures by means of pfn', *Der Unfallchirurg*, vol. 105, no. 10, pp. 881–885, 2002.
- [10] H. Handoll, C. Sherrington and M. Parker, 'Mobilisation strategies after hip fracture surgery in adults', *Cochrane Database Syst Rev*, vol. 1, 2007.
- [11] R. Simmermacher, J. Ljungqvist, H. Bail, T. Hockertz, A. Vochteloo, U. Ochs, C. vd Werken *et al.*, 'The new proximal femoral nail anti-rotation (pfna®) in daily practice: Results of a multicentre clinical study', *Injury*, vol. 39, no. 8, pp. 932–939, 2008.
- [12] J. Zou, Y. Xu and H. Yang, 'A Comparison of Proximal Femoral Nail Antirotation and Dynamic Hip Screw Devices in Trochanteric Fractures', *Journal of International Medical Research*, vol. 37, no. 4, pp. 1057–1064, Aug. 2009.
- [13] K. De Bruijn, D. den Hartog, W. Tuinebreijer and G. Roukema, 'Reliability of predictors for screw cutout in intertrochanteric hip fractures', *The Journal of Bone and Joint Surgery. American Volume*, vol. 94, no. 14, pp. 1266–1272, Jul. 2012.
- [14] A. S. Kumar, V. Parmar, J. Bankart, S. Williams and W. Harper, 'Comparison of accuracy of lag screw placement in cephalocondylic nails and sliding hip screw plate fixation for extracapsular fractures of the neck of femur', *International orthopaedics*, vol. 30, no. 5, pp. 320–324, 2006.

- [15] A. S. Kumar, V. Parmar, S. Kolpattil, S. Humad, S. Williams and W. Harper, 'Significance of hip rotation on measurement of tip apex distance during fixation of extracapsular proximal femoral fractures', *Injury*, vol. 38, no. 7, pp. 792–796, 2007.
- [16] M. B. Sommers, C. Roth, H. Hall, B. C. Kam, L. W. Ehmke, J. C. Krieg, S. M. Madey and M. Bottlang, 'A laboratory model to evaluate cutout resistance of implants for pertrochanteric fracture fixation', *Journal of orthopaedic trauma*, vol. 18, no. 6, pp. 361–368, 2004.
- [17] A. Lobo-Escolar, E. Joven, D. Iglesias and A. Herrera, 'Predictive factors for cutting-out in femoral intramedullary nailing', *Injury*, vol. 41, no. 12, pp. 1312–1316, 2010.
- [18] A. Lenich, H. Vester, M. Nerlich, E. Mayr, U. Stöckle and B. Füchtmeier, 'Clinical comparison of the second and third generation of intramedullary devices for trochanteric fractures of the hip-blade vs screw', *Injury*, vol. 41, no. 12, pp. 1292–1296, 2010.
- [19] G. H. Van Lenthe and R. Müller, 'Ct-based visualization and quantification of bone microstructure in vivo', *IBMS BoneKEy*, vol. 5, no. 11, pp. 410–425, 2008.
- [20] J. Vander Sloten and G. van Lenthe, 'Towards improved assessment of bone fracture risk', in *Advanced Technologies for Enhanced Quality of Life, 2009. AT-EQUAL'09.*, IEEE, 2009, pp. 21–21.
- [21] P. Prendergast, 'Finite element models in tissue mechanics and orthopaedic implant design', *Clinical Biomechanics*, vol. 12, no. 6, pp. 343–366, 1997.
- [22] Y. Bourgeois, Y. Petit and Y. G-Laflamme, 'Finite element model of a greater trochanteric reattachment system', in *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE*, IEEE, 2010, pp. 3926–3929.
- [23] M. Baumgaertner R, S. Curtin, D. Lindskog and J. Keggi, 'The value of the tip-apex distance in predicting failure of fixation of per-

References

- itrochanteric fractures of the hip', English, *Journal of Bone & Joint Surgery-American Volume*, vol. 77, no. 7, pp. 1058–64, Jul. 1995.
- [24] M. R. Baumgaertner and B. D. Solberg, 'Awareness of Tip–Apex Distance Reduces Failure of Fixation of Trochanteric Fractures of the Hip', en, *J Bone Joint Surg Br*, vol. 79–B, no. 6, pp. 969–971, Nov. 1997.
- [25] J. A. Geller, C. Saifi, T. A. Morrison and W. Macaulay, 'Tip–apex distance of intramedullary devices as a predictor of cut–out failure in the treatment of peritrochanteric elderly hip fractures', en, *International Orthopaedics*, vol. 34, no. 5, pp. 719–722, Jun. 2010.
- [26] H. Pervez, M. J. Parker and S. Vowler, 'Prediction of fixation failure after sliding hip screw fixation', *Injury*, vol. 35, no. 10, pp. 994–998, 2004.
- [27] M. Güven, U. Yavuz, B. Kadioğlu, B. Akman, V. Kılınçoğlu, K. Ünay and F. Altıntaş, 'Importance of screw position in intertrochanteric femoral fractures treated by dynamic hip screw', *Orthopaedics & Traumatology: Surgery & Research*, vol. 96, no. 1, pp. 21–27, 2010.
- [28] R. Fisker, *Making deformable template models operational*. Institute of Mathematical Modelling, Technical University of Denmark, 2000.
- [29] A. K. Jain, Y. Zhong and M.–P. Dubuisson–Jolly, 'Deformable template models: A review', *Signal processing*, vol. 71, no. 2, pp. 109–129, 1998.
- [30] T. Cootes, D. Cristinacce and V. Petrovic, 'Statistical Models of Shape and Texture for Face Recognition', *Toward Category–Level Object Recognition*, pp. 525–542, 2006.
- [31] P. Martins, J. F. Henriques, R. Caseiro and J. Batista, 'Bayesian constrained local models revisited', *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 4, pp. 704–716, 2016.
- [32] G. S. Muralidhar, A. C. Bovik, J. D. Giese, M. P. Sampat, G. J. Whitman, T. M. Haygood, T. W. Stephens and M. K. Markey, 'Snakules: A Model–Based Active Contour Algorithm for the Annotation

