FACULTY OF HEALTH SCIENCES

DIVISION OF BIOMEDICAL ENGINEERING

Development of a statistical shape and appearance model of the skull from a South African population

Author:

BRIAN LUGADILU

(LGDBRI001)

Supervisor:

Dr Tinashe Mutsvangwa

Co-Supervisor:

Professor Tania Douglas

Submitted to the Department of Human Biology at the University of Cape Town in partial fulfilment of the academic requirements for the degree of M.Sc. Biomedical Engineering by coursework and dissertation

December 2017

DECLARATION

I, Brian Lugadilu hereby declare that the work in which this dissertation is based is my original work (except where acknowledgements indicate otherwise) and that neither a part nor the whole work has been or is to be submitted to any other college or university for assessment or for award of any other academic degree.

Signed: _ Signed by candidate

Date: __19/01/2018_____

ABSTRACT

Statistical shape models (SSMs) and statistical appearance models (SAMs) have been applied in medical analysis such as in surgical planning, finite element analysis, model-based segmentation, and in the fields of anthropometry and forensics. Similar applications can make use of SSMs and SAMs of the skull. A combination of the SSM and SAM of the skull can also be used in model-based segmentation. This document presents the development of a SSM and a SAM of the human skull from a South African population, using the Scalismo software package.

The SSM development pipeline was composed of three steps: 1) Image data segmentation and processing; 2) Development of a free-form deformation (FFD) model for establishing correspondence across the training dataset; and 3) Development and validation of a SSM from the corresponding dataset. The SSM was validated using the leave one-out cross-validation method. The first eight principal components of the SSM represented 92.13% of the variation in the model. The generality of the model in terms of the Hausdorff distance between a new shape generated by the SSM and instances of the SSM had a steady state value of 1.48mm. The specificity of the model (in terms of Hausdorff distance) had a steady state value of 2.04mm.

The SAM development pipeline involved four steps: 1) Volumetric mesh generation of the reference mesh to be used in establishing volumetric correspondence; 2) Sampling of intensity values from original computed tomography (CT) images using the in-correspondence volumetric meshes; and 3) Development of a SAM from the in-correspondence intensity values. A complete validation of the SAM was not possible due to limitations of the Scalismo software. As a result, only the shapes of the incomplete skulls were reconstructed and thereby validated. The amount of missing detail, as represented by absent landmarks, affected the registration results. Complete validation of the SAM is recommended as future work, via the use of a combined shape and intensity model (SSIM).

ACKNOWLEDGEMENT

I would like to express my sincere appreciation to the people who supported my research in different ways.

I am grateful to my supervisor, Dr T. Mutsvangwa whose expertise, guidance and support enabled me to work on a research topic that was of great interest to me. It was a pleasure working with him as he was always available when I needed assistance in my work.

I would also like to express my gratitude to Prof. T. Douglas for the research supervision, guidance, support and provision of resources for the project.

I would like to also thank Mr. C. Reyneke for his guidance and consultations throughout my research work.

I would also like to thank Mr. J.S Rassaire for the consultations we had together during my research.

Lastly, I would like to thank my family members: my mother Mrs. P. E. Kaziga, my late step-father Dr. D. E. Lumumba, and my brother Mr. K. Amadiva for their continuous support, motivation and prayers.

TABLE OF CONTENTS

LIST OF FIGURES

ABBREVIATIONS

2D	Two dimensional
3D	Three dimensional
AAM	Active appearance model
ASM	Active shape model
BMD	Bone mineral density
CT	Computer tomography
DefR	Deformed reference
FE	Finite element
FFD	Free-form deformation
FitPost	Fitted posterior
$FitPost_{Alg}$	Aligned fitted posterior
$FitPost_{REF1}$	Fitted posterior reference 1
$FitPost_{REF2}$	Fitted posterior reference 2
$FitPost_{AlgREF}$	Mean shape of fitted posteriors
GPA	General Procrustes alignment
HREF	Human Research Ethics Committee
ICP	Iterative closest point
KL	Karhunen-Loeve
LAT	Local affine transformation
LDD	Log-Domain demons
MDL	Minimum description length
MDT	Metal deletion technique
MRFSR	Markov random field surface reconstruction
MRI	Magnetic resonance imaging

PCA	Principal Component analysis
R	Reference
RKHS	Reproducing kernel Hilbert space
SAM	Statistical appearance model
SEK	Squared exponential kernel
SSM	Statistical shape model
STL	Stereolithography
SVD	Singular value decomposition
SVF	Stationary velocity fields
TPS	Thin plate splines
$Vol.Meshes_{Reg}$	Registered volumetric meshes

1 INTRODUCTION

1.1 MOTIVATION

Statistical shape models (SSMs) and statistical appearance models (SAMs) describe the mean shape and mean density, respectively, of anatomical structures such as bones within a population under study, based on information obtained from medical images (Sarkalkan et al., 2014). Additionally, they describe the main modes of variation of shape and density from the mean shape and mean density, respectively. They can therefore be used to describe shape and density variation within a population. This is possible because the shape and density distributions are obtained by adding sets of principal modes of variation onto the average shape and density, respectively. Statistical shape models can be developed from either two dimensional (2D) images such as planar X-ray images or three dimensional (3D) images from modalities such as computed tomography (CT) and magnetic resonance imaging (MRI) (Cootes et al., 1995). First, correspondence is established across the dataset. This often involves registration. The dataset in correspondence is then aligned to remove effects of translation and rotation across the dataset. Principal component analysis (PCA) is applied on the covariance matrix which represents the variation in shape across the dataset thereby calculating the average shape and mode of variation from the average. SAMs are generated by capturing pixel intensity information within the training dataset (Cootes & Taylor, 2001). However, effects of lighting that may be caused by different imaging modalities are first removed by normalizing the grey values of the pixel intensity. After the normalization, the average density distribution and the principal modes of variation from the average pixel value are calculated.

Three dimensional CT imaging has been used extensively for diagnostic purposes (Offut, 1990; Ono et al., 1992). However, before the CT images can be used for diagnosis, they first need to be processed for accurate delineation and biological characterisation of individual tissues e.g. bone, fat or muscle (Huang & Tsechpenakis, 2009). For delineation purposes, segmentation protocols are applied to the CT images. However, most segmentation techniques are based on intensity values of pixels in the CT images (Pal & Pal, 1993). Using such techniques, it is difficult to differentiate tissues that have intensity values that are similar, such as different bones (Pal & Pal, 1993). Fortunately, medical experts can still delineate these objects based on their prior knowledge of what the objects should look like. Typically, medical experts manually identify the objects of interest and delineate them. However, this is time consuming especially in countries where resources are limited such as in South Africa (Coovadia et al., 2009). The prior knowledge of medical experts during segmentation forms the basis for model-based segmentation protocols where the information of the object to be segmented is modelled by SSMs and SAMs and used in the segmentation (Hiemann & Hans-Peter,

2009). The model-based segmentation can be automatically applied to CT images, reducing the segmentation time.

Statistical shape models and SAMs of the skull have been applied in medical analysis such as in surgical planning (Gateno et al., 2007), FE analysis (Neuert, 2016), model-based segmentation (Sinha et al., 2016) and anthropometry and forensics (Paysan et al., 2009). In the South African context, similar applications can make use of SSMs and SAMs. According to a report by the Institute of Security Studies of South Africa, in 2014-2015, there were 33 murders for every 100,000 people in the country (Altbeker, 2008). Of the reported murders in the country, only a negligible percentage is solved annually (Salfati et al., 2015). Statistical shape models can aid in forensic investigation where there is need for reconstruction of shattered skulls for investigative purposes. Craniofacial deformities are abnormalities of the face and skull resulting from abnormal growth patterns of the face/skull (David et al., 2012). They may be caused by trauma to the head, diseases and birth defects such as fetal alcohol syndrome (FAS), which is highly prevalent in South Africa (Viljoen et al., 2005). Statistical models of the cranium may be of use in capturing the morphological variation for computer aided interventions for craniofacial deformities.

Research on skull SSM and SAM is limited, and has mostly been applied on Caucasian populations (Zachow et al., 2005). Research on SSM and SAM of the skull from a South African population may provide a starting point for comparative analysis between different South African sub-populations; and between South African and other population groups. It should be noted that the South African population is heterogeneous. The population is composed of sub-populations of Caucasian, African and mixed-race ancestry, distinguishing it from other populations across most Sub-Saharan countries. Thus, a SSM or SAM developed from the South African population would capture shape variations of the skull characteristic of both the Caucasian and African ancestry. Such models cannot therefore be used for comparative studies with other SSMs of the skull developed from a purely Caucasian or African ancestry. They would however be a starting point for sub-population-specific South African skull models.

1.2 PROBLEM STATEMENT

Research in the field of SSM and SAM of the skull based on the South African population can aid in: comparative analysis between the skull of the South African and other population groups; 3D reconstruction from 2D images of the skull; skull reconstruction; automated skull segmentation from images; development of patient-specific FE models of the skull; skull implant design and; surgical planning specific to the South African population. However, currently, there exists no SSM or SAM of the skull from the South African population.

1.3 PROJECT AIM AND OBJECTIVE

The aim of this project was to develop a pipeline for building SSM and SAM of the skull based on the South African population. The SSM and SAM are to be used for various medical applications. To achieve the stated aim, the following objectives had to be met:

i. Develop a segmentation protocol for the CT volumes for generation of 3D mesh surfaces required for the statistical models.
ii. Develop a protocol for the identification of anatomical landmark points on the 3D mesh surfaces. These landmarks were to be used in guiding the registration process (see next objective).
iii. Develop a protocol for establishing dense correspondence across the 3D mesh surfaces.
iv. Develop a SSM and a SAM of the skull from the corresponding 3D mesh surfaces.
v. Validate the quality of the SSM and SAM models.

1.4 PROJECT SCOPE AND LIMITATION

The primary focus of the project was to develop an algorithmic pipeline for developing the SSM and SAM of the skull from the South African population. As such, the data requirements to achieve the focus were modest. CT image volumes of 16 subjects were used to develop the pipeline. With this small training set, the resultant statistical models are expected to be limited in their generalization capacity. However, the algorithmic pipeline that has been established, may be applied to different, and larger, data sets in future.

1.5 ETHICAL CONSIDERATION

Since the project involved the use of human remains, ethical clearance was sought from the Human Research Ethics committee, Faculty of Health Sciences, University of Cape Town (HREF REF: 264/2017). Under the approval, only CT images obtained from deceased individuals, who had donated their bodies for research purpose to the University of Cape Town, would be used.

1.6 DISSERTATION OVERVIEW

The rest of this document is structured as follows: Chapter 2 presents a literature review focusing on: the anatomy of the skull; SSMs and SAMs in general; and the validation of SSMs and SAMs. Chapter 3 outlines the segmentation of the CT images and the pre-processing of the segmented surfaces. Chapter 4 describes the registration of the 3D surfaces to establish dense correspondence. The detailed outline of SSM development from registered surfaces (in-correspondence) is provided in Chapter 5. Chapter 6 describes the development of a SAM from the SSM developed in Chapter 5. Finally, Chapter 7 presents a discussion and conclusion.

2 LITERATURE REVIEW

2.1 THE SKULL

The skull is a skeletal structure that surrounds the brain for protection. In adults, the skull is composed of 22 individual bones. All these bones are fixed apart from the mandible which is movable. The individual bones have different shapes which results in the complex shape of the skull (Drake et al., 2014). Apart from the bones, the skull also has orbits, septa, conches, fossa, teeth and sinuses. In studying the skull for either diagnostic or therapeutic purposes, all these features have to be considered. Computed tomography (CT) is among the imaging modalities used for the detection of cranial and intracranial injuries (Masters et al., 1987).

2.2 STATISTICAL SHAPE MODELLING OF BONE

Statistical shape models describe the mean shape of anatomical structures and the patterns of variation from this average shape (Sarkalkan et al., 2014). Such models are useful for: generation of new shape instances (Cootes et al., 1995); reconstruction of anatomical parts (Baka et al., 2011); automatic segmentation of images of anatomical parts (Heimann & Meinzer, 2009); representation of the shape variation of the anatomical parts (Dryden & Mardia, 2016) and for 3D shape reconstruction from 2D shapes (Baka et al., 2011). Generation of new shape instances of anatomical parts can be used in surgical planning where anatomical reconstruction is required (Hochfeld et al., 2014). These instances can be manufactured as implants and fixed onto the missing or deformed anatomical part (Paysan et al., 2009). Statistical shape models for model-based segmentation are useful, as conventional bony-structure segmentation uses intensity values for segmentation and is not efficient in segmenting one bone from other bones because of similar intensity values (Loubele et al., 2005). Model-based segmentation has the advantage of containing prior information about the structure of the bone to assist segmentation of one bone from another (Sinha et al., 2016).

Statistical shape models of the human skull have been developed and reported in recent literature for a variety of medical image analysis applications. Kim et al. (2012) used the SSM of the mandible for treatment planning. Computed tomography images from 23 Caucasian males and 23 Caucasian females were used to develop a 3D mandible model which was used for treatment planning in orthognathic surgery. Zachow et al. (2005) used CT images from 11 subjects to create a SSM of the mandible which was used for surgical planning of mandible deformities. Gateno et al. (2007) used a SSM of the skull for computer aided surgical simulation of five patient cases with complex craniomaxillofacial surgery. Hochfeld et al. (2014) developed a SSM of the skull from a dataset of 21 MRI images of healthy infants. The SSM was later used for objective and patient-specific guidance in the reshaping of skull deformities in infants suffering from craniosynostosis.

Semper-Hogg et al. (2017) used a SSM of the skull to virtually reconstruct a fracture of the lateral midface. Their SSM was developed from a dataset of 175 CT image volumes of Caucasian ancestry. They illustrated that SSM-based fracture reconstruction was more precise compared to standard procedure (mirroring technique-based) reconstruction. Paysan et al. (2009) used SSM of the skull and a face model developed from 20 CT images and 840 3D surface scans, respectively, to establish a relationship between the two models which were used in facial reconstruction. However, they had to include the age, weight and sex of a subject to improve the reconstruction process and make it subject specific. In product design for biomechanics, finite element (FE) analysis is often used (Neuert, 2016). However, the process of developing a FE mesh from CT images is tedious since most researchers focus on highly validated FE meshes at the expense of variability of anatomical geometry. Neuert (2016) used a SSM of the skull developed from CT images of 22 subjects to solve this issue. Sinha et al. (2016) used a SSM for segmentation of the paranasal sinuses of three patients on whom functional endoscopic sinus surgery was to be performed. Their SSM was based on CT images of the head from 53 subjects.

To obtain a 3D SSM of the skull, volumetric images of the skull are required. Such volumetric images contain information on the skull, the skin, flesh and the teeth. Apart from the skull, information on the other tissues need to be segmented out before the images can be used in the SSM development pipeline.

2.2.1 SKULL SEGMENTATION FROM CT IMAGES

The volumetric image data of the skulls obtained from CT require segmentation, to separate the skull from other tissues and leave out sections of the skull that are not required in the statistical models (Hristidis, 2009).

There exist different techniques of segmentation i.e. edge-based, region-based, feature-based, model-based and threshold-based segmentation (Yogamangalam & Karthikeyan, 2013). In edge-based segmentation, boundaries are identified to detect discontinuities in images (Pardhi et al., 2016). This is achieved by dividing the image into a foreground and a background using changes in pixel intensity of an image. Edge-based segmentation assumes that edges in an image represent object boundaries. Therefore, to identify objects in an image, edges in the image are detected (Pardhi et al., 2016). In region-based segmentation, pixels related to a particular object are grouped together for segmentation (Yogamangalam & Karthikeyan, 2013). Feature-based segmentation is based on the characteristics of an image such as texture, motion, or colour (Pardhi et al., 2016). In textured images, pixels of raised objects may be clustered together and segmented (Pardhi et al., 2016). Model-based segmentation utilises *a-priori* knowledge of the object to be segmented from the image together with a fitting algorithm to segment the object from the image (Heimann & Meinzer, 2009). Model-based segmentation is not dependent on intensity and is therefore applicable in low contrast images or in

images containing artefacts. However, model-based segmentation requires different *a-priori* information for the different objects to be segmented from the image (Heimann & Meinzer, 2009). Threshold-based segmentation is based on threshold values obtained from histograms of the image to be segmented (Gobindchandra & Santhosh, 2015).

The results of edge-, region-, feature-, model- and threshold-based segmentation can be evaluated using empirical goodness tests, empirical discrepancy tests or analytical tests (Zhang, 1996). Empirical goodness tests measure the performance of segmentation algorithms by measuring the quality of the resulting segmented images. Empirical discrepancy tests evaluate the differences between an ideal segmented image and an actual segmented image. Analytical tests analyse segmentation algorithms to evaluate their complexity.

Segmentation may be used to remove structures such as the teeth and the mandible from the volumetric skull image data (Berar et al., 2006).

2.2.2 LANDMARK IDENTIFICATION ON A TRAINING DATASET

Landmarks are points on anatomical objects within a population upon which correspondence across the objects is preserved (Dryden & Mardia, 2016). As such, landmark points on the 3D surfaces are identified to aid in establishing correspondence throughout the surfaces. Craniometric landmarks (see **Figure 1**) include bilateral and median landmarks on the skull (Storm & Alyson, 2010).

While it is possible for a trained professional with a background in anatomy to manually identify these landmarks, the task is tedious if the dataset has many skulls. Automatic craniometric landmark detection uses algorithms to register a reference skull with known craniometric landmarks onto a target skull (Kun, 2014). However, these algorithms are sensitive to artefacts and noise, which affect the integrity of the registration results later on in the statistical shape modelling process (Zhang et al., 2013). Manual land-making addresses artefacts much better and also ensures that there is consistency in assignment of landmark locations where the shapes of the skull in the training dataset vary widely. Once the landmarks are assigned on all the instances in the dataset, they are used for establishing correspondence.

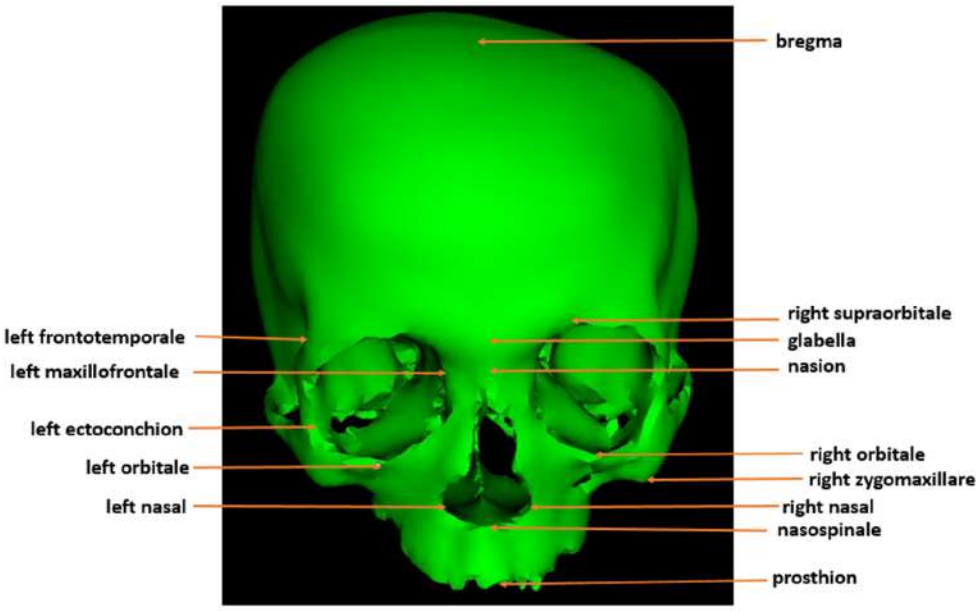

***Figure 1**: A skull with some anatomical landmarks.*

2.2.3 CORRESPONDENCE ACROSS A TRAINING DATASET

Correspondence is the close similarity that exists between instances in two different shapes or images of the same class (Tau & Hassner, 2016) whereas registration involves determining a transformation that relates corresponding positions between two or more shapes or images (Crum et al., 2004). The registration is determined using registration algorithms which are made up of three components: a similarity measure of how well the shapes or images being compared match; a transformation model specifying ways in which one shape or image can be changed to match another one using numerical parameters that specify instances of the transformation and; an optimization process that alters the transformation model parameters to improve the matching criterion (Crum et al., 2004).

For similarity measures where registration is based on anatomical images, the content of the image can be divided into two: intensity approaches or feature approaches (Tau & Hassner, 2016). In intensity approaches, intensity patterns in images are matched using statistical or mathematical criteria. The intensity similarity is measured between a reference and a source image and the transformation adjusted until a maximized similarity measure is achieved. The assumption in this approach is that at the correct registration, the two images are at their maximum similarity. This can be validated by tests that include calculation of squared differences in correlation or intensities. In the squared differences test, the images are assumed to be identical at registration with the only exception being Gaussian noise (Crum et al., 2004). Because intensity approaches do not use identifiable anatomical features, the registration lacks biological validity which is important for medical applications. Feature approaches use identifiable anatomical features such as anatomical landmarks

which ensure the registration has biological validity and is easily interpretable anatomically and physiologically (Crum et al., 2004). Reliably identified landmarks are used to establish explicit correspondence between paired landmark points in the two images and the two images are interpolated to infer the correspondence to the rest of the image volume.

A transformation model defines how an individual shape or image can be deformed to match onto another (Crum et al., 2004). The transformation model controls how features of one shape or image can be moved with respect to those of a second shape or image to improve the similarity between the two. How the transformation is adjusted to refine the image similarity is referred to as optimization. Hence a good optimizer is one that establishes the quickest transformation reliably.

In feature-based registration, sparse or dense correspondences are possible. In sparse correspondence, only a small set of features is used for registration whereas in dense correspondence, a higher number of features is used. The higher the number of features used, the better the global smoothness in the registration results and hence the preference for dense correspondence over sparse correspondence. Dense correspondence between shape surfaces can be achieved by surface parameterization (Floater & Hormann, 2005), group representation (Davies et al., 2002) or mesh registration (Hajnal & Hill, 2001).

2.2.3.1 Surface parameterization in establishing dense correspondence

Surface parameterization involves one-to-one mapping between a base domain and an input mesh (Floater & Hormann, 2005; Hormann et al., 2007). The base domain can be 2D such as a circle or 3D such as a sphere. 3D-based surface parameterization is more complex because of the shape of the base domain involved. Surface parameterization is limited to simple mesh surfaces (Bennis et al., 1991; Hormann et al., 2007; Sheffer et al., 2007) and is not suitable for complex shapes such as the skull.

2.2.3.2 Group representation in establishing dense correspondence

Group representation involves the development of SSMs together with an objective function that is used to assess how good and efficient the resulting model is (Davies et al., 2002; Davies et al., 2008). The objective function can be the determinant of a covariance matrix in the training dataset (Kotcheff & Taylor, 1998) or be based on the minimum description length (MDL) which describes the representation with the least information possible (Davies et al., 2002). In the former, the goodness and efficiency analysis of the shape model is based on the compactness and specificity of the model whereas in the latter, goodness and efficiency analysis is based on specificity and compactness of the model together with the generalizability of the model (Davies et al., 2002; Kotcheff & Taylor, 1998). A model is said to be compact if it has as little variance as possible and if it requires only a few parameters to define an instance. Specificity of a model is its ability to produce

shapes that are valid in terms of their resemblance to the shapes in the training dataset while generalization is the ability of a model to represent new shapes of the class similar to that of the training dataset (Davies et al., 2002). To ensure the shape models are compact and specific, studies have optimised the determinant objective function by using various techniques during the establishment of correspondence: a piecewise linear parameterization scheme (Davies et al., 2002) during landmark placement on the dataset; or a genetic algorithm (Jones & Forrest, 1995). The same can be achieved where the MDL objective function is used by continuous re-parameterization and optimization using a Cauchy kernel (Vishwanathan & Smola, 2004) and Nelder Mead algorithm, which are numerical functions used to find either the maximum or minimum of objective functions in multidimensional spaces (Gavin, 2013). Group representation is not suitable for complex shapes such as the skull or on large datasets because it is computationally expensive (Davies et al., 2002).

2.2.3.3 Mesh registration in establishing dense correspondence

Mesh registration involves the registration of mesh properties from a reference surface onto a target surface (Hajnal & Hill, 2001). First, the reference surface is registered onto the target surface before the vertex connectivity on the reference are transferred onto the target. The process can be conducted through rigid registration and non-rigid registration.

Rigid registration entails the registration of two surfaces without changing their shapes (Hajnal & Hill, 2001). The commonly used algorithms in rigid registration include the iterative closest point algorithm (ICP) (Besl & McKay, 1992), and the soft-assign Procrustes algorithm (Rangarajan et al., 1997) since these can work with meshes that have different numbers of vertices as inputs and reliably output the best similarity transformation between them. However, the two algorithms are only suitable for targets that are complete and have no artefacts. Additionally, the two algorithms cannot be used in cases where registration involves complex surfaces such as the skull because of the varying curvatures of the skull and the presence of teeth and holes (Drake et al., 2014). To overcome these challenges, variations of the ICP algorithm which include the functional ICP algorithm (Phillips et al., 2007) data aligned rigidity constrained exhaustive search (DARCES) (Chen et al., 1999) and K-D tree based ICP algorithm (Nutcher et al., 2007) can be used. These algorithms allow the use of incomplete targets as they only use a subset of the vertices located on the targets and normally, vertices that are close to the reference surface are the ones used.

Non-rigid registration, also known as deformable registration, is the non-uniform mapping between surfaces and involves measuring, varying and correcting discrepancies between the surfaces by deforming the reference surface and matching it onto a target surface (Galeotti, 2016). If there are two surfaces *A* and *B*, the two can be registered non-rigidly by establishing a deformation function that maps points of *A* onto *B* (Galeotti, 2016; Rueckert & Aljabar, 2010). This deformation function is able to map points of *B* back to *A*. The registration is established using three steps:

I. Overlap is detected between the source surface and the target surface.
II. Correspondence is established between the two surfaces.
III. The source surface is deformed to fit onto the target surface.

The deformation entails establishing closest points between the two surfaces, pruning, and performing global optimization and repeating the whole process sequentially until the surfaces converge (Li et al., 2008).

Algorithms used for non-rigid registration include Thirion demons algorithm (Dedner et al., 2007; Thirion, 1998; Vercauteren et al., 2007) graph and manifold matching (Zeng et al., 2010), mass-spring model (Boehler & Peitgen, 2009), local affine transformation (LAT) (Park et al., 2010), octree-splines (Szeliski & Lavallée, 1996), thin plate splines (TPS) algorithm (Bookstein, 1989; Brown & Rusinkiewicz, 2004; Hu et al., 2012), log-domain demons (LDD) algorithm (Vercauteren et al., 2007), free-form deformation (Abdelmunim & Farag, 2011; Myronenko et al., 2007) and Gaussian process registration (Gerig et al., 2014).

The Thirion demons algorithm is based on image registration concepts and tackles non-rigid registration as a diffusion problem (Thirion, 1998). Given two images *M* and *S*, "demons" are introduced onto the membrane of one surface, *M* and these demons enable the surface of *S* to diffuse inside or outside of *M*, depending on whether *M* is labelled inside or outside respectively. Graph and manifold matching approaches non-rigid registration as a graph matching problem by combining appearance and geometric similarities with intrinsic embedded information (Hotta, 2008). The spring-mass system algorithm models deformation using masses interconnected by springs (Shen et al., 2006). External forces generated from the surface content, while under constraints from the internal forces of the system, are used to guide the spring-mass system to the best matching configuration. Local affine transformation (LAT) is used in non-rigid registration of surfaces which have affines that move independently of each other. The LAT assigns prior definitions to each affine before deforming the source onto the target. In Octree-splines, deformation is represented as a volumetric transformation by performing a least square minimization of the distance between two surfaces of the shapes to be registered. One octree-spline is used to effect the minimization while another one is used to model the deformation (Szeliski & Lavallée, 1996).

In TPS, deformation is represented as an algebraic problem where the physical bending energy of a thin metal plate is expressed as a dependence on point constraints. For the interpolation of one shape surface onto another over a fixed set of nodes, the bending energy is taken as a quadratic form in the height of the stated surface (Bookstein, 1989). The log domain demons (LDD) algorithm (Cahill et al., 2009; Vercauteren et al., 2007) is used to find deformation vector fields (von Funck et al., 2006). These fields are usually smooth and invertible, and thus useful for creation of new

instances. It should be noted that LDDs are generated via exponential stationary velocity fields (SVF) (Vercauteren et al., 2007) which are a result of the registration process. The LDDs repeatedly minimise the energy function, which is composed of correspondence and regularization terms. The correspondence term is defined by SVF, while the regularization term adds smoothness. Free-form deformation is based on a grid of control points and works by moving and deforming control points of the source in a smooth and uniform way under some shape constraints until the shape of the source exactly matches the boundary of the target (Schwarz, 2007). The placement of the control points can be at any variable distance or location enabling the control of the deformation precision.

A recently reported non-rigid registration approach is the Gaussian Process registration. Traditional registration methods assume that deformation models remain constant over the image or object domain. However, Gaussian process registration allows for different regularization properties for different regions of the domain (Gerig et al., 2014). Therefore, Gaussian process registration can be used for differentiating two dissimilar tissues or regularizing regions in an object domain that has noise. Hence, Gaussian process registration is better suited at handling complex structures such as the skull or structures with outliers and noise. Registration is approached as a deformation problem where a reference shape is identified, and deformation fields from the reference shape to the target shapes induce correspondence (Gerig et al., 2014). Prior assumptions of the deformation fields, u, are modelled by a Gaussian process GP which is defined by a mean function, $m(x)$,and a kernel, $k(x, x')$, equation [2.1] (Williams & Rasmussen, 2006).

$$\boldsymbol{u \sim GP(m, k)} \qquad \textbf{[2.1]}$$

where u are the deformation fields modelled as a Gaussian process, with mean function m and kernel k.

A Gaussian process can be used in defining a prior distribution $p(u) \sim GP(\mu, k)$ of all the possible deformation fields of a registration process (Gerig et al., 2014). Hence, with the likelihood function of $p(I_T | I_R, u) \propto \exp(-D[I_R, I_T, u])$ where I_R and I_T are the reference and target shape respectively and D is the distance between the two surfaces, the registration task can be framed as a MAP estimation problem (equation [2.2]).

$$\boldsymbol{\arg max\, p(u | I_R, I_T) = \arg max\, p(u) p(I_T | I_R, u)} \qquad \textbf{[2.2]}$$

This MAP estimation problem can be solved by solving a minimization problem in the reproducing kernel Hilbert space (RKHS) defined by k (equation [2.3]) (Gerig et al., 2014).

$$\boldsymbol{\arg min\, D[I_R, I_T, \mu] + \gamma ||u||_k^2} \qquad \textbf{[2.3]}$$

where D is the distance between the reference and target shape and $\gamma \in \mathbb{R}^+$is a regularization weight.

The minimization of equation [2.3] can be achieved by performing a low rank approximation of k from the first n Eigen functions in equation [2.4].

$$\boldsymbol{k(x, x') = \sum_{i=1}^{\infty} \lambda_i \emptyset_i(x) \emptyset_i(x')^T} \quad \textbf{[2.4]}$$

where $(\lambda_i \emptyset_i)$ are the eigenvalue and eigenvector pair of the integral operator $\tau_k f(.) \coloneqq \int_{\Omega} k(x, .) f(x) dx$.

2.2.4 SHAPE ALIGNMENT OF DATASETS IN DENSE CORRESPONDENCE

After establishing correspondence within the examples in the training dataset, the various examples need to be aligned. This is to remove variations between the training dataset caused by factors other than the variation in shape. The variations include location, scale and rotation of the instances in the training dataset. However, if the investigation is looking at shape changes as well, the scale is not filtered out. In SSMs of the skull, the training data is likely to contain skulls from specimens of different age, sex or ancestry which are characterised by variation in size and in such a case, the scale need not be filtered out. Shape alignment is achieved using measurements between features within the dataset such as Hausdorff distance (Huttenlocher et al., 1993); strain energy (Sclaroff & Pentland, 1995) or Procrustes distance (Dryden et al., 1997). The first two metrics are applicable in datasets where the shapes have unequal number of landmark points whereas Procrustes distance can only be used if the examples in the dataset have equal number of landmark points. Where the Procrustes distance metric is used, Procrustes analysis is employed, which involves matching of one observation onto another by performing a least square matching between two configurations using a similarity transformation (Stegmann & Gomez, 2002). The estimation of the similarity parameter is done by minimizing the squared Euclidean distance (Goodall, 1991). Procrustes analysis is of two forms: ordinary Procrustes analysis and generalized Procrustes analysis (GPA). The former entails the matching of only two shapes at a time whereas GPA involves several shapes simultaneously. GPA follows the steps outlined below:

i. Selection of one example in the dataset as a reference shape.
ii. Alignment of the rest of the examples in the dataset to the reference shape using the centroid of each of the examples and moving them to the origin before normalizing the examples to unit sizes and rotating them to a new approximate mean shape.
iii. Calculation of a new approximate mean shape from the aligned ones.
iv. Repeating the same process from step (ii) to step (iii) if the result of step (ii) are different from that of step (iii)
v. However, if the results of step (ii) are similar to those in step (iii) then the shapes in the dataset are said to be aligned.

2.2.5 SHAPE MODEL BUILDING FROM A CORRESPONDING DATASET

After alignment of the data in the training dataset, the difference in locations of corresponding landmarks across the examples are assumed to represent only a variation in shape. However, in the case of a skull dataset, there is also a variation in scale encoded within the dataset (Albrecht et al., 2013; Gerig et al., 2014). Models developed using a dataset whose scale metric has been removed are called shape models whereas those developed with dataset that still encode scale are referred to as form models (Bouabene, 2017). However, in this document and indeed in most of the literature, the statistical models are referred to as SSM even though the training dataset still encode the scale metric.

Traditional SSMs, also known as point distribution models are developed by modelling distributions over meshes in correspondence. Gaussian process can be adopted onto point distribution models by modelling the distribution over functions that define the in-correspondence meshes (Gerig et al., 2014). The meshes are converted into a set of deformation fields with respect to a reference mesh. To retrieve a particular mesh, the reference mesh is warped with the deformation fields of that mesh. A SSM developed using the GP is just a Gaussian process model of deformation fields which captures the mean shape and the variation in shape from the mean shape across the training dataset. The SSM can be represented as in equation [2.5] (Williams & Rasmussen, 2006). Hence, every deformation (u), is represented by the Gaussian process i.e. $u \sim GP(\mu, k)$. Therefore by using Gaussian processes, distributions over functions can be modelled without choosing a discretisation in advance. To model a 3D vector field, the Gaussian process $GP(\mu, k)$ can be defined where $\mu: \Omega \to \mathbb{R}^3$ represent the mean function and $k: \Omega * \Omega \to \mathbb{R}^{3*3}$ represent the covariance function. These functions represent mean deformation $\mu(x)$ for all the points $x \in \Omega$ and the covariance $k(x, x')$ between the deformations for any pair points x and x'. The Karhunen-Loeve (KL) expansion is used to rewrite the Gaussian process model $u \sim GP(\mu, k)$ as illustrated by Kittle and Young (1973) (equation [2.5]).

$$\boldsymbol{u = \mu + \sum_{i=1}^{\infty} \sqrt{\lambda_i}\, \emptyset_i \alpha_i,\ \alpha_i \sim N(0, 1)} \qquad \textbf{[2.5]}$$

where $\emptyset_i, \lambda_i, i = 1, \ldots, \infty$ are eigenfunction and eigenvalue pairs of the operator associated with the covariance function k.

For discretised deformation fields, each discretised function $\breve{u}$ is represented as a vector $\vec{u} = (u_1, \ldots, u_n)^T$ whose distribution over $\breve{u}$ is defined as (equation [2.6]).

$$\boldsymbol{\vec{u} \sim N(\vec{\mu}, K)} \qquad \textbf{[2.6]}$$

where K is a symmetric, positive semi-definite matrix which can thus admit eigen decompositions (Gerig et al., 2014).

$$\boldsymbol{K} = \phi \boldsymbol{D} \phi^{\boldsymbol{T}} = \begin{pmatrix} \overrightarrow{\emptyset_1} & \cdots & \overrightarrow{\emptyset_n} \\ \vdots & \ddots & \vdots \\ & \cdots & \end{pmatrix} \begin{pmatrix} \boldsymbol{d_1} & \cdots & \\ \vdots & \ddots & \vdots \\ & \cdots & \boldsymbol{d_n} \end{pmatrix} \begin{pmatrix} \overrightarrow{\emptyset_1} & \cdots & \overrightarrow{\emptyset_n} \\ \vdots & \ddots & \vdots \\ & \cdots & \end{pmatrix}^{\boldsymbol{T}} \quad \textbf{[2.7]}$$

where $\overrightarrow{\emptyset_i}$ is the i-th column of ϕ and it represents the i-th eigenvector of K, while $d_i$is the eigenvalue that correspond to $\overrightarrow{\emptyset_i}$. Decomposing equation [2.7] using a Singular Value Decomposition (SVD) results in an expansion of the eigenpairs $d_i, \overrightarrow{\emptyset_i}$.

.

$$\vec{\boldsymbol{u}} = \vec{\boldsymbol{\mu}} + \sum_{\boldsymbol{i=1}}^{\boldsymbol{n}} \sqrt{\boldsymbol{d_i}}\, \overrightarrow{\emptyset_1} \boldsymbol{\alpha_i} \sim \boldsymbol{N(0,1)} \quad \textbf{[2.8]}$$

where the anticipated value of $\vec{u}$ is $E[\vec{u}]$= $\vec{\mu}$ and the covariance matrix $E[(\vec{u} - E[\vec{u}])(\vec{u} - E[\vec{u}])^T] = K$. Therefore, $\vec{u} \sim N(\vec{\mu}, K)$.

Principal component analysis is a KL expansion but for discrete representation of data with their covariance matrix being based on example shapes. Therefore, given a set of discrete deformation fields $\widetilde{u_1}, \ldots\ldots, \widetilde{u_m}$ which are also representable as vectors $\overrightarrow{u_1}, \ldots\ldots, \overrightarrow{u_m}, \overrightarrow{u_i} \in \mathbb{R}^n$, where $\overrightarrow{u_1}$ is a vector representing a full deformation field and n is a large value, the assumption of the PCA is that the covariance function is estimated from the example shapes (Williams & Rasmussen, 2006).

$$\sum = \frac{\boldsymbol{1}}{\boldsymbol{m}} \sum_{\boldsymbol{i=1}}^{\boldsymbol{m}} (\overrightarrow{\boldsymbol{u_i}} - \overline{\boldsymbol{u}})(\overrightarrow{\boldsymbol{u_i}} - \overline{\boldsymbol{u}})^{\boldsymbol{T}} =: \frac{\boldsymbol{1}}{\boldsymbol{m}} \boldsymbol{X} \boldsymbol{X}^{\boldsymbol{T}} \quad \textbf{[2.9]}$$

where the data matrix X is defined as $X = (\overrightarrow{u_1} - \bar{u}, \ldots\ldots, \overrightarrow{u_m} - \bar{u}) \in \mathbb{R}^{n*m}$, and $\bar{u}$ is the sample mean i.e. $\bar{u} = \frac{1}{m} \sum_{i=1}^{m} \overrightarrow{u_i}$.

The rank of Σ in equation [2.9] is at most m as opposed to that in equation [2.8] which is n. This allows for efficient computation of decomposition by performing a SVD on the smaller data matrix X. As a result, the expansion reduces to $\vec{u} = \bar{u} + \sum_{i=1}^{m} \sqrt{\succ_i}\, \overrightarrow{\emptyset_1} \alpha_i, \alpha_i \in N(0,1)$. Hence, any deformation $\vec{u}$ can be completely specified by a coefficient vector $\vec{\alpha} \in \mathbb{R}^m$. The eigenvectors $\overrightarrow{\emptyset_i}$ derived from the coefficient vectors are the principal components, with the first principal component being referred to as main mode of variation. The variation represented by an instance of the principal component can be visualised by varying entities of the coefficient vector to correspond to any standard deformation away from the mean (Bouabene, 2017).

2.3 STATISTICAL APPEARANCE MODEL

Statistical appearance models are a natural extension of SSMs since they describe the combination of a variation in shape and image intensity (Bonaretti et al., 2014). In the case of radiological images, this intensity can be used to study bone mechanical properties as it corresponds to bone mineral density (BMD) in calibrated radiological images. Hence, SAMs can be used to obtain

bone properties in 3D from CT images, for example, which can be used to simulate finite elements for assessing bone quality and strength (Bonaretti et al., 2014). SAMs are developed based on either mesh or image models.

2.3.1 MESH BASED APPEARANCE MODEL

In mesh-based SAMs, examples in the training dataset are represented by meshes in the first step of the SAM development pipeline (Bonaretti et al., 2014). From the SSM as documented in Section 2.3, a reference volumetric mesh is generated from the surface mesh of the mean of the SSM via planar parameterization (Saboret et al., 2006). The volumetric mesh of the reference is then registered onto the surface mesh of all examples in the training dataset creating iso-topological tetrahedral finite element meshes. During the registration process, pre-selected landmarks are used as constraints and mesh morphing is performed on the volumetric meshes by using the vertex connectivity of the surface meshes as a guide. The iso-topological tetrahedral finite element meshes are however too dense and rough and need to be improved so as not to affect the outcome of the appearance model. To improve the quality of these meshes, each mesh is smoothed and decimated using the Laplacian operator (Belkin & Niyogi, 2002). Smoothing results in node removal, which can be re-established using the Markov random field surface reconstruction (MRFSR) as proposed by Paulsen et al. (2009).

Information related to the grey values of the meshes is then determined by warping examples in the training dataset onto the mean shape in the SSM using either thin-plate spline or piecewise affine as demonstrated by Stegmann (2000). However, this information is affected by different grey values when more than one CT scan machine is used for scanning different specimens of one dataset. To eliminate the varying grey value effect, normalisation of the values is performed. Normalization is achieved by first recording the grey values into a texture vector then applying an offset β and scaling α, to the values (equation [2.10]).

$$\boldsymbol{g = (g_{in} - \beta 1)/\alpha} \qquad \textbf{[2.10]}$$

where g is normalised intensity value and g_{in} is original intensity value (Stegmann, 2000). β and α are determined using equation [2.11] where n is the number of elements in the vectors and $\overline{g}$ is the mean of the normalised data.

$$\boldsymbol{\beta = (g_{in} * 1)/n \text{ and } \alpha = g_{in} * \overline{g}} \qquad \textbf{[2.11]}$$

The normalization is defined by the mean of the normalised data which is obtained via a recursive process. Hence, to identify a stable solution, one of the intensity value sets is used as the mean and the rest of the sets are aligned to it and the mean estimated iteratively (Stegmann, 2000).

Once the texture has been normalised, PCA is applied on the normalised data to generate a linear model (equation [2.12]).

$$\boldsymbol{g = \overline{g} + P_g b_g} \quad \textbf{[2.12]}$$

where $\overline{g}$ is the mean of the normalised grey-level vector, b_g is a set of grey-level parameters and P_g is a set of orthogonal modes of variation.

2.3.2 IMAGE BASED APPEARANCE MODEL

In image-based SAM, dataset objects are represented by volumetric images for the whole SAM development pipeline. Hence, images are transformed to finite element meshes just in the final step of the development pipeline. Anatomical correspondence is established by image registration (Brown, 1992). Image registration is achieved by registering other examples in the training dataset onto the reference using registration algorithms documented in Section 2.3.3. The reference is obtained using an iterative procedure where one example is selected as reference from a training dataset and all other dataset objects are registered to it. Afterwards, the average transformation is calculated and the example in the training dataset closest to the average transformation is selected as the new reference. This process is repeated until convergence is achieved. After establishing correspondence, finite element meshes for the examples in the training dataset are created and the steps involved include: the extraction of bone surface using the matching cube algorithm (Lorensen & Cline, 1987). These surfaces are then enhanced using the MRFSR (Paulsen et al., 2009) which is followed by the creation of volumetric meshes via planar parameterization (Saboret et al., 2006). Intensities warped to the reference by the examples in the training dataset are used to generate the SAM. Here, PCA is performed on these intensities to generate covariance matrix which are then used to generate the SAM.

As reported by Bonaretti et al. (2014), mesh-based SAMs have a greater compactness when compared to image-based ones. Additionally, mesh-based SAMs perform better in recreating bone shapes than image-based SAM. They also are computationally less expensive to develop than their image-based counterpart; taking a shorter time in constructing a SAM. However, image-based SAMs produce better results when constructing original intensities and perform better in finite element simulation than mesh-based SAMs.

2.4 STATISTICAL SHAPE AND APPEARANCE MODEL VALIDATION

Validation is crucial in scientific studies as it enables researchers to assess the credibility of proposed methodologies. In SSM, the leave-one-out cross validation is normally used (Kohavi, 1995).

Here, one example is removed in turn from a training dataset and used as a test set. The remaining examples are used to create the statistical model as described in Section 2.3. A fitting procedure is then used to fit the model to the dataset not included in the model. The resulting reconstructed surface of the test case is compared to the surface of the ground truth of that test case. This validation process is often used in cases where there are few samples in the data set making it difficult to create a separate training set and validation set (Luo et al., 2013).