- of Spicules on Mammography’, *IEEE Transactions on Medical Imaging*, vol. 29, no. 10, pp. 1768–1780, Oct. 2010.
- [33] T. F. Cootes, M. C. Ionita, C. Lindner and P. Sauer, ‘Robust and accurate shape model fitting using random forest regression voting’, in *European Conference on Computer Vision*, Springer, 2012, pp. 278–291.
- [34] T. Mutsvangwa, V. Burdin, C. Schwartz and C. Roux, ‘An automated statistical shape model developmental pipeline: Application to the human scapula and humerus’, *IEEE Transactions on Biomedical Engineering*, vol. 62, no. 4, pp. 1098–1107, 2015.
- [35] M. Kass, A. Witkin and D. Terzopoulos, ‘Snakes: Active contour models’, *International journal of computer vision*, vol. 1, no. 4, pp. 321–331, 1988.
- [36] T. F. Cootes, C. J. Taylor, D. H. Cooper and J. Graham, ‘Training models of shape from sets of examples’, in *BMVC92*, Springer, 1992, pp. 9–18.
- [37] T. Albrecht, M. Lüthi and T. Vetter, ‘Deformable models’, *Encyclopedia of Biometrics*, pp. 337–343, 2015.
- [38] T. McInerney and D. Terzopoulos, ‘Deformable models in medical image analysis: A survey’, *Medical image analysis*, vol. 1, no. 2, pp. 91–108, 1996.
- [39] T. F. Cootes and C. J. Taylor, ‘Active shape models—’smart snakes’.’, in *BMVC*, vol. 92, 1992, pp. 266–275.
- [40] G. J. Edwards, T. F. Cootes and C. J. Taylor, ‘Face recognition using active appearance models’, in *European conference on computer vision*, Springer, 1998, pp. 581–595.
- [41] — —, ‘Advances in active appearance models’, in *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on*, IEEE, vol. 1, 1999, pp. 137–142.

References

- [42] T. F. Cootes, G. J. Edwards and C. J. Taylor, 'Active appearance models', *IEEE Transactions on pattern analysis and machine intelligence*, vol. 23, no. 6, pp. 681–685, 2001.
- [43] D. Cristinacce and T. F. Cootes, 'Feature detection and tracking with constrained local models.', in *BMVC*, vol. 1, 2006, p. 3.
- [44] D. Cristinacce and T. Cootes, 'Automatic feature localisation with constrained local models', *Pattern Recognition*, vol. 41, no. 10, pp. 3054–3067, 2008.
- [45] Y. Wang, S. Lucey and J. F. Cohn, 'Enforcing convexity for improved alignment with constrained local models', in *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, IEEE, 2008, pp. 1–8.
- [46] C. Lindner, S. Thiagarajah, J. Wilkinson, T. Consortium, G. Wallis and T. Cootes, 'Fully automatic segmentation of the proximal femur using random forest regression voting', *IEEE transactions on medical imaging*, vol. 32, no. 8, pp. 1462–1472, 2013.
- [47] A. Asthana, S. Zafeiriou, S. Cheng and M. Pantic, 'Robust discriminative response map fitting with constrained local models', in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 3444–3451.
- [48] T. Baltrusaitis, P. Robinson and L.-P. Morency, 'Constrained local neural fields for robust facial landmark detection in the wild', in *Computer Vision Workshops (ICCVW), 2013 IEEE International Conference on*, IEEE, 2013, pp. 354–361.
- [49] A. Zadeh, T. Baltrusaitis and L.-P. Morency, 'Convolutional experts network for facial landmark detection', in *Proceedings of the International Conference on Computer Vision & Pattern Recognition (CVPRW), Faces-in-the-wild Workshop/Challenge*, vol. 3, 2017, p. 6.
- [50] S. Lucey, Y. Wang, M. Cox, S. Sridharan and J. F. Cohn, 'Efficient constrained local model fitting for non-rigid face alignment', *Image and vision computing*, vol. 27, no. 12, pp. 1804–1813, 2009.