2.5 SUMMARY OF LITERATURE REVIEW

This chapter reviewed the skull and the steps involved in the development of both a SSM and a SAM. Volumetric image data of skulls are segmented to obtain tissues needed for the SSM and SAM development pipeline. Correspondence (across 3D surfaces obtained from the segmentation process) is established based on intensity or features such as landmarks. In the former, an entire image is mapped onto another whereas in the later, points are used for mapping where landmark points are placed on the crania. The more landmarks used, the better the correspondence accuracy hence dense correspondence preference to sparse correspondence. Dense correspondence can be achieved by surface parameterization, group representation or mesh registration. Surface parameterization and group representation are not suitable for complex shapes like the skull; mesh registration is better suited. Mesh registration can be achieved either through rigid or non-rigid registration. Rigid registration cannot deform shapes during the registration, but non-rigid registration does, and hence is best suited where examples in a training dataset need to be varied, corrected and deformed to achieve the best mapping across the training dataset.

Non-rigid registration can be achieved using the Thirion demon algorithm, free-form deformation, graph and manifold matching, the mass-spring model, local affine transformation, octree splines, TPS, the LDD algorithm or Gaussian process regression. Gaussian process regression, unlike the other non-rigid registration techniques, first generates an instance (posterior model) that is close to the target being registered. This posterior model is then deformed to fit onto the target hence ensuring the final fit is as close as possible to the target. Upon establishing correspondence, the examples are aligned either using Hausdorff distance, strain energy or Procrustes distance. Where the examples in the training dataset have an equal number of landmarks, Procrustes distance is used. General Procrustes analysis is preferable to the ordinary Procrustes analysis because after every iteration of the alignment process, a mean shape is calculated and therefore the mean shape achieved at the end of the iteration is unbiased. A SSM can then be developed using the process documented in Section 2.3.5 followed by the development of a SAM. Statistical appearance model development can be based on meshes or images. If the SAM is to be built on top of a SSM, a mesh-based SAM is preferable since the meshes reproduce shapes more accurately. After a SSM and a SAM have been built, the models are validated using the leave-one-out validation process which is applicable where the number of

examples in the training dataset is not enough to warrant the development of a separate validation model from the statistical model.

3 DATA COLLECTION AND PROCESSING

Data collection is the process through which a researcher collects information needed to answer a research question or hypothesis (Sapsford & Jupp, 2006). For effective data collection, the researcher decides on the type of data they require, how they collect this data, who will collect this data and when the data is to be collected, based on the research question. Pre-processing is the process of altering data to either enhance, change or remove certain features in the data (Hedley & Yan, 1992).

Many biological imaging applications utilise volumetric image datasets which aid in the visualization of 3D dataset using 2D projections of the dataset (Goshtasby et al., 2000). These images can be viewed by extracting their iso-surfaces and rendering them on mesh grids. However, before the iso-surface extraction from the volumetric images, segmentation is required to label-out sections of the image that need not be viewed.

This chapter details the data collection and pre-processing of the dataset that was used for the project. The first section outlines the image acquisition. The second section details the image segmentation and pre-processing steps. The third section details the post-processing steps. Afterwards, results for the post-processing are provided and a summary on the entire chapter is provided at the end. The Amira v6.2.0 software (Fei Imaging- www.fei.com) was used in the segmentation and post-processing of the dataset and the procedures were conducted on an Intel (R) Core (TM) i7-4790 CPU @3.60GHz running a 64-bit operating system on a 32 GB RAM.

3.1 CT DATA ACQUISITION

The development of the SSM and SAM was based on CT images of the skull. The CT images contained information on the name, age, sex, ethical group, occupation and birthdate all of which were anonymised. Access to the dataset was restricted to authorised personnel. The CT images had been acquired from a total of 16 cadaveric heads from deceased individuals who had donated their bodies for scientific research. Eleven of the cadavers were of Caucasian ancestry, two were from African ancestry and three from Mixed-race ancestry. Of the 16 cadaveric heads, six had been preserved by embalming using formaldehyde and 10 had been preserved by freezing. The specimens were imaged at the University of Cape Town Private Academic Hospital using a Toshiba-64 slice scanner (Mather, 2005). The scan thickness was 0.9 millimetre, the data collection diameter was 500 millimetres; exposure time was 1114 millisecond with an exposure of 400 milliamphere seconds. Before the cadaveric heads were scanned, they were sealed inside wrapped plastic to prevent contamination to the scanning room and to allow for hygienic manipulation of the heads by the

radiographer during the scanning. The CT scan images were stored in DICOM format for easier digital transfer and storage.

The number of skull CT scan images was limited because of the availability of cadaveric samples. While it would have been possible to add more samples in the training dataset derived from dry bone, dry bone has a different density and would not be representative of the bone-density found within a cadaveric sample (Mohiuddin, 2013).

Ethics approval for using these images was obtained from the Human Research Ethics Committee (HREC REF: 264/2017).

3.2 PRE-PROCESSING AND SEGMENTATION OF CT IMAGES OF THE SKULL

The CT images were composed of varying pixel intensities which characterised the different tissues of the imaged cadaveric heads. The various tissues included the muscle, bone and fat tissues. However, in this project, only the bone was required in the development of the statistical models. Therefore, segmentation was needed to label out the fat and muscle tissues. Through segmentation, these tissues would be removed from the dataset leaving the bone (skull) to be used in subsequent steps. Before the images could be used for the development of the SSM and SAM, they were pre-processed. The pre-processing involved removal of metallic dental filing artefacts from the CT images that had such artefacts.

Dental fillings in some of the cadaveric samples created streaks in the CT images which degraded the image quality. This made it difficult to effectively interpret the boundaries between tissues in the images. Different tissue types could therefore not be correctly labelled using the grey scale values of the pixels in the image. The dental fillings had attenuated the x-ray beams during image acquisition creating the streaks. Of the 16 sets of CT scans, three had metallic dental filing artefacts and the metal deletion technique (MDT) (Boas & Fleischmann, 2011) for CT metal artefact reduction was used to remove the artefacts from the CT images. The MDT discards inaccurate metal data in the CT images. In reconstructing the metallic part of the images, the MDT iteratively replaces the metal data with forward projected values.

The processed CT images were segmented using Amira v6.2.0 software (Fei Imaging- www.fei.com) to remove teeth and the mandibles from the skulls. Thresholding is the segmentation technique used in the Amira v6.2.0 software (Fei Imaging- www.fei.com) and which was used in this project. The thresholding technique is expressed in equation [3.1].

$$\boldsymbol{T = T[x, y, p(x, y), f(x, y)]} \qquad \textbf{[3.1]}$$

where T is the threshold value; x, y are the coordinates of the threshold value point and; $p(x, y)$and $f(x, y)$ are points located in the gray level image pixels. Therefore, an image thresholding is defined as equation [3.2].

$$g(x, y) = f(x) = \begin{cases} 1, & if\ f(x,y) > T \\ 0, & if\ f(x,y) \leq T \end{cases} \quad \textbf{[3.2]}$$

The thresholding technique in equation [3.1] can be either local, Otsu (Otsu, 1979) or basic global thresholding (Gobindchandra & Santhosh, 2015). In local thresholding, segmentation is conducted based on local thresholding such that the threshold for each pixel in the image is compared to the average of the surrounding pixels. If the value of the current pixel under comparison is lower than the average, the pixel is set to white, or black otherwise (Gobindchandra & Santhosh, 2015). Otsu's thresholding aims to reduce and maximise the within-class variance and inter-class variance respectively, within an image (Otsu, 1979). Basic global thresholding segmentation is based on a single pixel in the image and is applicable where pixel values of the foreground and background of the image are approximately consistent over the entire image. First, an initial estimate of the threshold value (*T)* is selected, and this threshold value is used to segment the image. The segmented image is composed of two groups of pixels: pixels whose values are above the threshold value *(if(x,y)> T)* and those whose value is equal or below the threshold value *(f(x,y) ≤ T)*. For accurate results, the threshold value can be iteratively determined and used until the best segmentation result is achieved (Gobindchandra & Santhosh, 2015).

In Amira, the DICOM files for each of the 16 CT images sets were loaded onto the software. A multi-thresholding module was applied on the grey value of the CT images generating five regions separated by four different thresholds as shown in **Figure 2**.

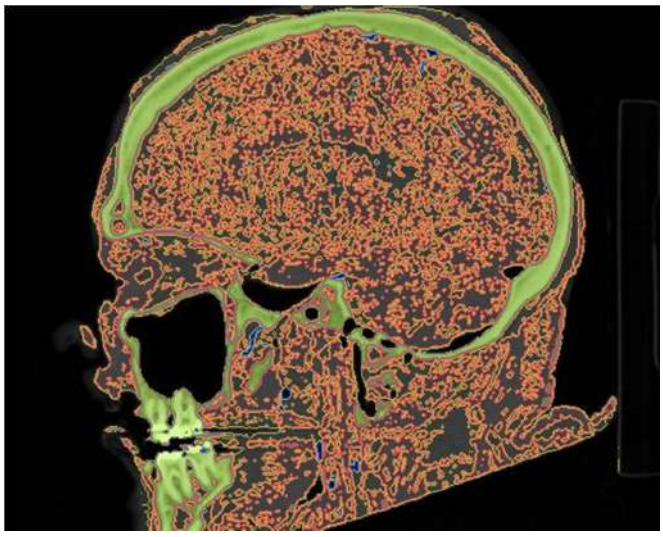

Figure 2: *Results of preliminary segmentation using the multi-thresholding module in Amira. Green regions are bone and teeth, yellow regions are fat tissues, orange-red regions are muscle tissues, black region is the exterior of the CT image*

Since only bone information was required in this project, the regions representing fat and muscle tissues were included in the exterior region thereby remaining with only two regions: Bone and teeth (green) and exterior (black) as shown in **Figure 3**.

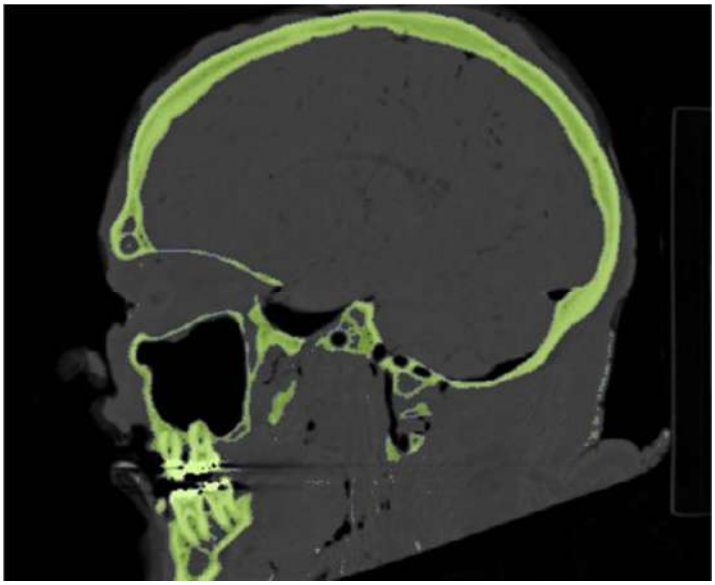

***Figure 3**: Results of segmentation after using the multi-thresholding module in Amira and after segmenting the fat and muscle tissues from the CT image. Green region represents bone and teeth and the black and grey regions represent the exterior of the CT image.*

The bone contained the cranium, the teeth and mandible. However, only the cranium was required in the SSM and SAM generation. The mandible had the same intensity value as the cranium and could therefore not be segmented based on grey level values. Hence, the mandible bone was selected manually and segmented out. To ensure only the mandible was segmented, manual correction was carried out over every CT slice containing both the crania and mandible. For the teeth, a similar process of manual segmentation was used. To ensure only the exterior surface of the segmented crania would be generated, the "Fill hole" command in Amira was used to fill holes on the segmented crania (Guide, 2009). The command worked by generating a complement of the image in every slice that made up the volumetric CT images. The edges of each image were then subjected to Geodesic dilation into the inverted image and using the complement of the former image, the holes in the slices were filled up (Guide, 2009). Holes that were not filled using this command were manually selected and added to the bone tissue.

The results of the segmentation process were evaluated to ensure they were representative of the objects in the CT images. The evaluation of the segmentation used empirical goodness evaluation methods that assessed the quality of the segmented images (Zhang, 1996). However, because there was no ground truth upon which to base the evaluation, human intuition was used to determine what conditions were to be satisfied for a segmentation to be accepted as ideal. The "goodness measure"

assessed the segmentation on intra-region uniformity, inter-region contrast and region shape (Zhang, 1996). A segmentation was deemed to be good if, for a CT slice, having objects (cranium bone) and a background (exterior region), the cranium was uniform rather than having holes (object and background) (**Figure 4**).

***Figure 4**: Left - a segmented CT slice of the cranium with the background (black spots) encroaching on the cranium (white). Right - a segmented CT slice with a uniform cranium (white) without the background within it (no background spots within the crania).*

After the evaluation of the segmentation, the segmented (labelled) images were used to generate 3D surfaces. The generation of the 3D surfaces (**Figure 5**) was done using the "*Generate surface*" module in Amira.

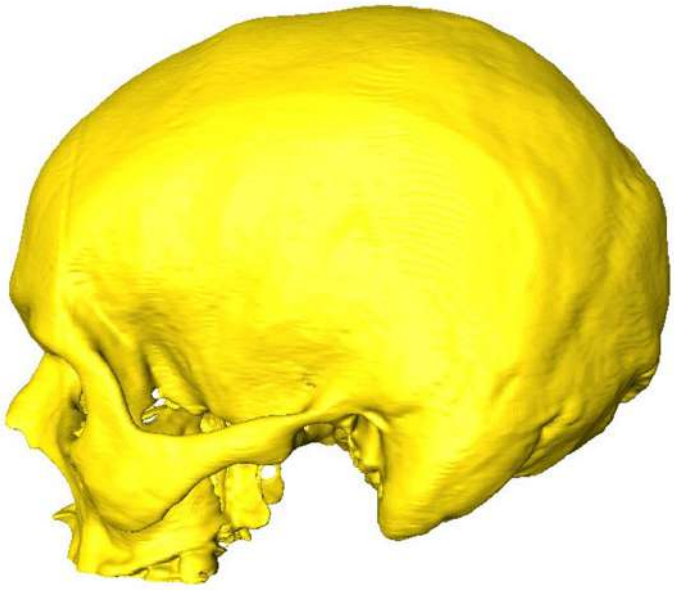

***Figure 5**: A 3D surface of one of the segmented skulls in the dataset.*

Each 3D surface was composed of approximately 1.5 million nodes and 3 million triangles. Across the 3D surfaces in the dataset, the nasal region and nasal sinuses had wide variations with some surfaces missing sections of the nasal region. Such a dataset could pose correspondence challenges during the registration process. The large number of nodes and triangles that represented the 3D surfaces resulted in large StereoLithography (.STL) files with no gain in topology and therefore, the files were post-processed to reduce their sizes.

3.3 MESH POST-PROCESSING

The 3D surfaces (meshes) generated in Amira contained too many vertices/nodes (with no associated gain in topological representation of the skull). The surfaces were on average represented by 1.5 million nodes which could have increased the processing time when calculating the PCA during the model development later on in the project. Post-processing allows for the simplification of meshes represented by excess features such that only the features that encode crucial characteristics of the meshes are maintained. The 3D surfaces were cleaned to remove outliers. The cleaning involved manually segmenting out the small fragments of would-be surfaces around the nasal-orbital area. This was to ensure that during the registration process, the best correspondence could be established across the 16 cranium surfaces around the nasal optic area. After cleaning the 3D surfaces, the surfaces were re-meshed and smoothened reducing the number of points and triangles representing them. The features were reduced from an average of 1.5 million nodes and 3 million irregular triangles to approximately 100,000 nodes and 200,000 isotropic triangles for each 3D surface (**Figure 6**).

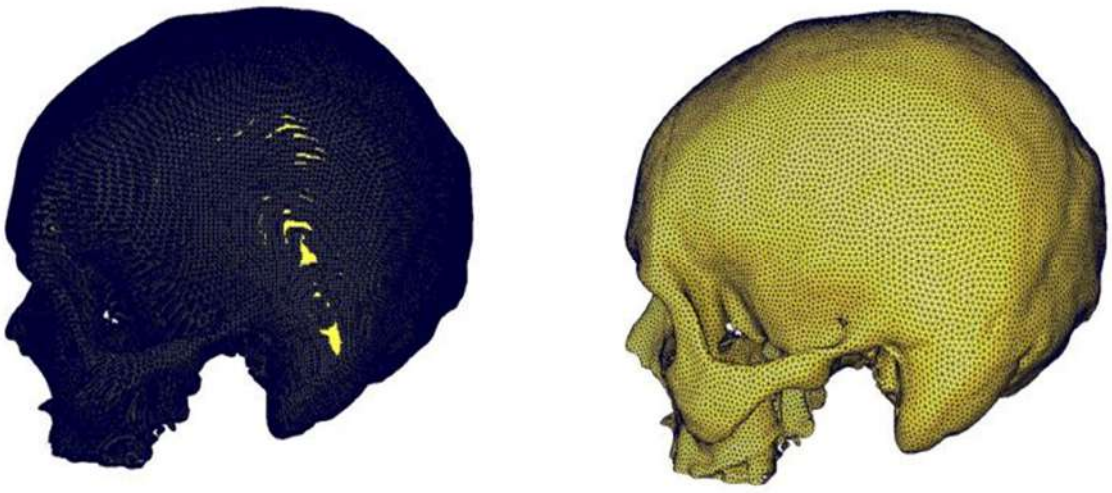

***Figure 6**: One of the 3D surface meshes before (left) and after (right) re-meshing. The surface on the left is composed of approximately 3 million irregular triangles. The surface on the right is composed of approximately 200,000 isotropic triangles.*

Care was taken to find a balance between number of nodes and topological representation of important features on the skull. An assumption was made that, provided the features of the crania were maintained in the down-sampled 3D surfaces, the meshes were acceptable. This is because it is only the features that are required during the registration process (Tam et al., 2013). The re-meshed and smoothened surfaces were saved as StereoLithography .STL files (Szilvi-Nagy & Matyasi, 2003).

3.4 SUMMARY

Metal deletion technique was used for pre-processing of the CT images to remove artefacts caused by dental fillings. The pre-processed CT images were manually segmented to separate the mandible and teeth from the cranium. Because the mandible and teeth had been assigned similar intensity values to the cranium in the Amira software, it was difficult to distinguish them from the cranium. Therefore, the manual segmentation depended on intuition of the observer doing the segmentation. As such, the process was not repeatable. The evaluation of the segmentation results was also based on human intuition. The goodness measure used to evaluate the segmented images was subjective. This can be considered a limitation of the protocol.

During the post-processing, the choice of the number of features to retain after the re-meshing was determined based on the assumption that, provided the features of the crania were maintained in the down-sampled 3D surfaces, the meshes were acceptable. There was no test conducted to determine the number of nodes and triangles needed after the re-meshing process. After the registration process, the number of nodes and triangles of the 3D surfaces in the training dataset would be replaced by features of the reference mesh. These re-meshed (down-sampled) 3D surfaces were used in the next step of registration (Chapter 4).

4 REGISTRATION OF 3D SURFACE MESHES

Registration is the process of determining a transformation that relates corresponding positions between two or more shapes or images (Crum et al., 2004). Thus it is the transformation of multiple instances of 3D data into a similar coordinate system (Tam et al., 2013). During the generation of 3D surfaces from volumetric images, objects are represented by meshes. The number of nodes that make up the meshes depend on the size and shape of the object being represented. Anatomical objects such as the skull have at least some variation across a given population. Hence, no two 3D surfaces of the skull from a population can be represented by the same number of nodes. The data has to be registered to the same coordinate system to establish correspondence across the dataset. After registration and assignment of correspondence, the surfaces in a dataset have similar numbers of nodes and triangles which can be analysed to obtain variations across the dataset - an important step in SSM and SAM development. This chapter details the registration of the 3D surfaces that were re-meshed in the previous chapter in order to establish dense correspondence across the mesh dataset. The original 3D mesh surface is referred to as target throughout the chapter. The first section of the chapter outlines landmark identification on the targets. The second section details the establishment of dense correspondence across the targets. A summary is provided at the end of the chapter reviewing key points and observations. The Scalismo software (Anon, 2017b) was used in landmark identification and in the registration process which were both conducted on an Intel (R) Core (TM) i7-4790 CPU @3.60GHz running a 64-bit operating system on a 32 GB RAM.

The registration process adopted here is dependent on features which constraint the movement of one shape to the other during the registration process. In anatomical shapes, anatomical features can often be used and in the work presented here, these features are referred to as landmarks.

4.1 LANDMARK IDENTIFICATION ON THE TARGETS

Landmarks are points on anatomical objects within a population upon which correspondence is preserved no matter the variation in size, pose or rotation across the objects (Dryden & Mardia, 2016). In the registration process presented below, a similarity measure of landmark points between a reference (***R***) and targets is used. A reference is an ideal object belonging to the same class as the targets and whose features are registered onto the targets**.** This similarity measure is applied to ***R*** therefore aligning and deforming ***R*** to fit targets. Hence, it is essential to identify anatomical landmarks across a dataset that can be reliably identified across all the targets.

A total of 33 easily identifiable anatomical landmarks were selected from both the median and bilateral landmarks of the skull. These landmarks were used as constraints in the registration process to establish correspondence across the targets. The anatomical landmarks included the:

bregma, glabella, nasion, nasospinale, prosthion, staphylion, basion, opistion, opithocranion, left and right euryon, left and right frontotemporale, left and right supraorbitale, left and right maxillofrontale, left and right ectoconchion, left and right orbitale, left and right nasal, left and right zygion, left and right zygomaxillare, left and right endomolare and, left and right foramen magnum (Rooppakhun et al., 2011). These landmarks were identified manually using the Scalismo software (Anon, 2017b) on each of the 16 surfaces as shown in **Figure 7**.

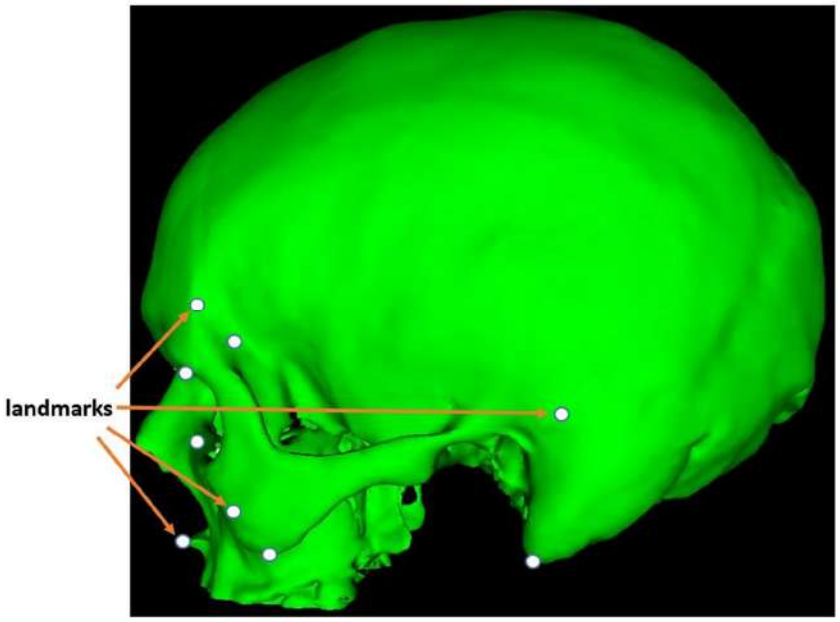

***Figure 7**: Examples of Landmarks (white) identified on one of the targets*

4.2 ESTABLISHING DENSE CORRESPONDENCE ACROSS THE TARGETS

In feature-based registration, the more features used the better the registration results in terms of the global smoothness of the registered surfaces. Hence, dense correspondence is preferred over sparse correspondence. In statistical models, the nodes of the meshes are used to establish correspondence. To achieve the correspondence, a reference mesh, ***R*** with known number of nodes and triangles is deformed to fit a target. The deformed ***R*** is used as a representation of the target; and can be referred to as ***DefR***. Because one ***R*** is used to establish correspondence across the entire dataset of targets, the deformed references (***DefRs***) have the same number of nodes and triangles and hence all ***DefRs*** are in correspondence with one another. In this document ***DefRs*** are also called in-correspondence objects. The reference needs to be flexible enough to deform onto the targets. A simple way to do this is to use ***R*** to develop a free-form deformation (FFD) model. The FFD can be parameterized to be flexible enough to deform onto the targets and used as such, in the registration process.

4.2.1 Developing a free form deformation model

Free-form deformations are geometric techniques used to model deformations of rigid objects (Barr, 1987). Free-form deformation works by enclosing an object within a hull and transforming the object as the hull is deformed. In this project, FFD deformation was used to establish dense correspondence. A FFD model was developed using the Gaussian process in Scalismo. In the GP process, a FFD model is developed from a mean mesh – an average mesh belonging to the same class as the targets that are to be registered and a kernel (Anon, 2017b). The mean mesh defines the average geometry of the shape and the kernel defines how the features of the shape deform from the mean shape. In this project, a mesh among the 16 targets with all the features expected for a human cranium was selected as the mean mesh. The kernel was based on the squared exponential kernel (SEK) (Babaud et al., 1986) as indicated in equation [4.1].

$$k_{SE}(x, x') = S \exp\left(\frac{-(x,x')^2}{l^2}\right) \quad \textbf{[4.1]}$$

where the value of *l* and *s* determine the span and degree of deformation of the features on the mean shape (Anon, 2017b).

The span of deformation refers to the radius from a point (x) on the mesh within which other points also undergo deformations when a deformation is applied to the point (x). The degree of deformation refers to the magnitude of deformation that a point (x) is subjected to (Anon, 2017b). In the case of a human face, if the point (x) is on the tip of the nose, a span of radius 50mm means that when the tip of the nose is deformed, the effects of the deformation affect the region around the nose within the 50 mm radius. If the tip of the nose is subjected to 3mm and 10mm of deformation, the former results in a smaller elongation of the tip of the nose than the later. Therefore, the two values of the SEK determine how and by what magnitude the mean shape of the FFD deforms. The FFD should be flexible enough to deform and fit as much as possible to all the targets. This can be ensured by using values of *l* and *s* that encode the maximum span and degree of deformation expected in the targets. These values were determined empirically by varying them one at a time and sampling the FFD model built from them by fitting it onto the targets.

Initially, a value of *s = 2mm* was used based on visual observation when the mean mesh was aligned to one of the targets i.e. ***TG1***. Approximately, 2mm of deformation would be required to fit the mean onto ***TG1*** (**Figure 8**). With *2 mm*, the value of *l* was varied and a FFD model built. The model was fit onto ***TG1*** to assess how well it represented that surface. At $s = 2mm$ and *l = 35mm*, the fitting recorded the lowest registration error (**Figure 9**). Therefore, *l = 35mm* was selected and used to determine the optimal value of *s*. The value of *s* was varied between *1 mm* and *3 mm*. At *s =*

3mm the lowest registration error was recorded when its FFD model was fit onto ***TG1*** as done previously with *l* (**Figure 10**).

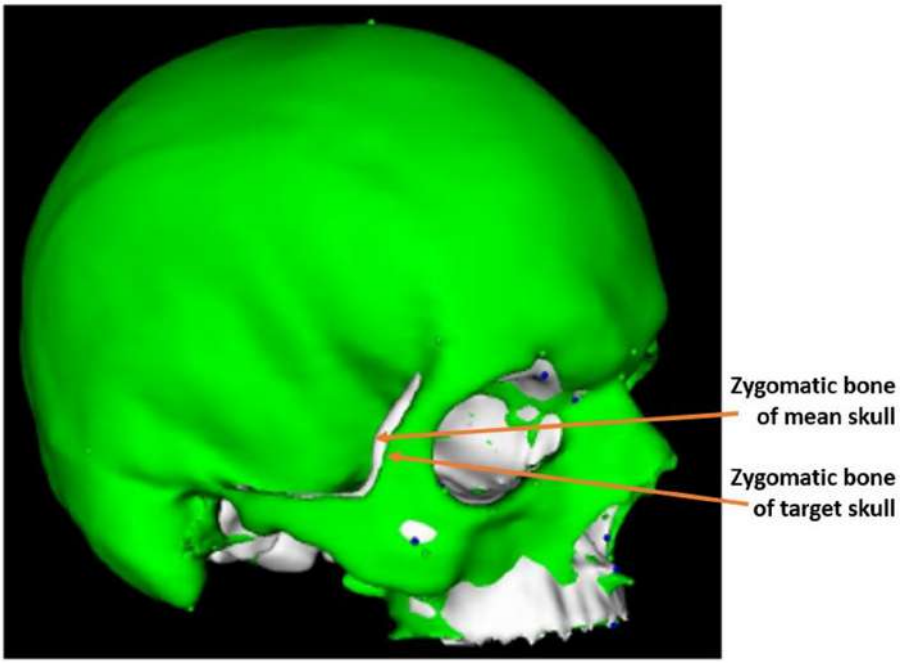

Figure 8*: A target skull- TG1 (green) overlaid onto the reference skull- R (white). TG1's zygomatic bone requires approximately 2mm of deformation to correspond to the zygomatic bone of R.*

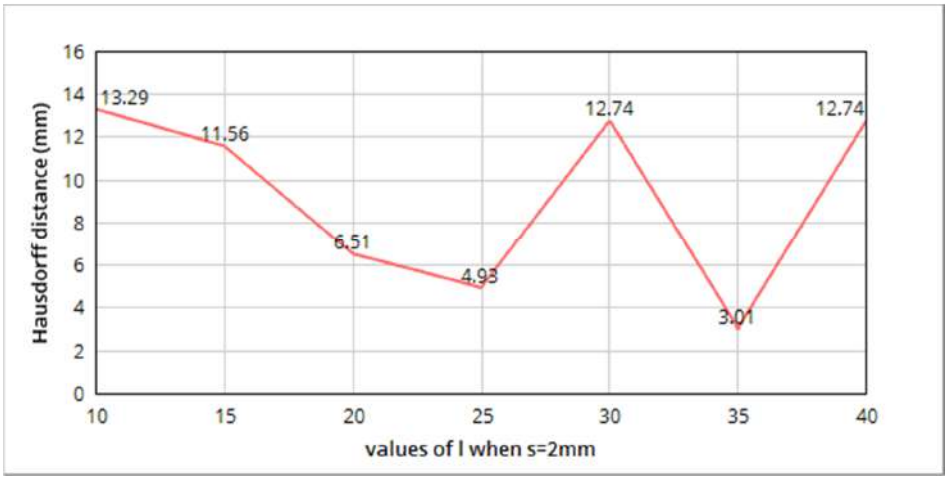

Figure 9*: Registration error when l is varied, and s is kept constant at 2mm.*

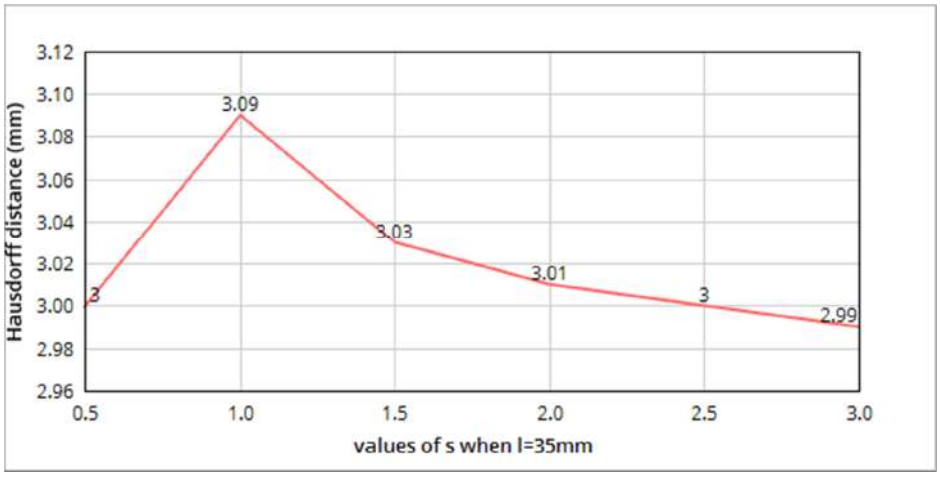

***Figure 10**: Registration error when s is varied and l is kept constant at 35mm.*

However, the value of s was above *2mm* which is recommended as the upper limit for deformations of anatomical shapes (Anon, 2017a). Hence the value of $s = 0.5\,mm$, which was the second lowest in terms of the registration error, and $l = 35\,mm$ were selected and used to construct a SEK which together with $\boldsymbol{R}$, was used to develop a FFD model.

4.2.2 Establishing correspondence between the FFD model and the targets

Using the FFD model, the targets and the identified landmark points, the Gaussian process registration (Gerig et al., 2014) was used to establish dense correspondence between the FFD model and the individual targets. The landmark points were used as constraints as the FFD model was rigidly aligned to the individual targets. To achieve the rigid alignment, Procrustes analysis (Stegmann & Gomez, 2002) was performed on the two sets of landmark points of the target and the FFD model and the rigid transformation that rotated and translated the first set of landmark points onto the second set was obtained. This transformation was applied to the FFD model, rigidly aligning the model to the individual targets as shown in **Figure 11**.

With the FFD model and the targets rigidly aligned, mesh points were uniformly sampled on the entire surface of the FFD model. Mesh points on the surface of the aligned targets that were closest to the sampled points on the FFD model were retrieved and used for regression where a posterior model (Albrecht et al., 2013) of the targets was generated. A posterior model was an instance of the FFD model that closely resembled an individual target. Similar to the FFD model, the posterior model still encoded deformability. While the posterior model closely resembled the individual targets, it was an average deformation from the FFD model that best resembled the targets. Hence, on a global scale, it best resembled the targets but at a local scale, there were still sections that required further deformation. Therefore, the posterior model was fit onto the targets. Using the posterior model, its respective target and the corresponding mesh points between the two, a Gaussian process regression was performed recursively until the posterior model fit as best as possible to the

targets (**Figure 12**). The Gaussian process regression was based on the mean-square metric which worked by calculating the average estimation error between the posterior model and the target. Hence, the final position of the mesh point of the fitted posterior (***FitPost***) were determined by the average estimation error (Härdle et al., 2012). The fitting ideally involved finding an instance of the posterior model that closely resembled the target being fitted. However, where the targets contain outliers, the ***FitPost*** generated by a mean –square metric estimator has sections of its surface (on the areas containing the outliers) with poor registration. To ensure the ***FitPost*** was not affected by presence of outliers, the fitting process was repeated using the robust metric estimator (Huber & Ronchetti, 1981). The robust metric estimator is resistant to errors (outliers) in the target. This is because, mesh points located far away from the rest of the mesh are truncated out. The fitted posterior (***FitPost***) was a representation of the target it fitted and was used in farther analysis. Because the same FFD model was used for the registration across the targets, all ***FitPosts*** were in correspondence.

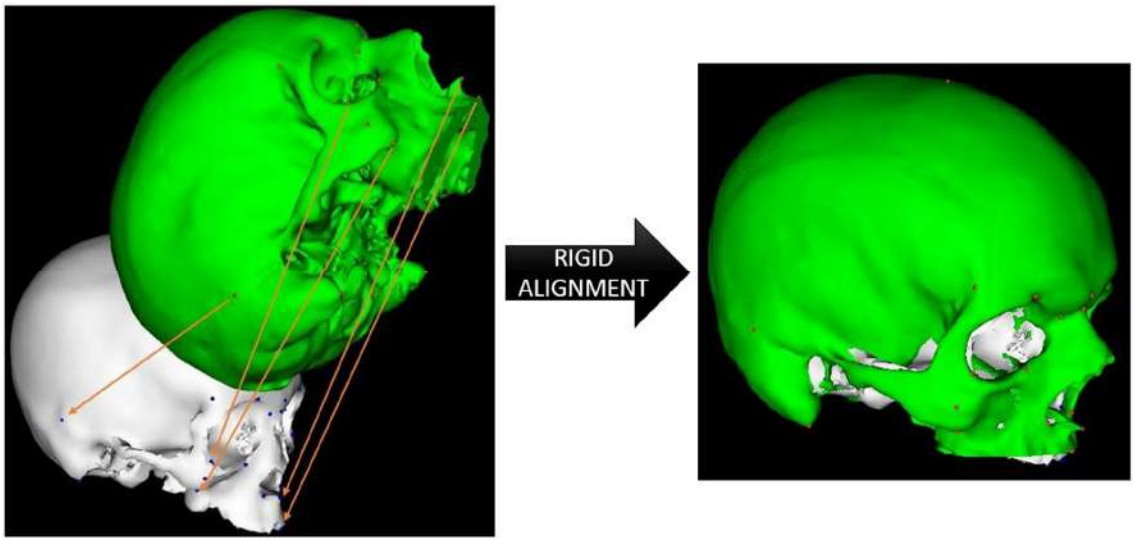

Figure 11: *Rigid alignment achieved through Procrustes analysis. The target (green) and the FFD model (white) before (left) and after (right) rigid alignment. The orange arrows on the image on the left illustrate the rigid transformation of the landmarks from the target to the model.*

After registration, ***FitPost*** is ideally used as a representation of the target. However, ***FitPost*** has to resemble the target it is representing as best as possible if one is to replace the target with the ***FitPost*** further downstream in the analysis. If not, the SSM generated from such a dataset is not representative of the targets. ***FitPost*** can be considered to be representative of target if the surface-to-surface distance in millimetres between the ***FitPost*** and its respective target is as low as possible. The surface-to-surface distance between one of the ***FitPost*** and its respective target is illustrated in the temperature map in **Figure 13**.

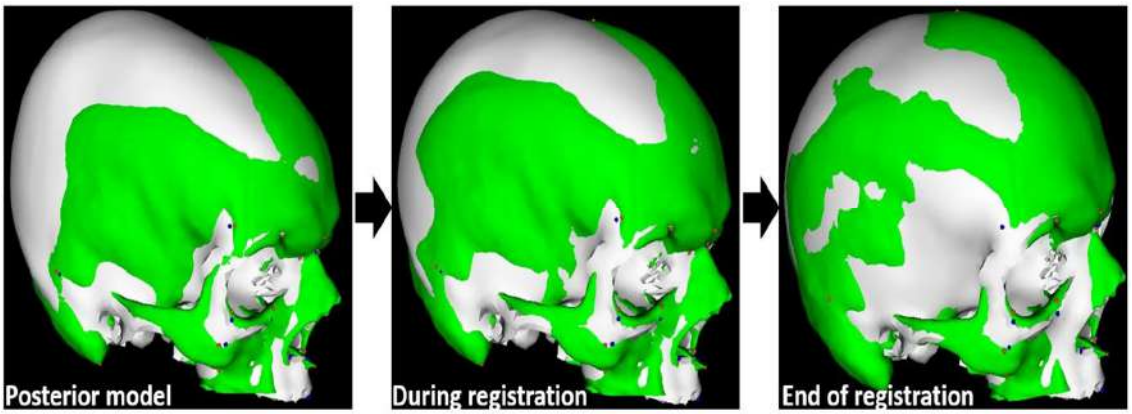

***Figure 12**: The non-rigid registration process. left: A posterior model (white) for the target (green) is generated, centre: The posterior model (white) being deformed to fit onto target (green), right: The fitted posterior–FitPost (white) closely resembling* target *(green) at the end of the registration process.*

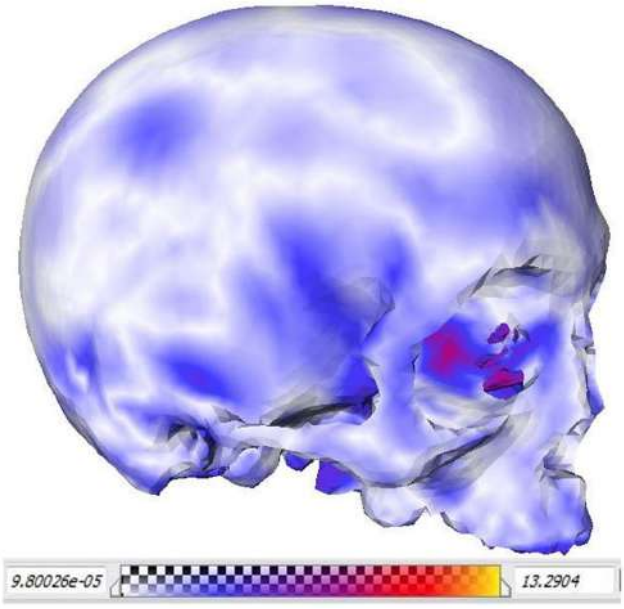

***Figure 13**: Results of the surface-to-surface distance between the first FitPost (based on the robust metric estimator) and its respective target. The mean distance is 1.86mm and the Hausdorff distance is 13.32mm.*

The surface-to-surface distance between the ***FitPosts*** and their respective targets showed that the fitting based on the robust metric estimator produced better registration results as compared to those based on the mean-square estimator. The average mean distances between the ***FitPosts*** and their respective targets were reduced in all cases but four instances where the distance remained the same or was slightly increased (**Figure 14**). Hence, the ***FitPosts*** produced using the robust metric estimator were used in further analysis.

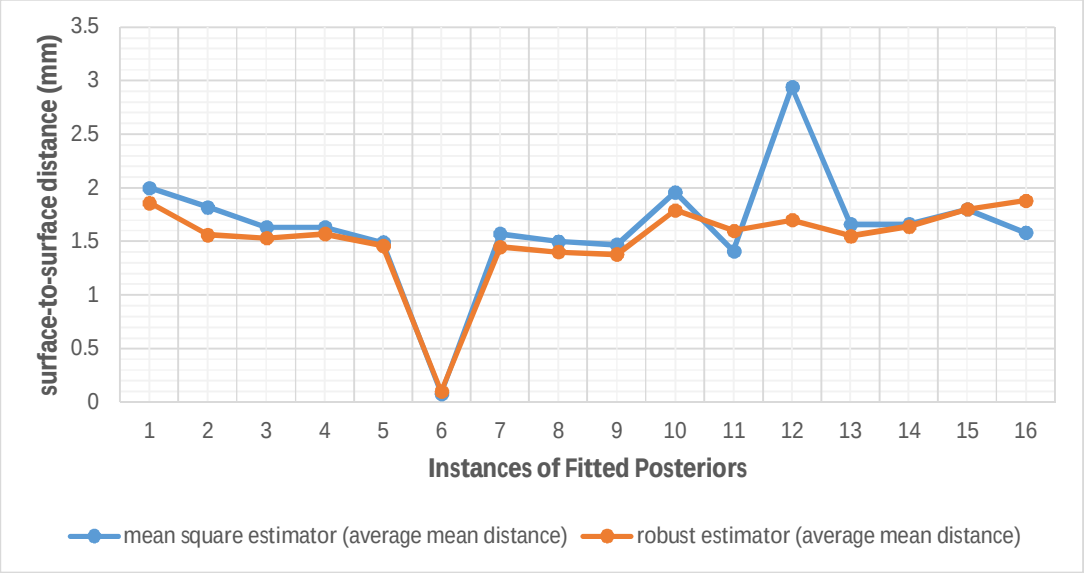

Figure 14: *A line graph illustrating the difference in the average mean distance between the fitted posteriors and their respective targets for the fitting process based on the mean square estimator versus the robust estimator.*

The surface-to-surface distance (based on the robust metric estimator) between the ***FitPosts*** and their respective targets are represented in the scatter graph in **Figure 15**.

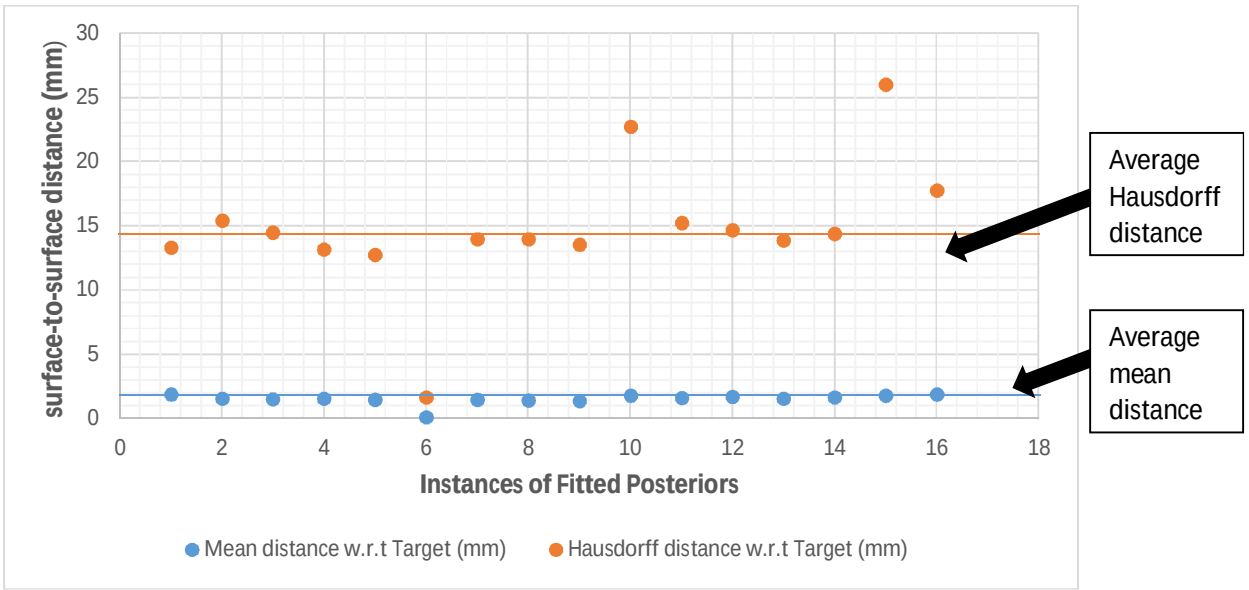

Figure 15*: A scatter plot illustrating the mean and Hausdorff surface-to-surface distance between the fitted posteriors and their respective targets together with their respective averages (solid lines).*

To avoid bias that could have been introduced because of the choice of whether to conduct the measurement from ***FitPost*** or target, the measurements were conducted twice (from both ***FitPost*** to target and from target to ***FitPost***). The mean distance was an average of the two means - measured from ***FitPost*** to target and vice versa, whereas the Hausdorff distance was the maximum value of the

two Hausdorff distances – measured from ***FitPost*** to target or vice versa. For all the ***FitPosts*** with respect to their associated targets, their associated surface-to-surface mean distance was below 2mm – the maximum deformation that anatomical structures such a bone can undergo (Anon, 2017a). This was an indication that the FFD model developed using the values of the SEK, performed a good registration. The large Hausdorff distances were due to the variations in the nasal regions between the ***FitPosts*** and their respective targets (**Figure 16**).