- [51] S. Lucey, Y. Wang, J. Saragih and J. F. Cohn, ‘Non-rigid face tracking with enforced convexity and local appearance consistency constraint’, *Image and vision computing*, vol. 28, no. 5, pp. 781–789, 2010.
- [52] J. M. Saragih, S. Lucey and J. F. Cohn, ‘Face alignment through subspace constrained mean-shifts’, in *Computer Vision, 2009 IEEE 12th International Conference on, Ieee*, 2009, pp. 1034–1041.
- [53] — —, ‘Deformable model fitting by regularized landmark mean-shift’, *International Journal of Computer Vision*, vol. 91, no. 2, pp. 200–215, 2011.
- [54] P. Martins, R. Caseiro, J. F. Henriques and J. Batista, ‘Likelihood-enhanced Bayesian constrained local models’, in *Image Processing (ICIP), 2014 IEEE International Conference on, IEEE*, 2014, pp. 303–307.
- [55] C. Lindner, ‘Statistical shape analysis of the proximal femur: Development of a fully automatic segmentation system and its applications’, PhD thesis, University of Manchester, 2014.
- [56] C. Lindner, P. A. Bromiley, M. C. Ionita and T. F. Cootes, ‘Robust and accurate shape model matching using random forest regression-voting’, *IEEE transactions on pattern analysis and machine intelligence*, vol. 37, no. 9, pp. 1862–1874, 2015.
- [57] P. A. Bromiley, J. E. Adams and T. F. Cootes, ‘Localisation of vertebrae on DXA images using constrained local models with random forest regression voting’, in *Recent Advances in Computational Methods and Clinical Applications for Spine Imaging*, Springer, 2015, pp. 159–171.
- [58] M. E. Tipping and C. M. Bishop, ‘Mixtures of probabilistic principal component analyzers’, *Neural computation*, vol. 11, no. 2, pp. 443–482, 1999.
- [59] L. Torresani, A. Hertzmann and C. Bregler, ‘Nonrigid structure-from-motion: Estimating shape and motion with hierarchical pri-

References

- ors', *IEEE transactions on pattern analysis and machine intelligence*, vol. 30, no. 5, pp. 878–892, 2008.
- [60] D. S. Bolme, B. A. Draper and J. R. Beveridge, 'Average of synthetic exact filters', in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, IEEE, 2009, pp. 2105–2112.
- [61] D. S. Bolme, J. R. Beveridge, B. A. Draper and Y. M. Lui, 'Visual object tracking using adaptive correlation filters', in *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, IEEE, 2010, pp. 2544–2550.
- [62] I. L. Dryden and K. V. Mardia, *Statistical analysis of shape*. Wiley, 1998.
- [63] D. G. Kendall, 'A survey of the statistical theory of shape', *Statistical Science*, pp. 87–99, 1989.
- [64] T. F. Cootes, A. Hill, C. J. Taylor and J. Haslam, 'The use of active shape models for locating structures in medical images', in *Biennial International Conference on Information Processing in Medical Imaging*, Springer, 1993, pp. 33–47.
- [65] T. F. Cootes, A. Hill, C. J. Taylor and J. Haslam, 'Use of active shape models for locating structures in medical images', *Image and vision computing*, vol. 12, no. 6, pp. 355–365, 1994.
- [66] T. F. Cootes, C. J. Taylor, D. H. Cooper and J. Graham, 'Active shape models—their training and application', *Computer vision and image understanding*, vol. 61, no. 1, pp. 38–59, 1995.
- [67] T. F. Cootes, C. J. Taylor *et al.*, *Statistical models of appearance for computer vision*. Technical report, University of Manchester, 2004.
- [68] D. Cristinacce and T. F. Cootes, 'A comparison of shape constrained facial feature detectors', in *Automatic Face and Gesture Recognition, 2004. Proceedings. Sixth IEEE International Conference on*, IEEE, 2004, pp. 375–380.
- [69] D. Barber, *Bayesian reasoning and machine learning*, eng. Cambridge: Cambridge University Press, 2012.

- [70] K. P. Murphy, *Machine learning: a probabilistic perspective*, eng, ser. Adaptive computation and machine learning. Cambridge, Mass. ; London: MIT Press, 2012.
- [71] W. Huizinga, D. H. J. Poot, J. .-.M. Guyader, R. Klaassen, B. F. Coolen, M. van Kranenburg, R. J. M. van Geuns, A. Uitterdijk, M. Polfliet, J. Vandemeulebroucke, A. Leemans, W. J. Niessen and S. Klein, 'PCA-based groupwise image registration for quantitative MRI', *Medical Image Analysis*, vol. 29, pp. 65–78, Apr. 2016.
- [72] J. E. Iglesias and N. Karssemeijer, 'Robust Initial Detection of Landmarks in Film-Screen Mammograms Using Multiple FFDM Atlases', *IEEE Transactions on Medical Imaging*, vol. 28, no. 11, pp. 1815–1824, Nov. 2009.
- [73] G. Quellec, S. R. Russell and M. D. Abramoff, 'Optimal filter framework for automated, instantaneous detection of lesions in retinal images', *IEEE transactions on medical imaging*, vol. 30, no. 2, pp. 523–533, 2011.
- [74] A. L. Samuel, 'Some studies in machine learning using the game of checkers', *IBM Journal of research and development*, vol. 44, no. 1.2, pp. 206–226, 2000.
- [75] G. Isabelle, 'Feature extraction foundations and applications. pattern recognition', 2006.
- [76] T. M. Mitchell, 'Machine learning. 1997', *Burr Ridge, IL: McGraw Hill*, vol. 45, no. 37, pp. 870–877, 1997.
- [77] T. M. Mitchell, 'Does machine learning really work?', *AI magazine*, vol. 18, no. 3, p. 11, 1997.
- [78] T. M. Mitchell, *The discipline of machine learning*. Carnegie Mellon University, School of Computer Science, Machine Learning Department, 2006, vol. 3.
- [79] C. Bishop, 'Pattern recognition and machine learning (information science and statistics), 1st edn. 2006. corr. 2nd printing edn', *Springer, New York*, 2007.