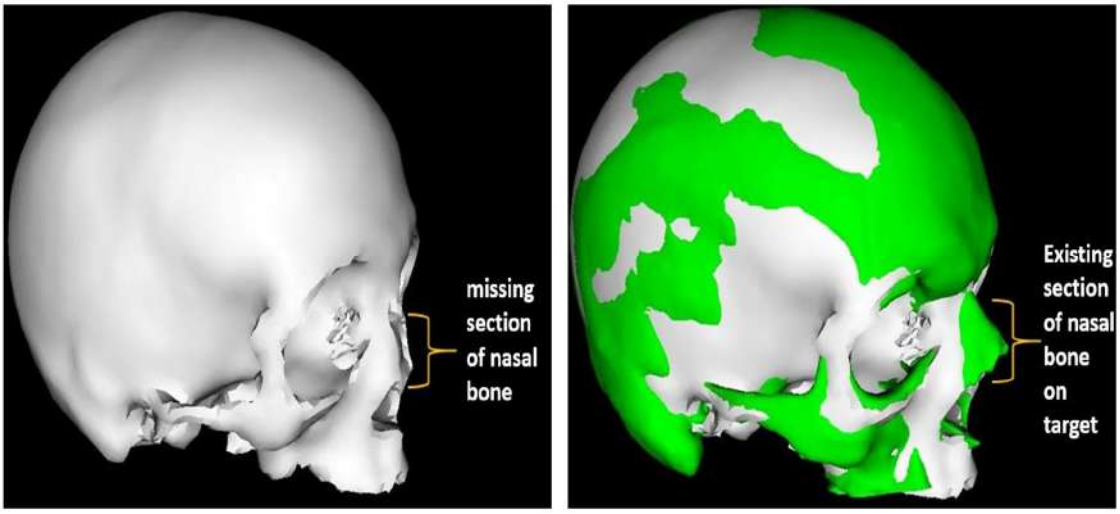

Figure 16*: Illustration of the variation in the nasal region between one of the fitted posteriors (white) and its respective target (green). Left – fitted posterior and right – fitted posterior overlaid on target.*

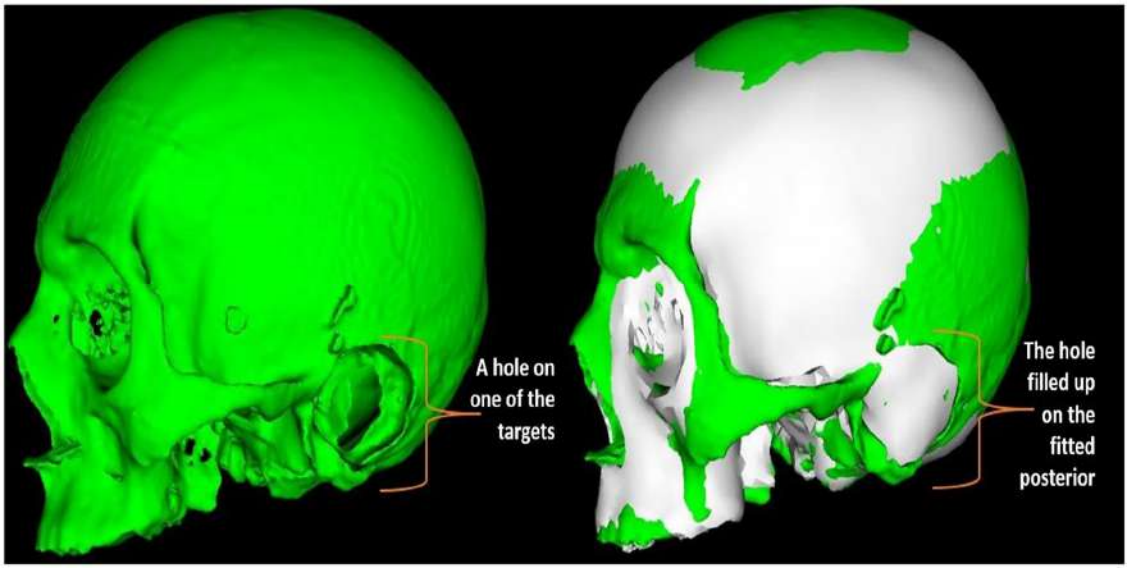

Figure 17: *Left: A hole on one of the targets -TG15 (green) right: The fitted posterior – FitPost15 (white) with the hole filled up overlaid against the target (green).*

The low values of the mean and Hausdorff surface-to-surface distance between ***FitPost***6 and its respective target was because this was the 3D surface selected in developing the FFD model. Therefore, its target was similar to the mean of the FFD model. The high value of the Hausdorff surface-to-surface distance between ***FitPost***15 and its respective target was because the target had a hole between the left temporal bone and the occipital bone which was "filled-up" during the registration process. Hence the large surface-to-surface distance in this region (**Figure 17**).

4.3 SUMMARY

In the development of the FFD model, a squared exponential kernel was used as the model's kernel. The span (l) and degree (s) of deformation values for the kernel were determined empirically. First, the value of the degree of deformation was determined based on visual observation of the deformation required to deform the mean to one of the targets (**Figure 8**). This value was then used in the determination of the value of the span of the deformation. Hence, the two values were correlated therefore influencing the integrity of the values obtained at the end. However, using the FFD model built from the two values for registration of the targets, the results were not adversely affected by the process used in determining the two values. The registration errors (in terms of mean distance) were small. The mean registration errors were below 2mm - a good result for a complex shape such as the crania. The large Hausdorff distance errors were not due to the deformability of the FFD model but rather the constraints of the features of the reference shape. The nasal region between the reference shape and the targets varied hence the large Hausdorff distances. The ***FitPosts*** were used in developing a skull SSM as documented in Chapter 5.

The registration process is summarised in **Figure 18**.

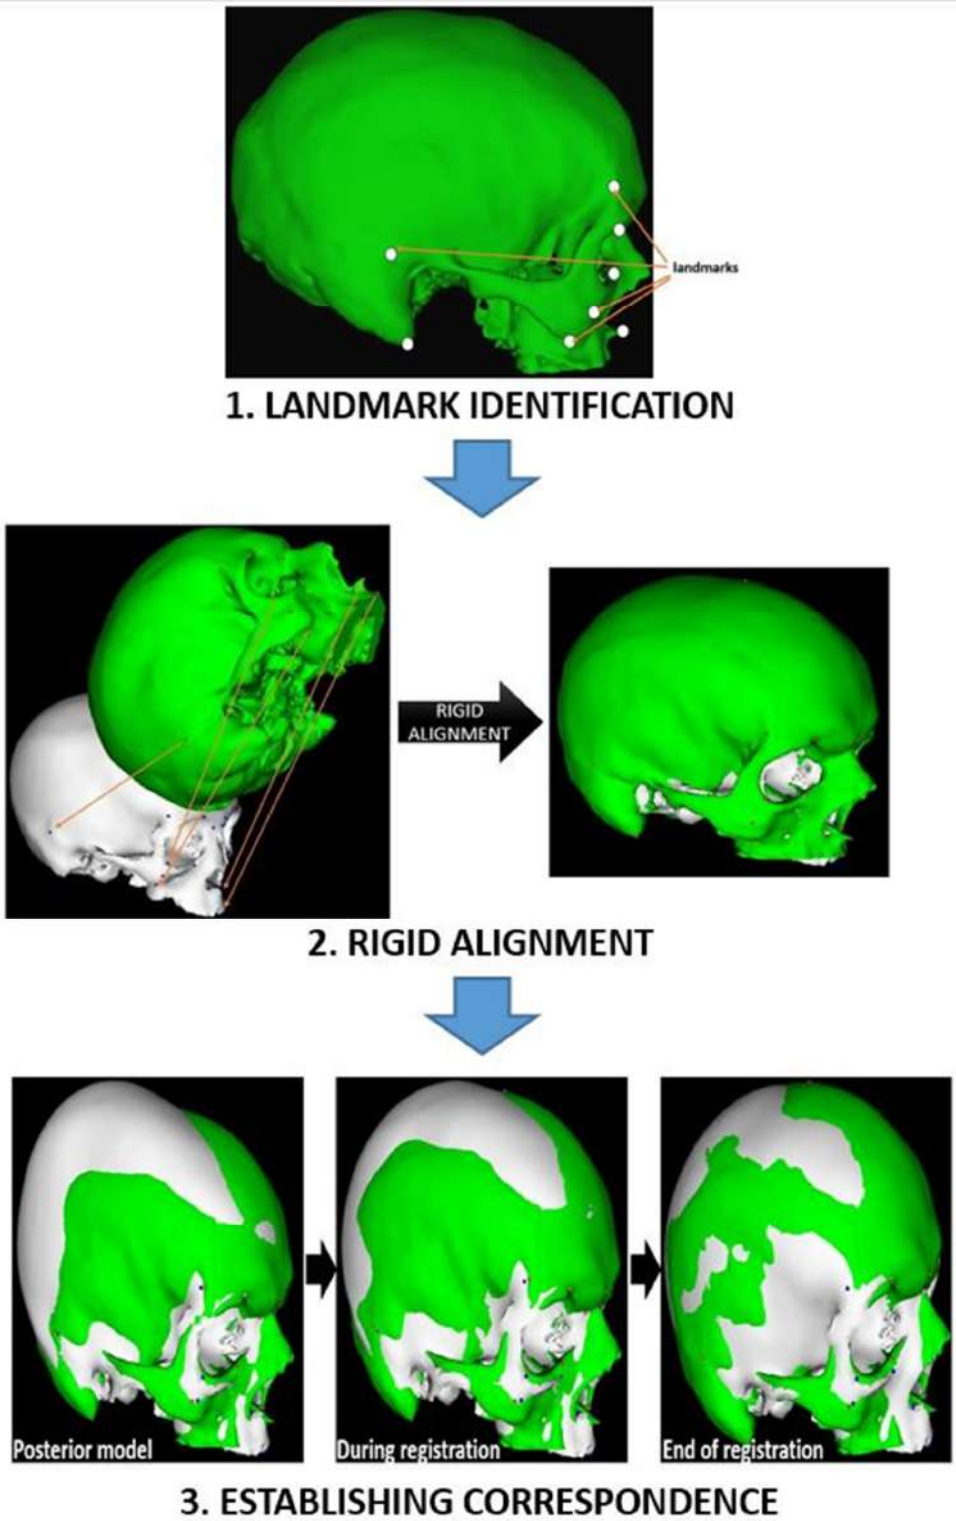

***Figure 18**: Summary of the registration process.*

5 DEVELOPING A SKULL SSM USING IN-CORRESPONDENCE DATA

Statistical shape models describe shape variation across a dataset belonging to the same shape family e.g. faces, hands or skulls. This is achieved by applying principal component analysis (PCA) to the covariance matrix of the in-correspondence data. The PCA captures the mean shape and modes of variation from the mean shape in the dataset. This chapter details the development of a skull SSM from in-correspondence data in the form of fitted posteriors (***FitPosts***) which were generated in the previous chapter. In-correspondence data mean that the ***FitPosts*** have the same number of nodes and that the nodes represent the same anatomical features on the skull. However, before the PCA could be applied on the ***FitPosts***, they needed to be in the same coordinate system. Statistical shape models describe shape variation only and effects of rotation and translation in the in-correspondence data need to be eliminated first. The first section of this chapter discusses the rigid alignment of the ***FitPosts***. The second section describes the skull SSM building process, followed by the model visualization and validation. The final section summarise key points identified in the process. The Scalismo software (Anon, 2017b) was used in the processes of rigid alignment, SSM development, visualization and validation - all of which were conducted on an Intel (R) Core (TM) i7-4790 CPU @3.60GHz running a 64-bit operating system on a 32 GB RAM.

5.1 ALIGNMENT OF THE IN-CORRESPONDENCE DATA (***FitPosts***)

The registration process in Chapter 4 was implemented by first rigidly aligning the FFD model to the different targets before a registration was conducted. Hence, even though the ***FitPosts*** were in-correspondence with one another, they were in different coordinate systems. Rigid alignment of the ***FitPosts*** was required to eliminate variations due to relative spatial displacements of the ***FitPosts*** from both rotation and translation. This is because SSMs only capture a variation in shape. Aligned in-correspondence data are in the same coordinate system and applying a PCA on such a dataset ensures the process will capture corresponding points across the dataset and calculate only the variation in shape.

Generalized Procrustes alignment was used to establish rigid alignment across the ***FitPosts***. This was achieved by selecting one of the ***FitPosts*** as a reference, $\boldsymbol{FitPosts_{ref1}}$ for the purpose of alignment. A set of corresponding points between $\boldsymbol{FitPosts_{ref1}}$ and each of the remaining fifteen ***FitPosts*** was selected. Each of the fifteen ***FitPosts*** and their corresponding points with $\boldsymbol{FitPosts_{ref1}}$ were used to align the ***FitPosts*** to $\boldsymbol{FitPosts_{ref1}}$. An approximate mean shape from the aligned ***FitPosts*** was calculated and used as a reference ($\boldsymbol{FitPosts_{ref2}}$) for the next iteration of the alignment process.

Several iterations of the alignment process were done until subsequent estimates of the mean shapes were not changing (**Figure 19**). The resulting aligned ***FitPosts (FitPosts$_{Alg}$)*** were used in the skull SSM building process as described in Section 5.2.

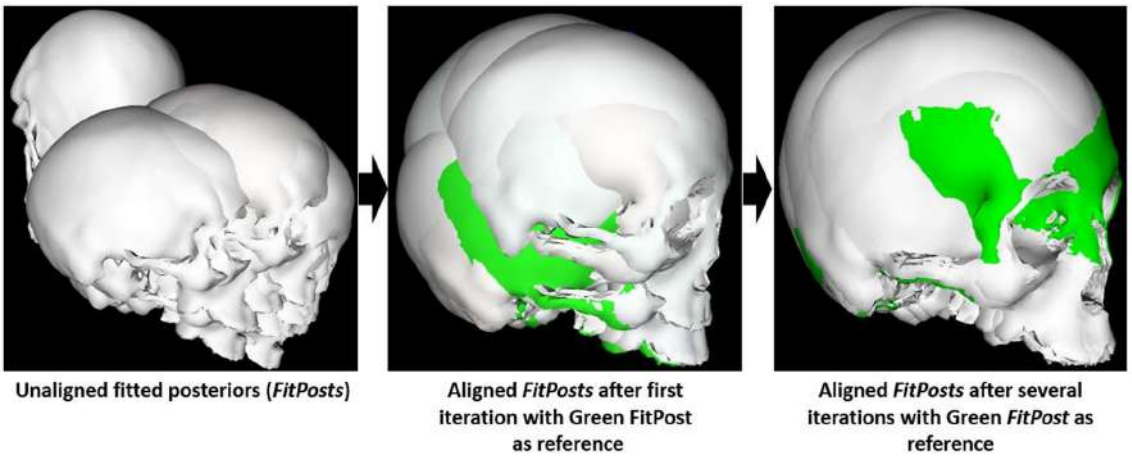

***Figure 19**: The alignment process using GPA. Left: unaligned fitted posteriors. Centre: aligned fitted posteriors after one iteration with the green fitted posterior as the reference. Right: aligned fitted posterior after the final iterations with green fitted posterior as reference.*

5.2 SKULL SSM BUILDING PROCESS

This section outlines the statistical analysis as carried out on ***FitPosts$_{Alg}$*** generated in section 5.1. The statistical analysis was conducted using the Gaussian process and PCA to capture the mean shape and principal modes of variation of the ***FitPosts$_{Alg}$***. In this section the mean shape of ***FitPosts$_{Alg}$*** is referred to as reference ***FitPosts$_{Alg}$*** i.e. ***FitPosts$_{AlgREF}$*** whereas the remaining fifteen ***FitPosts$_{Alg}$*** are still referred to as ***FitPosts$_{Alg}$***.

The ***FitPosts$_{Alg}$***, were meshes in-correspondence with each other. These ***FitPosts$_{Alg}$*** were converted into deformation fields by calculating the difference between each mesh point of the individual ***FitPosts$_{Alg}$*** with the respective mesh points on ***FitPosts$_{AlgREF}$***. The deformation fields of the fifteen ***FitPosts$_{Alg}$*** were converted into discrete vector fields. The discrete vector fields were converted into continuous vector fields by interpolating the regions between the values to cater for regions that did not have mesh points. Principal component analysis was performed on the continuous vector fields defined over the reference to generate a discrete low rank Gaussian process. The low rank Gaussian process encoded the normal distribution of the vector fields with a mean and a covariance function. The mean of the Gaussian process was the mean of all the deformations from ***FitPosts$_{Alg}$***. The covariance function was a matrix generated from the evaluation of the covariance between the deformation vectors of ***FitPosts$_{Alg}$*** at corresponding points on ***FitPosts$_{AlgREF}$***. The Gaussian process could yield deformation fields and these fields were warped onto ***FitPosts$_{AlgREF}$*** thereby building a

skull SSM. The skull SSM was composed of a mean shape and fifteen modes of variation from the mean shape. **Figure 20** illustrates the mean shape and the first three main modes of variation.

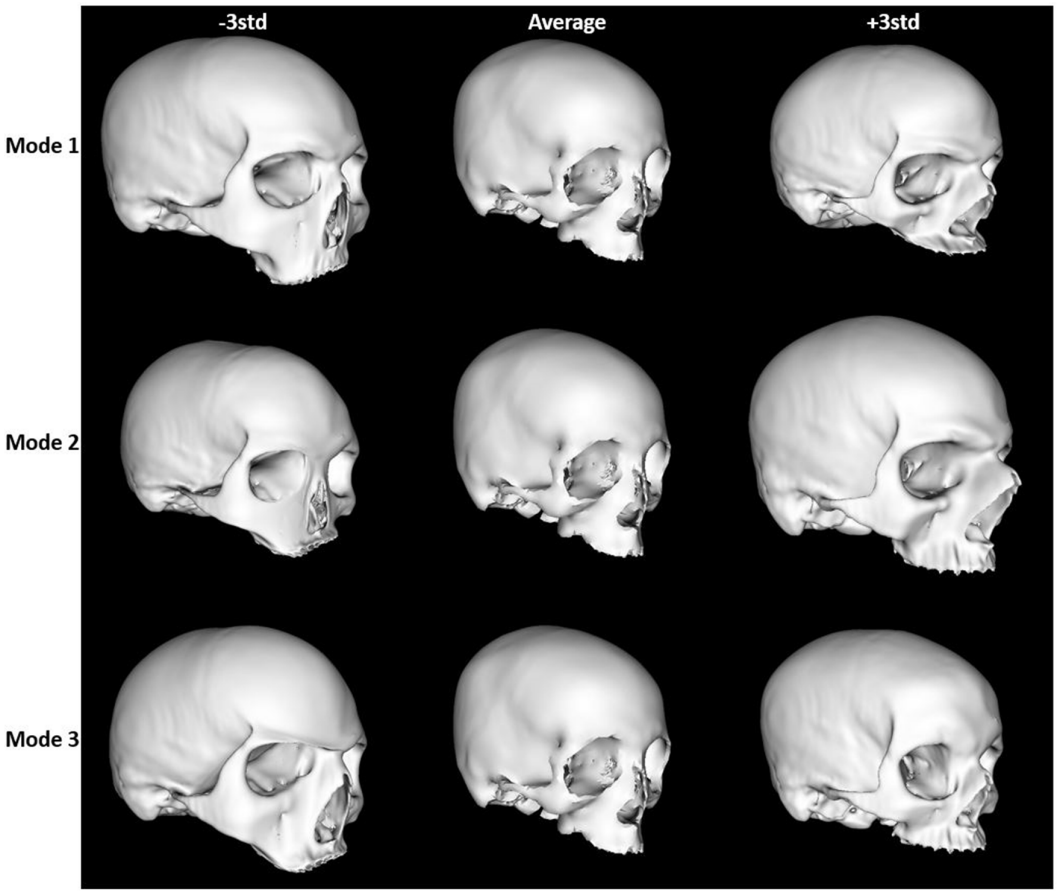

Figure 20: *The main modes of variation for the first three principal components of the skull SSM. First row = mode 1, second row = mode 2 and third row = mode 3. First column = -3 standard deviations, second column = average shape and third column = +3 standard deviations.*

The Mode 1 of the model represented a variation in the size of the skull. The size variation is illustrated across the figures in row 1 of **Figure 20**. Across the three skulls, the size variation is around the cheek bone region of the zygomatic and maxilla bones and the temporal process. The Mode 2 also represented some smaller size-related variation and variation around the optical region of the nasal, ethmoid and lacrimal bones (second row of **Figure 20**). Mode 3 represented variations around the frontal, parietal, occipital and temporal bones. The developed skull SSM was validated as outlined in Section 5.3 before it could be used further in the project or for other applications.

5.3 VALIDATING THE SKULL SSM

The skull SSM could be used for several medical applications such as in model-based segmentation of the skull, aid in the design of patient-specific implants of the skull or in the

development of SAMs and finite element analysis of the skull. However, before the model could be used for any application, it had to be validated to ensure that the instances generated from the model were representative of the shape of human skulls. The SSM of the skull was validated by conducting various tests on it. These tests included calculating the compactness, specificity and generality (using the leave-one-out cross validation process) of the model. The compactness test measured the amount of variation in the training dataset as explained by each mode (equation [5.1]) (Kohavi, 1995).

$$\boldsymbol{C(M) = \sum_{m=1}^{M} \lambda^{m}} \quad \textbf{[5.1]}$$

where λ^{m} is the m^{th} largest eigenvalue and $C(M)$ is the cumulative variance of the M^{th} mode.