References

- [80] Y. Anzai, *Pattern Recognition and Machine Learning*, English, 1 edition. Morgan Kaufmann, Dec. 2012.
- [81] M. I. Jordan and T. M. Mitchell, 'Machine learning: Trends, perspectives, and prospects', *Science*, vol. 349, no. 6245, pp. 255–260, 2015.
- [82] S. Theodoridis, *Machine learning: a Bayesian and optimization perspective*. Academic Press, 2015.
- [83] S. Belongie, J. Malik and J. Puzicha, 'Shape context: A new descriptor for shape matching and object recognition', in *Advances in neural information processing systems*, 2001, pp. 831–837.
- [84] M. Yang, K. Kpalma and J. Ronsin, 'A survey of shape feature extraction techniques', in. IN-TECH, 2008.
- [85] Y. Zhao and S. Belkasim, 'Multiresolution fourier descriptors for multiresolution shape analysis', *IEEE Signal Processing Letters*, vol. 19, no. 10, pp. 692–695, 2012.
- [86] N. Sebe and M. S. Lew, 'Texture features for content-based retrieval', in *Principles of visual information retrieval*, Springer, 2001, pp. 51–85.
- [87] M. S. Nixon and A. S. Aguado, *Feature extraction & image processing for computer vision*. Academic Press, 2012.
- [88] O. Van Kaick, H. Zhang, G. Hamarneh and D. Cohen-Or, 'A survey on shape correspondence', in *Computer Graphics Forum*, Wiley Online Library, vol. 30, 2011, pp. 1681–1707.
- [89] T. F. Cootes, C. J. Taylor *et al.*, 'Statistical models of appearance for medical image analysis and computer vision', in *Proc. SPIE medical imaging*, vol. 4322, 2001, p. 5.
- [90] S. C. Mitchell, B. P. Lelieveldt, R. J. van der Geest, J. Schaap, J. H. Reiber and M. Sonka, 'Segmentation of cardiac mr images: An active appearance model approach', in *Medical Imaging 2000*, International Society for Optics and Photonics, 2000, pp. 224–234.

- [91] M. Poon, G. Hamarneh and R. Abugharbieh, 'Efficient interactive 3d livewire segmentation of complex objects with arbitrary topology', *Computerized Medical Imaging and Graphics*, vol. 32, no. 8, pp. 639–650, 2008.
- [92] Y.-C. J. Hu, M. D. Grossberg and G. S. Mageras, 'Semi-automatic medical image segmentation with adaptive local statistics in conditional random fields framework', in *Engineering in Medicine and Biology Society, 2008. EMBS 2008. 30th Annual International Conference of the IEEE*, IEEE, 2008, pp. 3099–3102.
- [93] A. Saad, G. Hamarneh and T. Moller, 'Exploration and visualization of segmentation uncertainty using shape and appearance prior information', *IEEE Transactions on Visualization and Computer Graphics*, vol. 16, no. 6, pp. 1366–1375, 2010.
- [94] Y.-F. Chen, P.-C. Huang, K.-C. Lin, H.-H. Lin, L.-E. Wang, C.-C. Cheng, T.-P. Chen, Y.-K. Chan and J. Y. Chiang, 'Semi-automatic segmentation and classification of pap smear cells', *IEEE Journal of Biomedical and Health Informatics*, vol. 18, no. 1, pp. 94–108, 2014.
- [95] P. F. Felzenszwalb, R. B. Girshick, D. McAllester and D. Ramanan, 'Object Detection with Discriminatively Trained Part-Based Models', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 9, pp. 1627–1645, Sep. 2010.
- [96] T. Malisiewicz, A. Gupta and A. A. Efros, 'Ensemble of exemplar-SVMs for object detection and beyond', in *2011 International Conference on Computer Vision*, Nov. 2011, pp. 89–96.
- [97] D. Oliver, R. Griffiths, J. Roche and O. Sahota, 'Hip fracture', *BMJ clinical evidence*, vol. 2010, 2010.
- [98] J. Manninger and G. Kazár, 'Pathology of femoral neck fractures', in *Internal fixation of femoral neck fractures: An atlas*, J. Manninger, U. Bosch, P. Cserháti, K. Fekete and G. Kazár, Eds. Vienna: Springer Vienna, 2007, pp. 29–51.

References

- [99] C. Bulstrode, J. Wilson–MacDonald, D. M. Eastwood and J. Fairbank, *Oxford textbook of trauma and orthopaedics*. Oxford University Press, 2011.
- [100] P. Adam, ‘Treatment of recent trochanteric fracture in adults’, *Orthopaedics & Traumatology: Surgery & Research*, vol. 100, no. 1, S75–S83, 2014.
- [101] B. Gullberg, O. Johnell and J. Kanis, ‘World–wide projections for hip fracture’, *Osteoporosis international*, vol. 7, no. 5, pp. 407–413, 1997.
- [102] R. Osterkamp, ‘Population developments in germany until 2050’, *Der Chirurg; Zeitschrift für alle Gebiete der operativen Medizin*, vol. 76, no. 1, pp. 10–18, 2005.
- [103] M. Parker and A. Johansen, ‘Hip fracture’, *BMJ: British Medical Journal*, vol. 333, no. 7557, p. 27, 2006.
- [104] T. Goodman and C. Spry, *Essentials of perioperative nursing*. Jones & Bartlett Publishers, 2016.
- [105] J. Ahn and J. Bernstein, ‘In brief: Fractures in brief: Intertrochanteric hip fractures’, *Clinical Orthopaedics and Related Research®*, vol. 468, no. 5, pp. 1450–1452, 2010.
- [106] H. Andruszkow, M. Frink, C. Fromke, A. Matityahu, C. Zeckey, P. Mommsen, S. Suntardjo, C. Krettek and F. Hildebrand, ‘Tip apex distance, hip screw placement, and neck shaft angle as potential risk factors for cut–out failure of hip screws after surgical treatment of intertrochanteric fractures’, *International Orthopaedics*, vol. 36, no. 11, pp. 2347–2354, Nov. 2012.
- [107] V. Parmar and A. S. Kumar, ‘The importance of surgical education in the accuracy of implant placement during hip fracture fixation’, *Journal of Orthopaedics and Traumatology*, vol. 10, no. 2, pp. 59–61, 2009.
- [108] J. Rubio–Avila, K. Madden, N. Simunovic and M. Bhandari, ‘Tip to apex distance in femoral intertrochanteric fractures: A systematic

- review', *Journal of Orthopaedic Science*, vol. 18, no. 4, pp. 592–598, 2013.
- [109] M. J. Parker, 'Cutting-out of the dynamic hip screw related to its position', *Bone & Joint Journal*, vol. 74, no. 4, pp. 625–625, 1992.
- [110] L. J. Johnson, M. R. Cope, S. Shahrokhi and P. Tamblyn, 'Measuring tip-apex distance using a picture archiving and communication system (PACS)', *Injury*, vol. 39, no. 7, pp. 786–790, 2008.
- [111] C. C. Mainds and R. J. Newman, 'Implant failures in patients with proximal fractures of the femur treated with a sliding screw device', *Injury*, vol. 20, no. 2, pp. 98–100, 1989.
- [112] T. R. Davis, J. Sher, A. Horsman, M. Simpson, B. Porter and R. Checketts, 'Intertrochanteric femoral fractures. mechanical failure after internal fixation', *Bone & Joint Journal*, vol. 72, no. 1, pp. 26–31, 1990.
- [113] J.-H. Chung, S.-Y. Ko, D.-S. Kwon, J.-J. Lee, Y.-S. Yoon and C.-H. Won, 'Robot-assisted femoral stem implantation using an intramedulla gauge', *IEEE Transactions on Robotics and Automation*, vol. 19, no. 5, pp. 885–892, 2003.
- [114] T. Mendel, M. Hänni, B. Gueorguiev, D. Wohlrab and G. O. Hofmann, 'The virtual isocentric aiming device: A new mechanical targeting concept', *Archives of orthopaedic and trauma surgery*, vol. 131, no. 12, pp. 1655–1662, 2011.
- [115] E. J. Hazan, 'Computer-assisted orthopaedic surgery: A new paradigm', *Techniques in Orthopaedics*, vol. 18, no. 2, pp. 221–229, 2003.
- [116] C. T. Mehlman and T. G. DiPasquale, 'Radiation exposure to the orthopaedic surgical team during fluoroscopy: "how far away is far enough?"', *Journal of orthopaedic trauma*, vol. 11, no. 6, pp. 392–398, 1997.
- [117] J. A. Alonso, A. Maxwell, D. Shaw and G. Hart, 'Scattered radiation during fixation of hip fractures—is distance alone enough protection', in *Engineering in Medicine and Biology Society, 2000. Proceedings*

References

- of the 22nd Annual International Conference of the IEEE, IEEE, vol. 2, 2000, pp. 995–997.
- [118] G. Singer, ‘Occupational radiation exposure to the surgeon’, *Journal of the American Academy of Orthopaedic Surgeons*, vol. 13, no. 1, pp. 69–76, 2005.
- [119] D. M. Kahler, ‘Image guidance: Fluoroscopic navigation.’, *Clinical orthopaedics and related research*, vol. 421, pp. 70–76, 2004.
- [120] K. T. Foley, D. A. Simon and Y. R. Rampersaud, ‘Virtual fluoroscopy: Computer-assisted fluoroscopic navigation’, *Spine*, vol. 26, no. 4, pp. 347–351, 2001.
- [121] H. H. Handoll and M. J. Parker, ‘Conservative versus operative treatment for hip fractures in adults’, *The Cochrane Library*, 2008.
- [122] M. Müller, P. Belei, M. de la Fuente, M. Strake, O. Weber, C. Burger, K. Radermacher and D. Wirtz, ‘Evaluation of a 2d fluoroscopy-based navigation system for insertion of the dynamic hip screw (dhs): An experimental study’, *RoFo: Fortschritte auf dem Gebiete der Rontgenstrahlen und der Nuklearmedizin*, vol. 183, no. 6, pp. 536–542, 2011.
- [123] D. Mayman, E. Vasarhelyi, W. Long, R. Ellis, J. Rudan and D. Pichora, ‘Computer-assisted guidewire insertion for hip fracture fixation’, *Journal of orthopaedic trauma*, vol. 19, no. 9, pp. 610–615, 2005.
- [124] A. B. Okoli, J. B. Penny, W. L. Woo and S. S. Dlay, ‘Guidewire insertion planning for extracapsular hip fracture surgery’, in *Biomedical Circuits and Systems Conference (BioCAS), 2014 IEEE*, IEEE, 2014, pp. 105–108.
- [125] J. Dujardin and D. Slice, ‘Geometric morphometrics. contributions to medical entomology’, *Encyclopedia of Infectious Diseases. Modern Methodologies*, Wiley & Sons, pp. 435–447, 2007.
- [126] J.-P. Dujardin, ‘Morphometrics applied to medical entomology’, *Infection, Genetics and Evolution*, vol. 8, no. 6, pp. 875–890, 2008.