The specificity test, $S(R)$ measured the ability of the model to produce instances that are valid in their resemblance to other shapes in the training dataset (equation [5.2]) (Kohavi, 1995).

$$\boldsymbol{S(R) = \frac{1}{N} \sum_{I=1}^{N} |srandom_j(R) - s'_j|^2} \quad \textbf{[5.2]}$$

where $|srandom_j(R)$ are randomly generated shape examples using R principal modes and s'_j are the nearest example of the training dataset to the $srandom_j(R)$.

The generality test, $G(R)$ measured the ability of the model to represent new shapes belonging to the class object of the training dataset (equation [5.3]) (Kohavi, 1995).

$$\boldsymbol{G(R) = \frac{1}{N} \sum_{I=1}^{N} |S'_i(R) - S_i|^2} \quad \textbf{[5.3]}$$

where N is the number of instances, $S'_i(R)$ is the best model of reconstruction of S_i excluded from the model built with R principal components.

5.4 VALIDATION RESULTS

From the validation process, the first 8 principal components of the SSM represented 92.13% of the variance (**Figure 21**). Thus the SSM only required half the number of principal components to define an instance as described by Styner et al. (2003). The SSM was able to generate instances of the crania that were similar to those in the training dataset based on the specificity of thirteen instances generated by the SSM, which had a steady state value of approximately 2.04mm (**Figure 22**). Hence the average distance between the instances and their respective nearest member in the training dataset was at most approximately 2 mm. The SSM was able to effectively and accurately represent unseen instances of the crania through the leave-one-out method, where fifteen registered targets were left out of the SSM at different times and the SSM built was used to fit onto these targets. The approximation error (for the left-out registered targets) averaged over the number of trials i.e. the generality of the model had a steady state value of 1.48mm (**Figure 23**). The model effectively generalised onto unseen targets with approximation error of below 2mm across the fifteen targets.

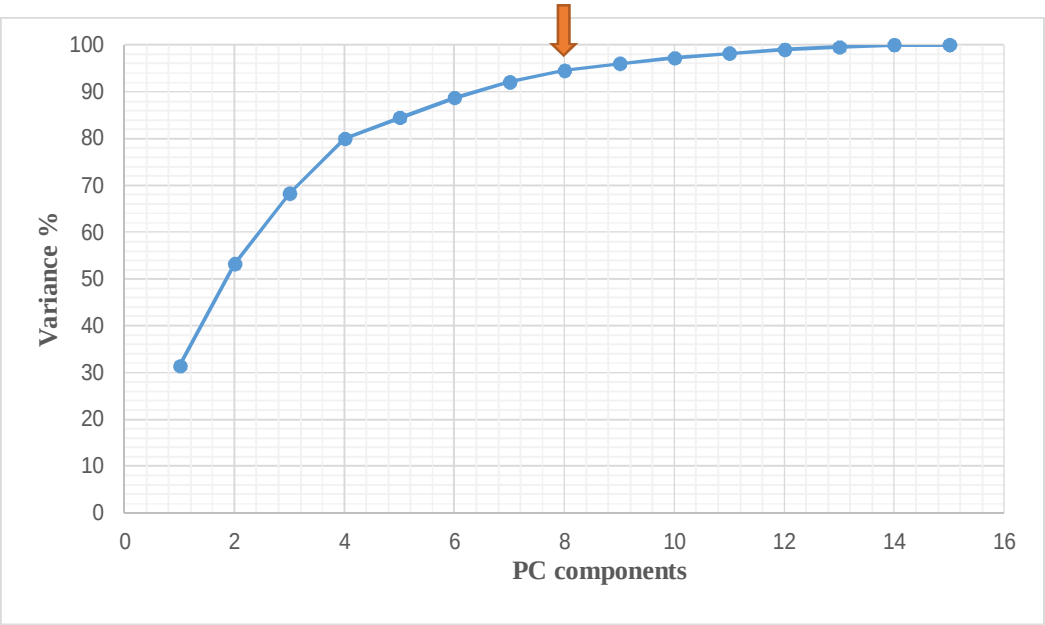

***Figure 21**: Compactness of the SSM (orange arrow points to the region where the first principal components of the SSM represents 92.13% of its total variance).*

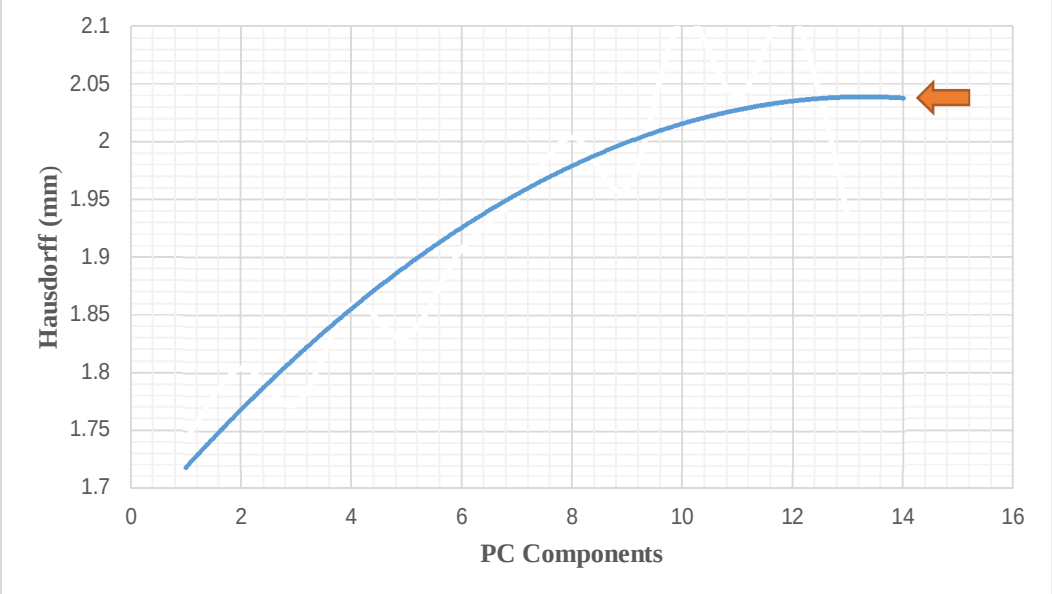

***Figure 22:** Specificity of the SSM. Note the steady state value of 2.04mm.*

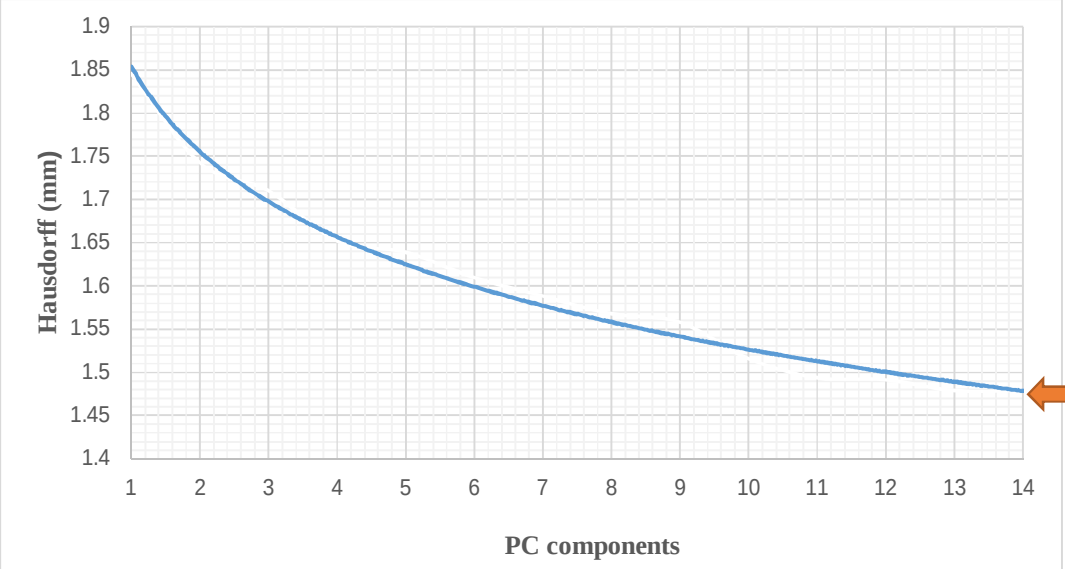

Figure 23: *Generality of the SSM. Note the steady state value of 1.48mm.*

5.5 SUMMARY

This chapter discussed the development and validation of the skull SSM from in-correspondence data. During the GPA, a random fitted posterior was initially selected as the reference ($\boldsymbol{FitPosts}_{\boldsymbol{ref1}}$) upon which the initial alignment iteration was based on. However, GPA overcomes this reference selection bias by subsequently selecting difference references for subsequent alignment iterations. Using an Intel (R) Core (TM) i7-4790 CPU @3.60GHz running a 64-bit operating system on a 32 GB RAM, the alignment and model building process took tens of seconds. If the dataset was large – in the 100s or if the computer used had less processing power, the process might have taken longer. The development of an initialization protocol that efficiently selects an initial reference which reduces the number of iterations in the alignment process and hence reduces the time taken for development of a SSM from in-correspondence data, would speed up processing.

The SSM generated from the in-correspondence, aligned data $\boldsymbol{FitPosts}_{\boldsymbol{Alg}}$ had a few limitations. The SSM was developed from a small sample size that would have affected its generality. This is because the principal modes of variation that defined the SSM were restricted to the features encoded on only the 16 samples ($\boldsymbol{FitPosts}_{\boldsymbol{Alg}}$) used.

The first principal mode of variation of the SSM represented a variation in the size of the skull. The variation was around the cheek bone region of the zygomatic and maxilla bones and the temporal process. All these variations are expected in a dataset of crania obtained across any population. The second principal mode of variation also represented smaller size-related variations and variations around the optical region of the nasal, ethmoid and lacrimal bones. These variations

were too wide to expect from a dataset obtained from a single population (similar ancestry). However, because the dataset used contained specimen with Caucasian, African and mixed-race ancestry such large variations around the optical region of the nasal, ethmoid and lacrimal bone were expected. A skull of Caucasian ancestry has a high projecting nasal bone with rectangular shaped orbits, while a skull of African ancestry has a broad nasal aperture with rounded nasal sills (Weinberg et al., 2005). The third mode of variation represented variations around the frontal, parietal, occipital and temporal bones. The validated SSM was used in the development of a SAM as detailed in Chapter 6.

The process of SSM development from in-correspondence meshes is summarised in **Figure 24**.

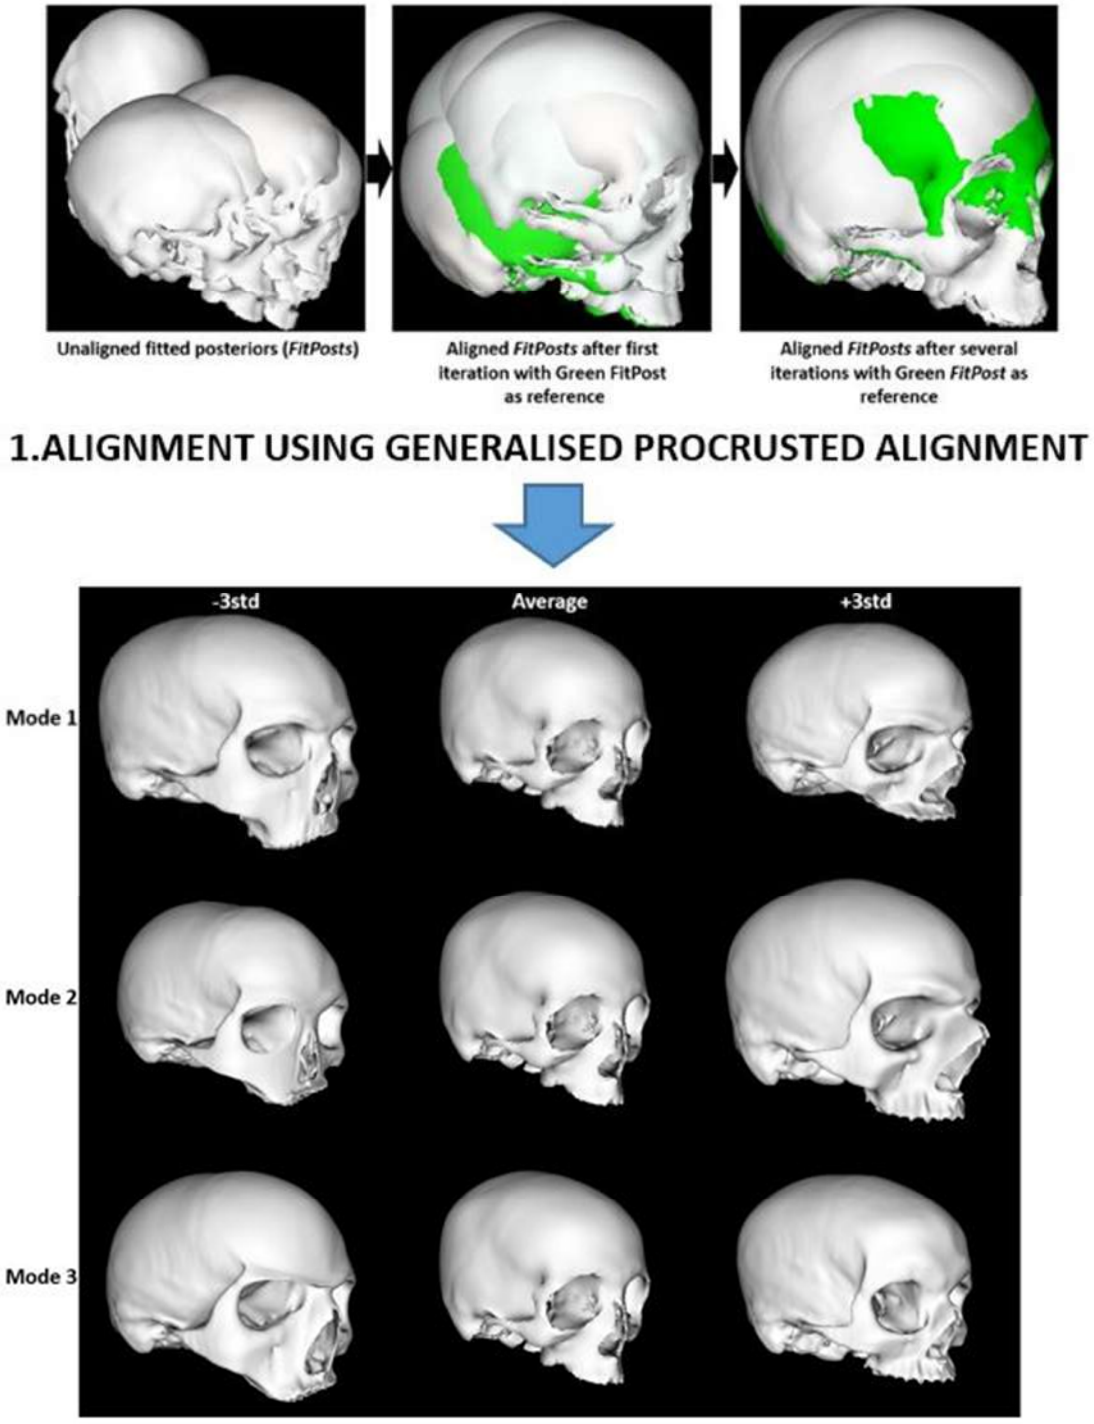

***Figure 24**: A summary of the process of SSM development from in-correspondence meshes.*

6 DEVELOPING A MESH-BASED SAM OF THE SKULL

Statistical appearance models are an extension of SSMs and they describe density and variation of density within a set of anatomical structures of a population. Statistical appearance models can be developed using either image-based or mesh-based processes. Image-based SAMs are developed directly from volumetric images such as CT images whereas mesh-based SAMs are developed from meshes hence they can be developed from SSMs. This chapter details the development of a skull SAM from the skull SSM developed in the previous chapter. The first section of this chapter documents the development of in-correspondence volumetric meshes of the crania. The second section details the development of a skull SAM using the in-correspondence volumetric meshes and CT images. The third section discusses the validation of the model followed by the final section which summarises key points and observations from the process. The Amira v6.2.0 software (Fei Imaging- www.fei.com) was used in developing a volumetric reference mesh and Scalismo software (Anon, 2017b) was used in establishing volumetric mesh correspondence and SAM development and validation – all of which were conducted on an Intel (R) Core (TM) i7-4790 CPU @3.60GHz running a 64-bit operating system on a 32 GB RAM.

6.1 DEVELOPING IN-CORRESPONDENCE VOLUMETRIC MESHES

To develop a mesh-based SAM, in-correspondence volumetric meshes are used to capture intensity values located at the nodes of the meshes on corresponding CT images. Principal component analysis is applied to this group of intensities to develop a SAM. This section details the development of in-correspondence volumetric meshes from the mean of the SSM ($\boldsymbol{SSM_{mean}}$) developed in Chapter 5 and fitted posteriors (***FitPosts***) developed in Chapter 4.

Developing in-correspondence volumetric meshes requires registration of ***FitPosts*** with a volumetric reference mesh. Unlike a reference developed from a surface mesh, a volumetric reference contains a large number of features which may limit its deformability. The restricted deformability may affect how the volumetric mesh fits across a dataset of ***FitPosts***. Using a volumetric reference that is approximately a mean of the dataset being registered reduces the degree of deformation the volumetric reference is required to undergo to fit onto each of the dataset; this addresses the deformability issue. In this project, $\boldsymbol{SSM_{mean}}$ was used to develop the volumetric reference. Before $\boldsymbol{SSM_{mean}}$ could be used to develop the volumetric reference, it was accessed to ensure the volumetric reference to be developed from it did not have intersecting meshes and tetrahedrons, poor orientations or out-of-bound aspect ratios (Ruppert, 1995). If a volumetric mesh has the mentioned defects, some of its nodes may be located in the same position. If such a mesh is used to sample intensities from a CT image, the sampling is un-even and at times from the same point, affecting the integrity of the SAM developed from such a dataset.

6.1.1 Testing SSM_{mean} against intersection, poor orientation, and out-of-bound aspect ratio

A surface mesh whose triangles are intersecting, poorly-oriented or have out-of-bound aspect ratios, generates a volumetric mesh with similar defects. It is paramount that only surface meshes without these defects are used to develop a volumetric mesh (Guide, 2009). SSM_{mean} was tested to ensure the triangles that made up its surface did not have the mentioned defects. The triangles were assessed with regard to intersections, orientation and aspect ratio. An orientation test was done to ensure that all surface triangles were properly oriented with respect to the exterior and interior materials of the mesh. The aspect ratio was computed as the ratio of the radius of the circum-circle to that of the in-circle for each triangle of the surface mesh. Large aspect ratio values (above 20) indicate irregular triangles which are prone to producing intersecting tetrahedrons (Guide, 2009). SSM_{mean} passed the intersection test. However, some of the surface triangles had poor orientation and large aspect ratio values and manual repair was performed on SSM_{mean}. The manual repair involved translating the vertices of the triangles with large aspect ratio to reduce the aspect ratio as illustrated in **Figure 25**. The repaired SSM_{mean} was composed of 20004 vertices (nodes) and 39902 nearly isotropic triangles. This repaired SSM_{mean} was used to develop a volumetric mesh (Section 6.1.2).

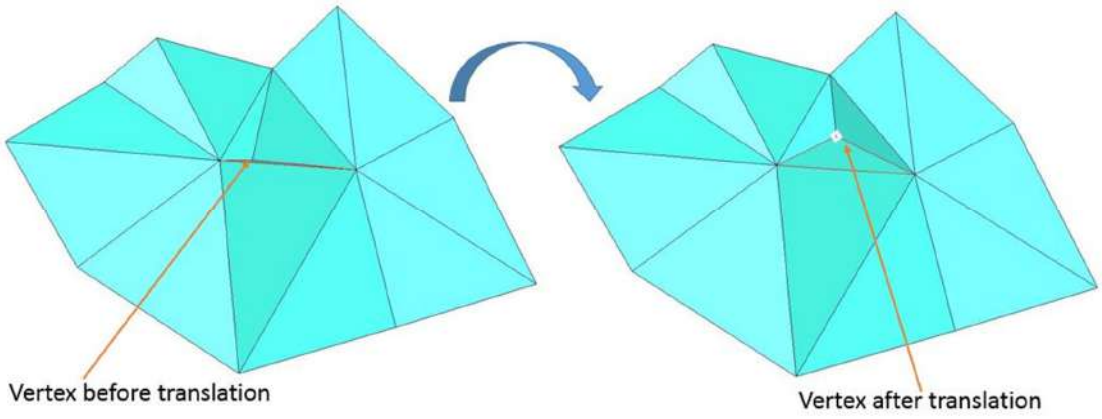

Figure 25: *Editing of a triangle with a large-aspect ratio using the translate vertex process to reduce the aspect ratio value and improve the quality of SSMmean.*

6.1.2 Developing a volumetric mesh from the repaired SSM_{mean}

The repaired SSM_{mean}- a surface mesh - was used to generate a volumetric mesh and establish volumetric correspondence across ***FitPosts***. The Tetra Grid generator in Amira which utilises the Delaunay triangulation technique (George & Borouchaki, 1998) was used to generate a volumetric mesh from the repaired SSM_{mean}. The Tetra Grid generator first tested the repaired SSM_{mean} against intersection, orientation, aspect ratio and dihedral angle tests. The repaired SSM_{mean} passed all these tests. The generator then created links between the outer and inner surfaces of the repaired SSM_{mean} according to the size of the triangles on the surfaces (**Figure 26**). A surface mesh has no internal

triangles/ tetrahedrons between the surface triangles whereas a volumetric mesh does (**Figure 27**). The links joined nodes on the opposite surfaces which were either on the same horizontal, vertical or diagonal plane. Where there was no common plane linking two nodes together, a node was created between two surface nodes and a link created between the two via the node. The collection of surface and internal triangles and the surface and internal nodes created a tetrahedral volumetric mesh of the crania. The volumetric mesh was used to develop a FFD volumetric model (Section 6.1.3).