- [127] C. F. Hester and D. Casasent, 'Multivariant technique for multiclass pattern recognition', *Applied Optics*, vol. 19, no. 11, pp. 1758–1761, 1980.
- [128] P. Viola and M. Jones, 'Rapid object detection using a boosted cascade of simple features', in *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, vol. 1, 2001, I–511–I–518 vol.1.
- [129] C. Lindner, S. Thiagarajah, J. M. Wilkinson, G. A. Wallis, T. F. Cootes, arcOGEN Consortium *et al.*, 'Accurate bone segmentation in 2d radiographs using fully automatic shape model matching based on regression-voting', in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, 2013, pp. 181–189.
- [130] T. Ojala, M. Pietikäinen and D. Harwood, 'A comparative study of texture measures with classification based on featured distributions', *Pattern recognition*, vol. 29, no. 1, pp. 51–59, 1996.
- [131] N. Dalal and B. Triggs, 'Histograms of oriented gradients for human detection', in *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, IEEE, vol. 1, 2005, pp. 886–893.
- [132] O. Alexa, B. Veliceasa, B. Puha and D. Cimpoesu, 'Digital templating in surgical treatment of trochanteric fractures', in *E-Health and Bioengineering Conference (EHB), 2013*, IEEE, 2013, pp. 1–4.
- [133] A. Herman, A. Dekel, I. B. Botser and E. L. Steinberg, 'Computer-assisted surgery for dynamic hip screw, using Surgix, a novel intraoperative guiding system', in *The International Journal of Medical Robotics and Computer Assisted Surgery*, vol. 5, no. 1, pp. 45–50, Mar. 2009.
- [134] J. Sauro, 'Estimating productivity: Composite operators for key-stroke level modeling', in *International Conference on Human-Computer Interaction*, Springer, 2009, pp. 352–361.

References

- [135] T. O. Ozanian and R. Phillips, 'Image analysis for computer-assisted internal fixation of hip fractures', *Medical Image Analysis*, vol. 4, no. 2, pp. 137–159, 2000.
- [136] A. R. Hsu, C. E. Gross, S. Bhatia and B. R. Levine, 'Template-directed instrumentation in total knee arthroplasty: Cost savings analysis', *Orthopedics*, vol. 35, no. 11, e1596–e1600, 2012.
- [137] R. Greiner-Perth, Y. Allam, H. El-Saghir, F. Röhl, J. Franke and H. Böhm, 'Analysis of reoperations after surgical treatment of degenerative cervical spine disorders: A report on 900 cases', *Central European Neurosurgery*, vol. 70, no. 01, pp. 3–8, 2009.
- [138] R. H. Davies, C. J. Twining, T. F. Cootes, J. C. Waterton and C. J. Taylor, 'A minimum description length approach to statistical shape modeling', *IEEE Transactions on Medical Imaging*, vol. 21, no. 5, pp. 525–537, 2002.
- [139] R. Pilgram, C. Walch, M. Blauth, W. Jaschke, R. Schubert and V. Kuhn, 'Knowledge-based femur detection in conventional radiographs of the pelvis', *Computers in Biology and Medicine*, vol. 38, no. 5, pp. 535–544, May 2008.
- [140] M. Lüthi, T. Gerig, C. Jud and T. Vetter, 'Gaussian process morphable models', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [141] Q. Zhao, K. Okada, K. Rosenbaum, D. J. Zand, R. Sze, M. Summar and M. G. Linguraru, 'Hierarchical constrained local model using ica and its application to down syndrome detection', in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, 2013, pp. 222–229.
- [142] T. Albrecht, M. Luthi, T. Gerig and T. Vetter, 'Posterior shape models', *Medical image analysis*, vol. 17, no. 8, pp. 959–973, 2013.
- [143] I. Matthews and S. Baker, 'Active appearance models revisited', *International journal of computer vision*, vol. 60, no. 2, pp. 135–164, 2004.

- [144] R. Tokola and D. Bolme, 'Ensembles of correlation filters for object detection', in *Applications of Computer Vision (WACV), 2015 IEEE Winter Conference on*, IEEE, 2015, pp. 935–942.
- [145] H. Kiani Galoogahi, T. Sim and S. Lucey, 'Correlation filters with limited boundaries', in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 4630–4638.
- [146] — —, 'Multi-channel correlation filters', in *Proceedings of the IEEE International Conference on Computer Vision*, 2013, pp. 3072–3079.
- [147] P. R. Kuzyk, R. Zdero, S. Shah, M. Olsen, J. P. Waddell and E. H. Schemitsch, 'Femoral head lag screw position for cephalomedullary nails: A biomechanical analysis', *Journal of orthopaedic trauma*, vol. 26, no. 7, pp. 414–421, 2012.
- [148] S. Li, S.-M. Chang, Y.-M. Jin, Y.-Q. Zhang, W.-X. Niu, S.-C. Du, L.-Z. Zhang and H. Ma, 'A mathematical simulation of the tip-apex distance and the calcar-referenced tip-apex distance for intertrochanteric fractures reduced with lag screws', *Injury*, vol. 47, no. 6, pp. 1302–1308, 2016.
- [149] J. M. Goffin, P. J. Jenkins, R. Ramaesh, P. Pankaj and A. H. Simpson, 'What is the relevance of the tip-apex distance as a predictor of lag screw cut-out?', *PloS one*, vol. 8, no. 8, e71195, 2013.

Appendices

Appendix A

A.1 Tip-Apex Distance Illustrated

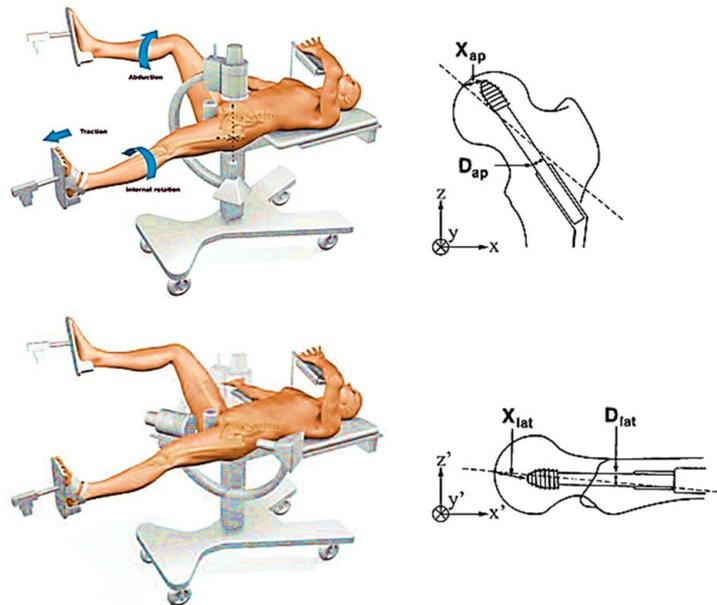


Figure A.1: This illustrates how TAD is computed taking into account the views and other parameters (*adapted from Goffin et al. [149]*).

$$TAD = \left(X_{ap} \times \frac{D_{ap}}{D_{true}} \right) + \left(X_{lat} \times \frac{D_{lat}}{D_{true}} \right) \quad (\text{A.1.1})$$

Where: D_{ap} = implant diameter in AP view. D_{lat} = implant diameter in lateral view. X_{ap} = distance between the tip of the implant and the apex in the AP view. X_{lat} = distance between the tip of the implant and the apex in the AP view. D_{true} = actual diameter of the implant.

A.2 Typical implants used in Extracapsular Hip Fracture Surgery

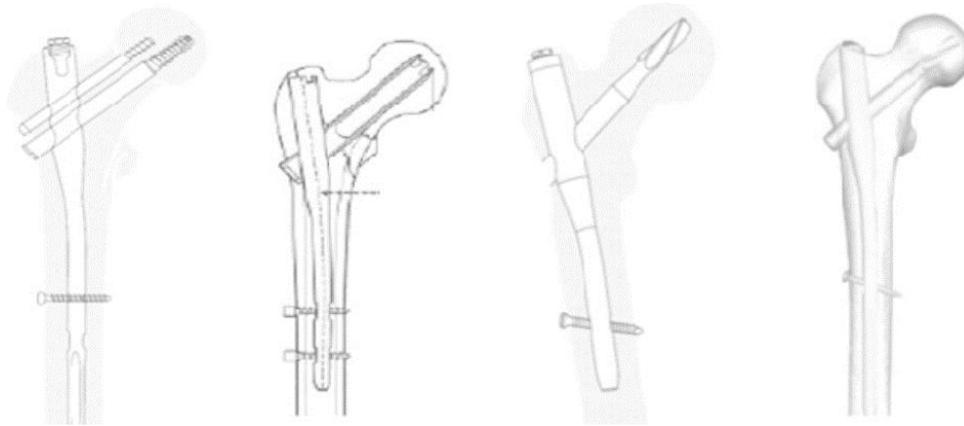


Figure A.2: Examples of implants used for intramedullary nailing.