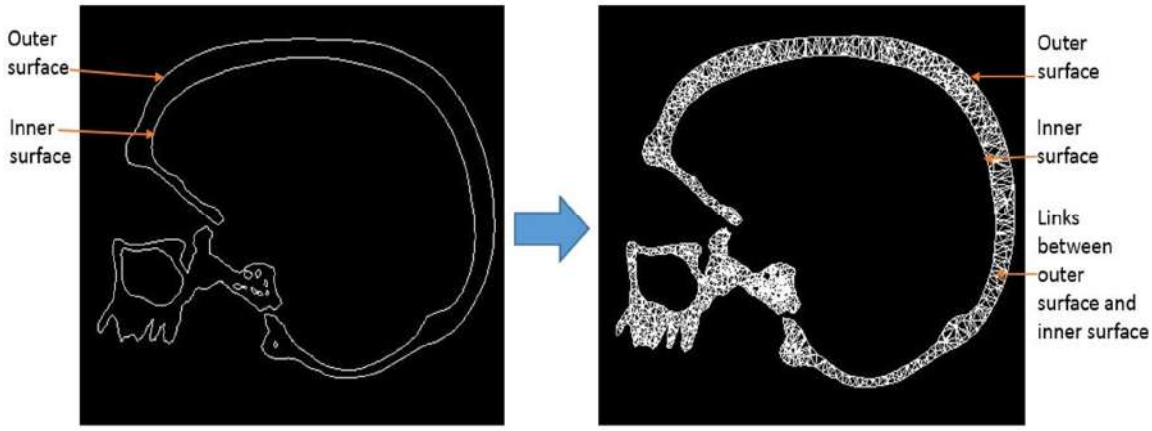

Figure 26: *Cross-section of a cranium showing the process of developing a volumetric mesh (right) from a surface mesh (left). The Tetra Grid generator created links between the nodes on the surface mesh creating triangles between the surfaces.*

Figure 27: *Cross-section of a part of the cranium showing absence of tetrahedrons (left) and presence of tetrahedrons (right) within a surface (left) and volumetric (right) mesh of the cranium.*

6.1.3 Developing a FFD volumetric model from a volumetric mesh

The volumetric mesh developed in section 6.1.2 was to be used to establish dense correspondence across ***FitPosts***. The volumetric mesh would be non-rigidly deformed to fit onto each

of the ***FitPosts***. However, because the volumetric mesh was a rigid body, deformability had to be introduced by developing a FFD model from the volumetric mesh.

As described in Section 4.2.1, FFD models are used to model deformations of rigid objects. A FFD volumetric model was developed using the Gaussian process in Scalismo. In the GP process, a FFD model is developed from a mean mesh and a kernel. The mean mesh of the FFD volumetric model was the volumetric mesh developed in Section 6.1.2. Unlike the use of only one squared exponential kernel (SEK) for the FFD model developed in Section 4.2.1, a combination of three SEKs was used for the FFD volumetric model. The volumetric mesh had a larger number of nodes and triangles unlike the mean mesh used in section 4.2.1. This increased number of features, together with the presence of tetrahedrons, had increased the rigidity of the volumetric mesh. Three SEKs were used to ensure that the kernel (equation 6.1) used could induce sufficient deformation to the volumetric mesh, enough to enable the volumetric mesh fit onto the ***FitPosts***.

$$k_{SE}(x, x') = S \exp\left(\frac{-(x,x')^2}{l^2}\right) \quad \textbf{[6.1]}$$

The first SEK was developed with the values of the span (*l)* and the degree (*s)* of the kernel as 70mm and 100mm respectively, to cater for global deformations (for details on *l* and *s*, refer to Section 4.2.1). Global deformations refer to deformations where large values of the span and degree of deformation affect large sections of the object being deformed (**Figure 28**). Medium and local deformations refer to deformations where the values of the span and degree of deformation affect medium and smaller sections of the object being deformed, respectively. The span and degree of deformation for the second and third SEKs were 30mm and 50mm and, 0.5mm and 35mm respectively (**Figure 29** and **Figure 30**). The three kernels were combined to form a master kernel (equation 6.2) which added global, medium and localised deformation to the developed FFD volumetric mesh as illustrated in **Figure 31**.

$$k_{MK}(x, x') = \left[S \exp\left(\frac{-(x,x')^2}{l^2}\right)\right] + \left[S \exp\left(\frac{-(x,x')^2}{l^2}\right)\right] + \left[S \exp\left(\frac{-(x,x')^2}{l^2}\right)\right] \quad \textbf{[6.2]}$$

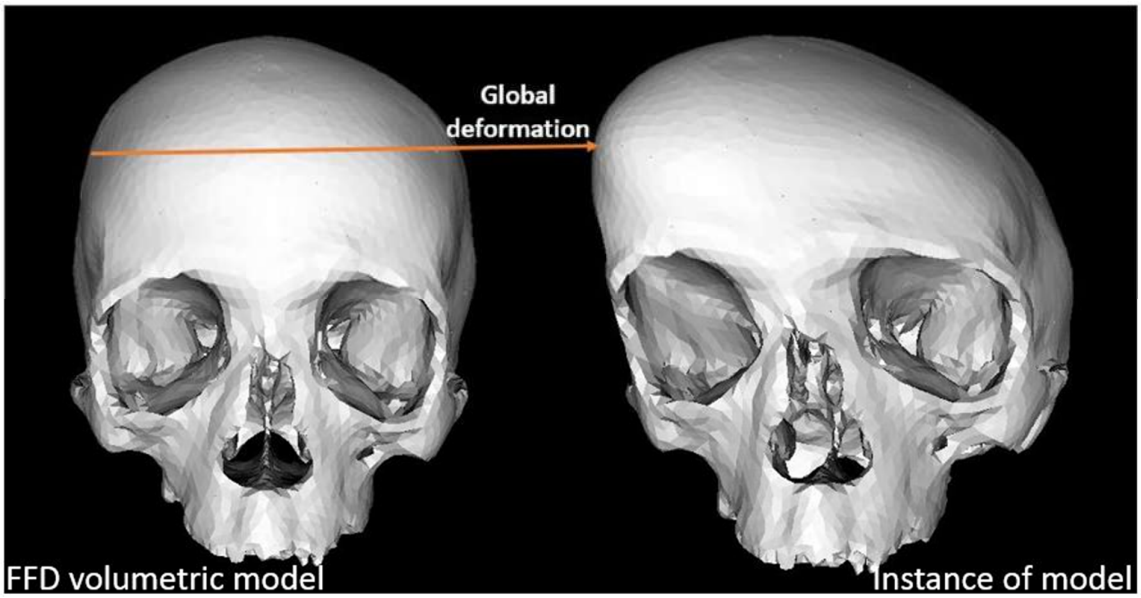

Figure 28: *An instance (right) generated from a FFD volumetric model (left) illustrating global deformation caused by large values of span (l) and degree (s) of deformation of the SEK.*

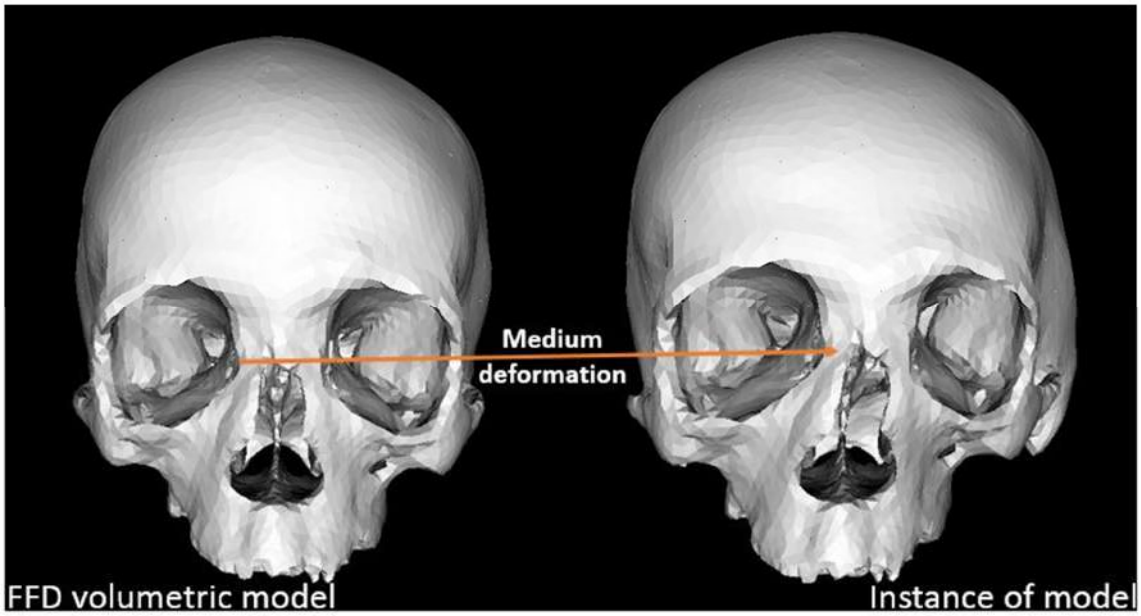

Figure 29: *An instance (right) generated from a FFD volumetric model (left) illustrating medium deformation caused by moderately small values of span (l) and degree (s) of deformation of the SEK.*

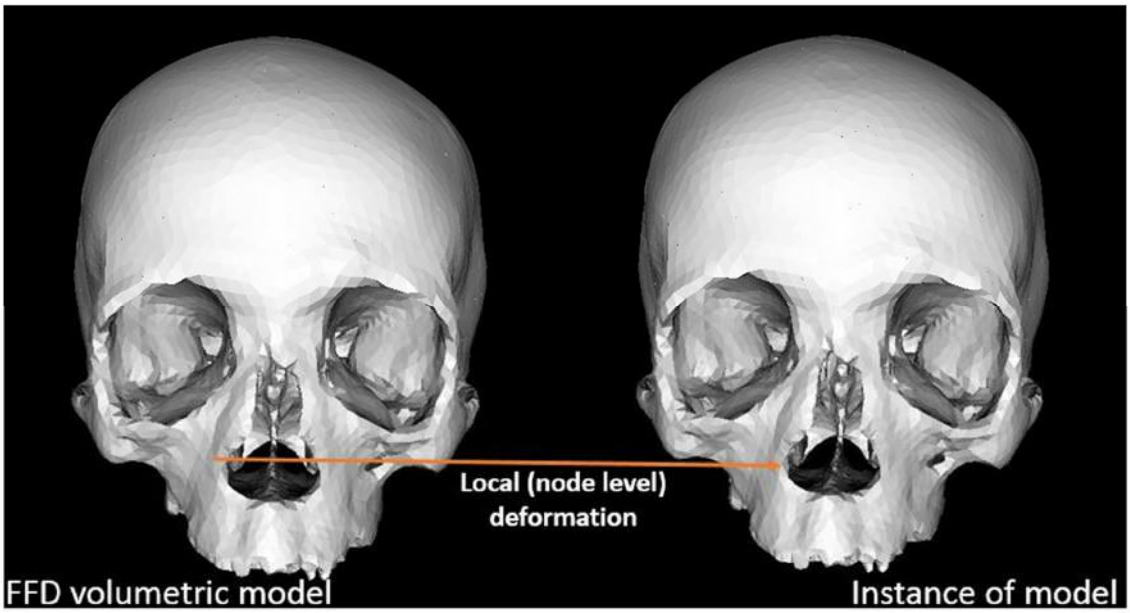

***Figure 30**: An instance (right) generated from a FFD volumetric model (left) illustrating localised deformation caused by very small values of span (l) and degree (s) of deformation of the SEK.*

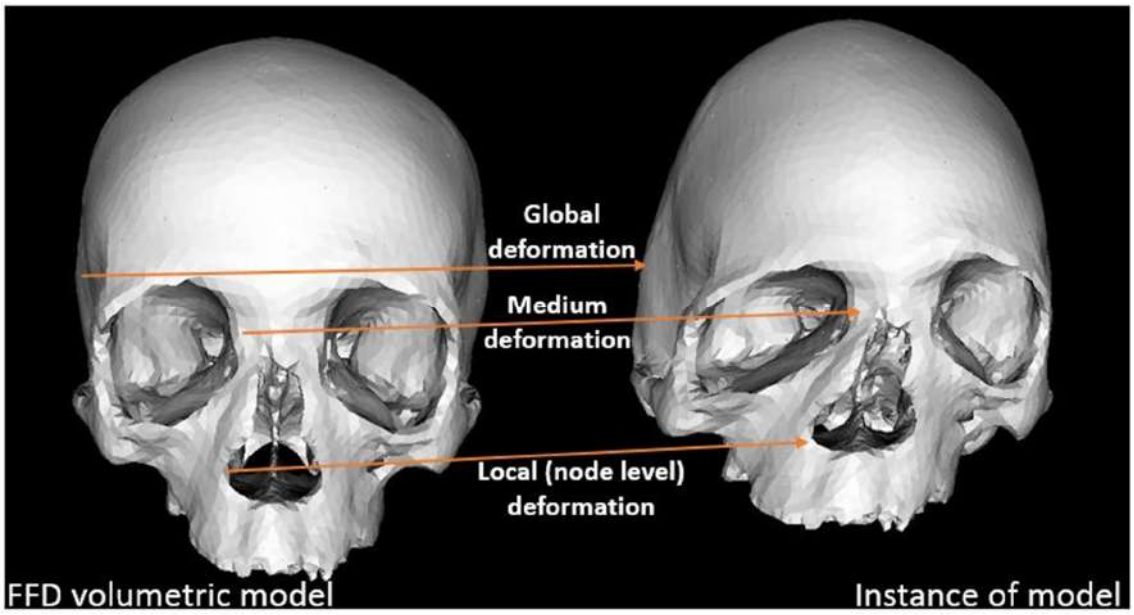

***Figure 31:** An instance (right) generated from a FFD volumetric model (left) illustrating global, medium and localised deformations caused by a combination of three kernels whose values of span and degree of deformation of the SEK vary from small to large values.*

The FFD volumetric model developed using the three kernels was used to establish volumetric correspondence across the surface meshes as detailed in section 6.1.4.

6.1.4 Establishing volumetric correspondence across surface meshes/ fitted posteriors (***FitPosts***)

A similar process to the one used to establish correspondence across the targets in section 4.2.2 was used to establish volumetric correspondence across ***FitPosts***. The FFD volumetric model, the ***FitPosts*** and corresponding landmark points were used in the Gaussian process registration to establish correspondence. Volumetric correspondence was required to ensure the intensity values captured from CT images by the various 3D surfaces were from corresponding points. Only intensities from corresponding points are used to develop a SAM.

As detailed in Section 4.2.2, the non-rigid registration process adopted here requires features to constrain the deformation to a predefined shape. During the generation of the ***FitPosts*** in Section 4.2.2, the landmarks on the FFD model were translated to get new landmark points that corresponded to expected landmarks on ***FitPosts***. These landmarks points were used in this section to constraint the FFD volumetric model to the individual ***FitPosts*** during the registration process. Using Procrustes analysis, the FFD volumetric model was rigidly aligned to each of the individual ***FitPosts*** using 33 corresponding anatomical landmarks (section 4.1) between the model and the ***FitPosts***. After the alignment, mesh points were uniformly sampled on the entire surface of the FFD volumetric mesh. Mesh points on the ***FitPosts*** closest to the sampled points on the FFD model were retrieved and used for regression where posterior models of individual ***FitPosts*** were generated (**Figure 32**).

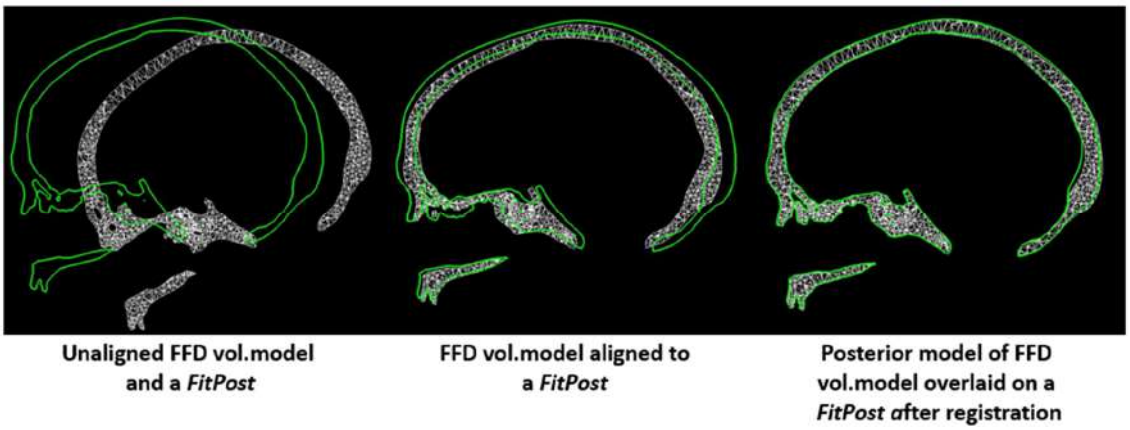

Figure 32*: A cross-section of one of the targets (green) and the FFD volumetric model (white) illustrating the process from alignment (centre) to generation of a posterior model (right).*

The regression was characterised by the movement of surface mesh points of the FFD volumetric model and similarly to the fitting process in Section 4.2.2, the robust metric estimator was used. The movement of the internal nodes and their final location on a posterior model was dependent on the vector and magnitudes of the movement of the surrounding surface mesh points. A posterior model was an instance of the FFD volumetric model that closely resembled a ***FitPost*** it was

deforming towards i.e. a posterior model was a deformation of the FFD volumetric model that best resembled its respective ***FitPost***.

The use of a FFD volumetric model developed from a combination of three SEKs produced good registration results at a global and local level. The registration errors (Hausdorff distance) for each of the registered volumetric meshes $\boldsymbol{Vol.Meshes_{reg}}$ were below 2mm as illustrated in the scatter plot in **Figure 33**. The small errors were expected since the topology of the FFD volumetric model and the ***FitPosts*** were developed from the same reference shape. This is in contrast to the previous registration step when the topology of the targets was not similar to the FFD model i.e. nasal regions. Thus, the $\boldsymbol{Vol.Meshes_{reg}}$ were representative of the ***FitPosts*** and were therefore used to sample intensity values from the corresponding CT images as detailed in Section 6.1.5.

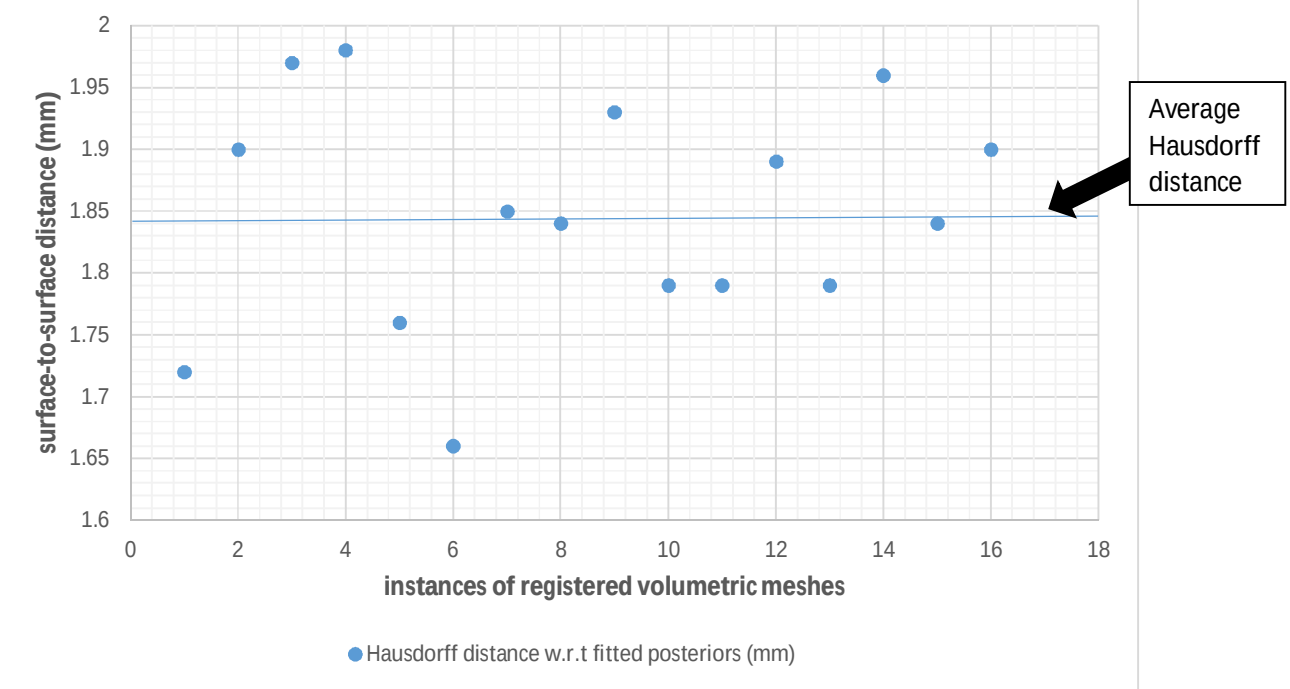

Figure 33*: A scatter plot illustrating the Hausdorff, and the average Hausdorff (horizontal blue-line) of the surface-to-surface distance between registered volumetric meshes and their respective fitted posteriors.*

6.1.5 Validating the volumetric meshes generated by the FFD volumetric model

The quality of the volumetric meshes $\boldsymbol{Vol.Meshes_{reg}}$ generated using the FFD volumetric model needed to be assessed to ensure that the tetrahedrons were not intersecting each other. If the tetrahedrons intersected, some nodes of the $\boldsymbol{Vol.Meshes_{reg}}$ would have been on the same location. Such volumetric meshes if used to sample intensity values from CT images, would give false intensity data because such nodes sample from the same location of the CT image. The aspect ratio of the volumetric mesh developed in section 6.1.2 ***Vol.Mesh***, and which was used to develop the FFD

volumetric mesh ranged between 0.0 and 1.50. Aspect ratio values of volumetric meshes generated in Amira and which are below 20 are an indication of volumetric meshes with no intersecting triangles (Guide, 2009). For the ***Vol.Meshes***$_{reg}$, the aspect ratio of their corresponding triangles to those in ***Vol.Mesh*** also varied between 0.0 and 1.50. Thus the quality of the tetrahedrons post-registration had been maintained.