(a) Well-aligned Implant



(b) Cutout Incident

Figure A.3: How an implant looks under an x-ray

A.3 Centre coordinates and radii plots for the Semi-automatic Tool

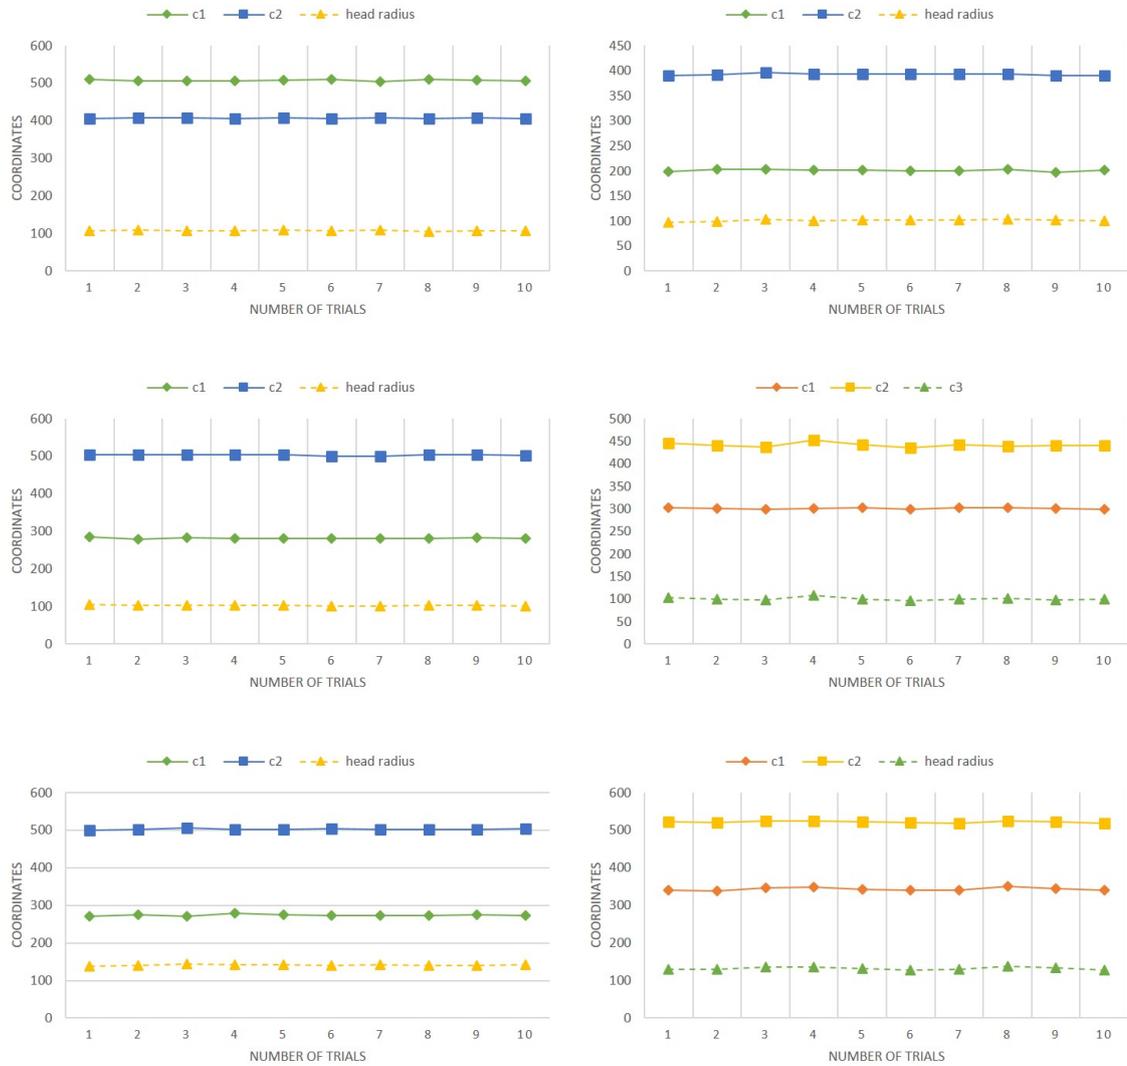


Figure A.4: Plots of the computed femoral head centre for 6 images shown along with the radius of the circumscribing circle.

A.4 Overview of Template-Based Object Matching Techniques

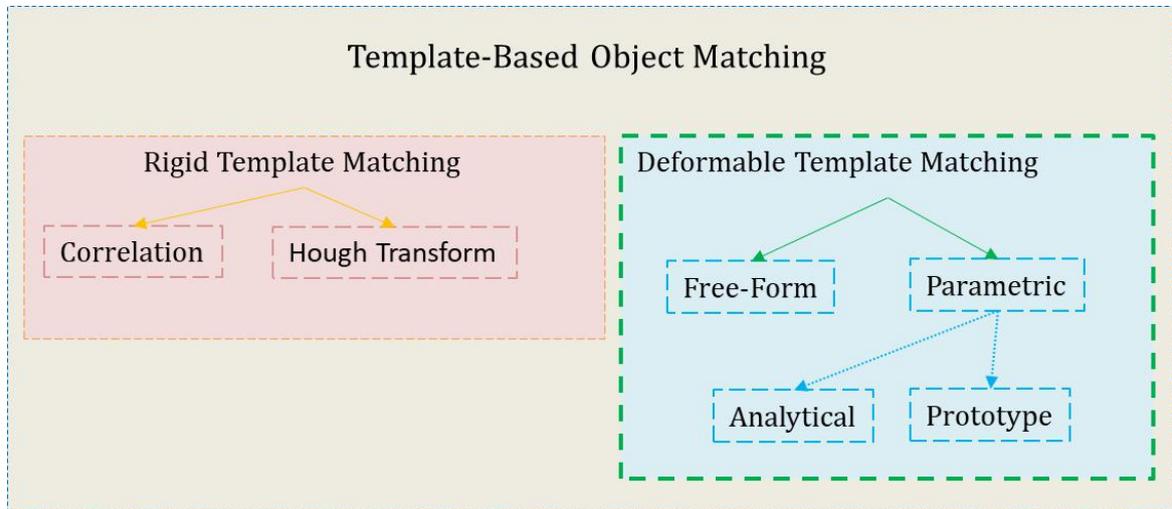


Figure A.5: An overview of Template-based Object-Matching Techniques
(*adapted from Jain et al. [29]*)