6.1.6 Sampling intensities from CT images using volumetric meshes and building a SAM

The registered volumetric meshes ***Vol.Meshes***$_{reg}$ were each overlaid onto their original CT images. The intensity values on the CT images corresponding to the location of the nodes were retrieved for each ***Vol.Meshes***$_{reg}$ (**Figure 34**). However, eleven of the CT images were "incomplete" – not corresponding to their ***Vol.Meshes***$_{reg}$ and thus, there were no intensity values to sample from where nodes on the ***Vol.Meshes***$_{reg}$ were located on the missing regions of the CT images. The original 3D surfaces (Section 3.2) developed from the CT images had nodes that corresponded to the CT images – their surfaces were within the region of the CT images. However, during the two registration processes (in developing the SSM and in establishing volumetric correspondence) part of the surface of the ***Vol.Meshes***$_{reg}$ used in place of the original 3D surface had either been slightly deformed inside (**Figure 34**) or outside (**Figure 35**) the region of the CT images i.e. part of the surface of the skull did not meet the 3D volume where the cortical bone region characterised with high intensity is located. During the sampling of intensities from CT images using the volumetric meshes, it was noted that some regions of the volumetric meshes were either slightly inside (**Figure 34**) or slightly outside (**Figure 35**) the region of the CT images. Where the regions were slightly inside and thus not aligned to the region of the cortical bone of the CT image, the nodes located in this region sampled from the trabecular rather than the cortical bone (evident by the yellow colour of the retrieved intensity values in sections of **Figure 34**. The entire surface of **Figure 34** should be green since cortical bone is ideally of the same bone density). Where the regions were outside the CT images, the nodes on the surface of the volumetric mesh in this region were located outside the region of the CT image and there was nothing to sample from. The slight deformation outside or inside the expected region might have been introduced by the two registration processes. It is highly likely that the first registration process introduced a larger percentage of these defect because the reference mesh used to establish correspondence across the targets was of a different topology to the targets as evident around its nasal region **Figure 16**. A mechanism of choosing a reference mesh that not only incorporates key features expected in a human skull but also one that has the closest resemblance to most of the targets onto which it registers, would be useful.

Only five sets of volumetric meshes and their corresponding CT images were used to create the SAM. The five sets were each used to retrieve intensity values for each of the 29545 nodes. The set of intensity values for each of the five volumetric meshes were converted into vectors. The multivariate normal distribution module in Scalismo was used to calculate the variations across the sets of intensities. This generated a SAM which was composed of a mean set of intensities across the five sets and a covariance. The model was a collection of scalar values which could not be visualised independently. To visualise the model required a sample of intensity values to be generated from the model. This sample of scalar values was then warped onto a mesh and visualised similarly to what was done when intensity values were being retrieved from CT images (**Figure 34**).

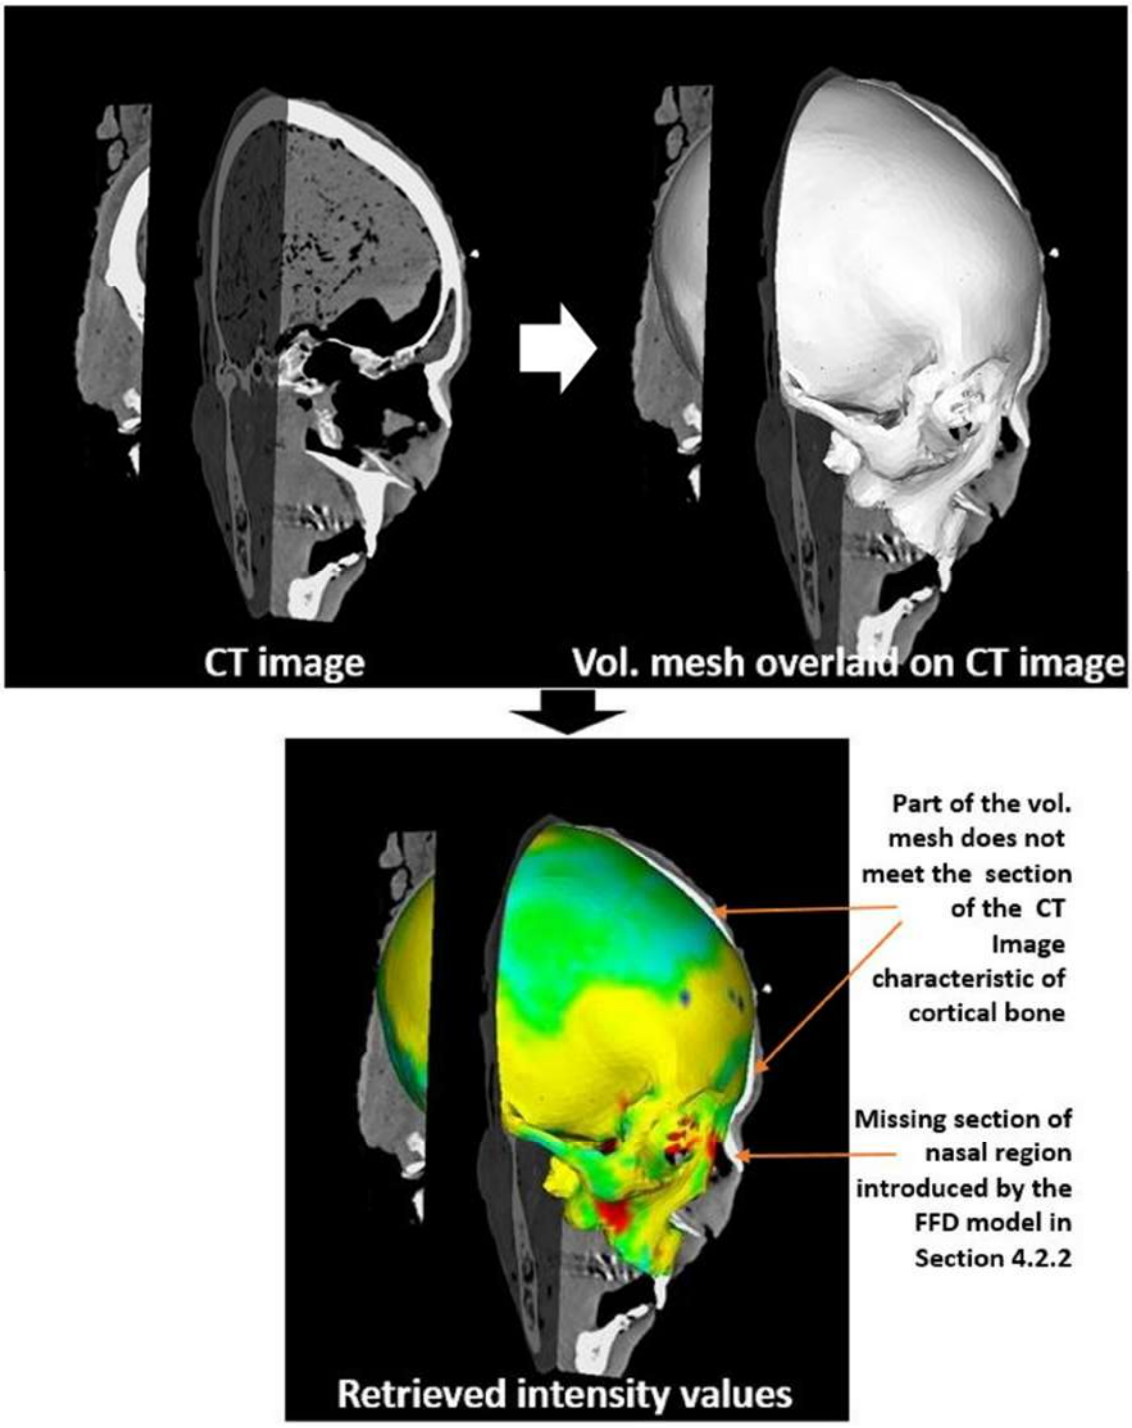

Figure 34: *Top-left: CT image of one of the volumetric meshes; top-right: the volumetric mesh overlaid on its respective CT image and; bottom: Intensity values retrieved and highlighted in various colors based on an absolute value range (green- intensity of cortical bone, yellow- trabecular bone and red - background). (NB: part of the surface of the volumetric mesh deformed slightly inside the CT image and does not meet part of the 3D volume characterised by cortical bone).*

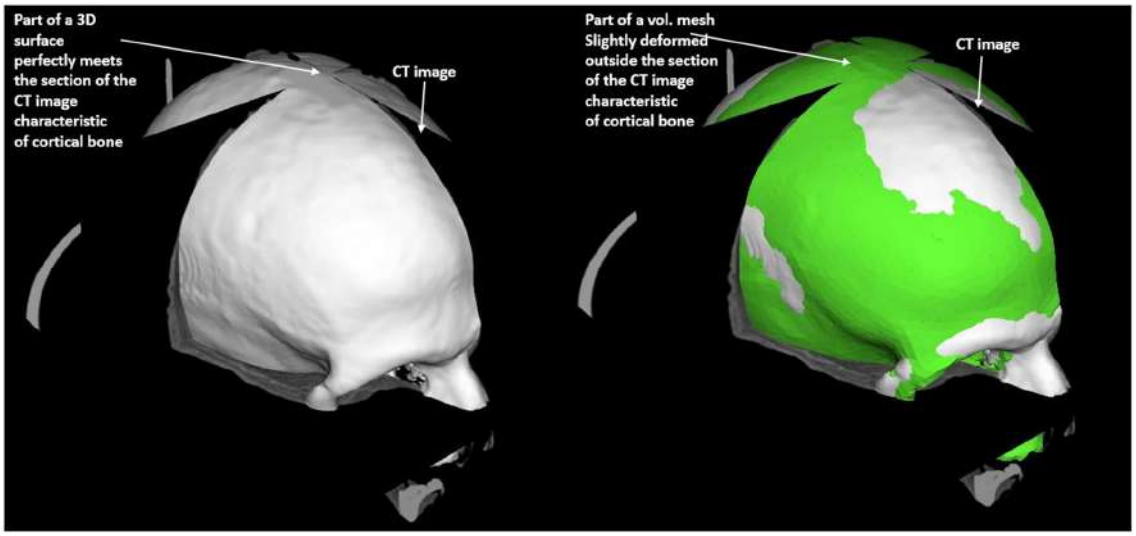

Figure 35: *Illustration of a 3D surface (white) located (left) just within and (right) a volumetric mesh (green) located slightly outside the region of a CT image. The area illustrated by the white arrow on the image on the right is just above/outside the region occupied by the CT image (section of the CT image characteristic of cortical bone). Nodes on this region do not have corresponding regions on the CT to sample intensities from.*

Before the SAM could be used for medical or research purposes, it had to be validated, as discussed in Section 6.2.

6.2 VALIDATING THE SKULL SAM

The SAM model was composed of a mean vector (composed of 29545 intensity values) and a vector of standard deviations (composed of 29545 values of standard of deviations from the mean vector). The mean of the standard deviation was 711.47 i.e. each of the 29545 mean intensity values making up the mean vector could be subjected to a standard deviation of 711.47 to generate new samples. This was validated by generating three vector samples of intensity values from the SAM. A standard deviation for each of the 29545 intensity values of the samples from the mean vector of intensity values of the SAM were calculated. The average standard deviation of the three samples from the mean vector of the SAM ranged between 294.85 and 593.82. The values were within the standard deviation (711.47) expected for sampled generated by the SAM. Hence, the model could generate intensity samples that were specific to the values used to develop it.

6.3 CHALLENGES OF COMBINED SSM-SAM MODEL

An attempt was made to combine the shape and appearance models to obtain a statistical shape and appearance model (SSIM). The combined model was to be used to fit onto incomplete skulls therefore reconstructing the missing surface together with the associated intensity values. Such a feature would increase the application of the statistical models. A skull SSM-SAM that can reconstruct shape and intensity values of missing parts of a skull can be resourceful in reconstructive

surgery, patient-specific implant design, automated segmentation and finite element analysis. The development of the combined model was done by building a model based on a tuple constituting vertex coordinates of points of the volumetric meshes and their respective intensity values (Anon, 2017b). The model was built by applying multivariateNormalDistribution (a module in Scalismo) to a sequence of the tuple from the dataset thereby generating a combined shape and appearance model. However, the combined model had some challenges which limited its applicability in reconstructing incomplete skulls and thus validate the SAM.

6.3.1 Visualization

By combining the vertex coordinates and the intensity values, the model could no longer be defined on a particular domain. The development platform (Scalismo) has two domains i.e. mesh domain and image domain that should be predefined before developing a model. The domains are required for visualization. While it was possible to generate samples from the developed combined model, visualization of the sample (constituting mesh points and associated intensities) was not possible.

6.3.2 Incompatibility with Scalismo

The motivation for combining the shape and appearance model was that it could be used in a fitting process (similar to the registration process detailed in Section 4.2.2). The assumption was that given an incomplete mesh, a combined model and corresponding mesh points of the two, one could use the Gaussian process regression to generate fitted posteriors defined by both mesh and intensity features. However, the GP regression module in Scalismo cannot be constrained by part of the value (mesh points) to generate values not constrained (which in this case were intensity values). To use the GP regression module in Scalismo together with a combined model would require modifying the combined model to introduce uncertainties to the mesh points and intensities. Additionally, the multivariateNormalDistribution component would need to be used together with the deformation fields of the mesh points. Therefore, this was not implemented and will form part of future work.

If the GP regression module of Scalismo had worked with the combined model (use the deformation obtained from mesh points to infer intensity values) it would have been used to fit onto two incomplete surface meshes of the skulls that were generated using a cutting tool module in Scalismo.

6.4 INCOMPLETE SURFACE MESHES

Incomplete surface meshes were required to test the fitting property of the skull SSM thereby further testing the validity of the SSM. Generating incomplete surface meshes from complete surface meshes ensure a ground truth that can be compared to a repaired incomplete surface meh is available.

6.4.1 Incomplete surface mesh generation

To generate an incomplete surface mesh, a landmark was identified on a surface mesh. A radius of a circle whose centre was the identified landmark was declared. A circle using the landmark and radius was cut off from the surface mesh creating an incomplete surface mesh (**Figure 36**).

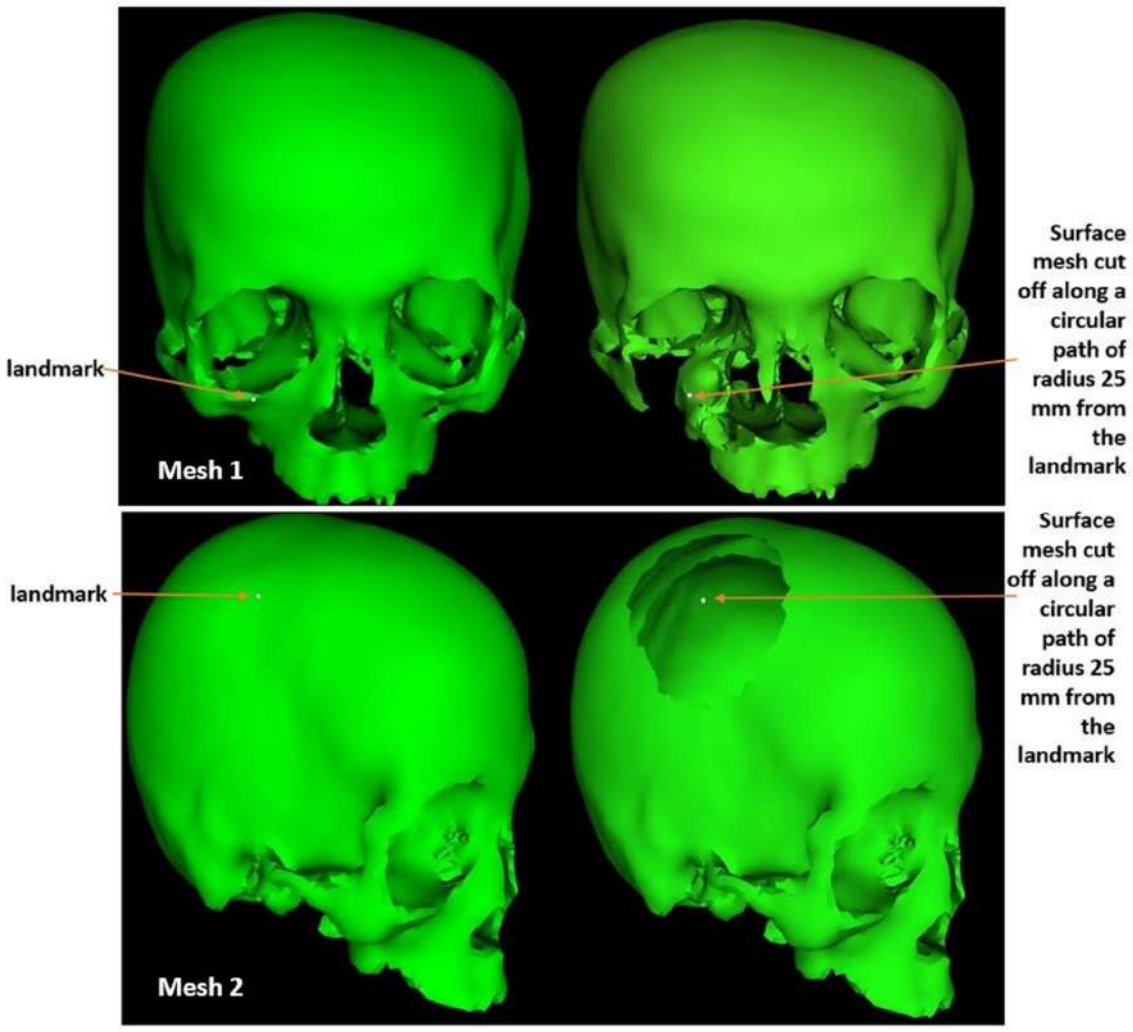

***Figure 36**: Generating incomplete surface meshes in Scalismo. Generating an incomplete mesh by cutting-off sections of the maxilla bone (Top row) and the parietal bone (bottom row).*

6.4.2 Fitting a volumetric SSM to incomplete meshes

Since the GP regression module in Scalismo could not work with the combined model, only a volumetric SSM generated from the in-correspondence volumetric meshes (Using the SSM development method discussed in Section 5.2) was used in the fitting process. The volumetric SSM was aligned to the incomplete mesh. A posterior of the model to the incomplete mesh was generated using the surface mesh-points of the volumetric model for the regression process as illustrated in Section 4.2.2. The mean of the posterior represented the "most-likely" representation of the incomplete mesh. Hence the posterior mean was taken as the reconstructed surface (**Figure 37**). The Hausdorff distance between the posterior mean and the respective original skulls from which the incomplete skulls were developed from was measured. From two reconstructions conducted, the Hausdorff distances were 23.41 mm and 9.01 mm for mesh1 and mesh2 respectively.

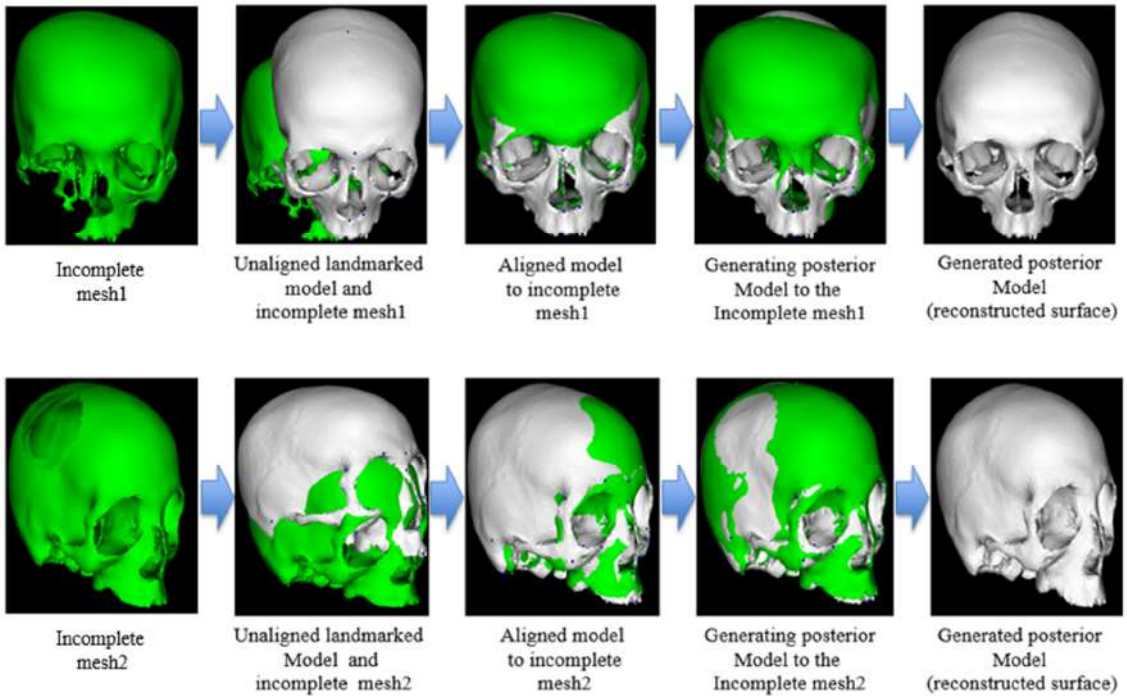

***Figure 37**: The reconstruction of incomplete skulls using the shape model generated from the combined shape and appearance model. First row: incomplete skull with hole on the maxilla bone and Second Row: incomplete skull with hole in the crania.*

The Hausdorff distance between the posterior mean model of mesh1 and its respective surface mesh was high (23.41mm) because of the high variability of the human skull around the maxilla, zygomatic and nasal bones. For registration/fitting, this region is normally constrained by approximately four landmarks i.e. the ectoconchion left, orbitale left, nasal left and zygomaxillare left (Section 2.2.2). In generating an incomplete surface with a missing section in this region, the fitting

results were affected as a result of the missing detail in this region. In contrast, there is limited variation around the parietal bone and during the registration/ fitting process of the skull, no landmark is used to constrain this section (Section 2.2.2). Hence, generating an incomplete surface mesh with a missing section in this region does not affect the registration results. This is the reason for the relatively low Hausdorff distance value for mesh2 in comparison to that of mesh1.

6.5 SUMMARY

This chapter discussed the development and validation of a skull SAM from in-correspondence volumetric meshes and CT images. The in-correspondence volumetric meshes were developed by registering a volumetric mesh onto surface meshes. The registration process was achieved by developing a free-form deformation (FFD) volumetric model from the volumetric mesh which was then fitted onto the surface meshes. To reduce the rigidity of the volumetric mesh, three squared exponential kernels were used in developing the FFD volumetric model thereby ensuring the model was flexible enough to deform and register onto each of the surface meshes. The registered volumetric meshes were overlaid onto their original CT images and intensity values located at the corresponding location of the meshes nodes were retrieved for each volumetric mesh.

Out of the sixteen sets of CT images, eleven were incomplete and could not be used to retrieve intensity values from them. Multivariate normal distribution was applied to the set of intensity values to generate a SAM which was defined by a mean set of intensity values and a standard deviation for each value. The SAM was validated by generating intensity values using the model and checking if the set of values was within the standard deviation. An attempt to combine the SSM and SAM and use it for further validation was not possible. In trying to use the combined model for validating the SAM, several challenges were encountered. There was a limitation on the GP regression module in Scalismo in generating posterior models of combined shape and intensities. Hence, the combined model, and therefore the SAM, could not be farther validated. It would have required introducing additional features (uncertainty to the mesh points and intensity values) to the model. As a result, only the shapes of the incomplete skulls were reconstructed and thereby validated. The validation of the combined model forms part of future work which will also focus on solving the visualization limitation.

The process of SAM generation from in-correspondence volumetric meshes is summarised in **Figure 38**.

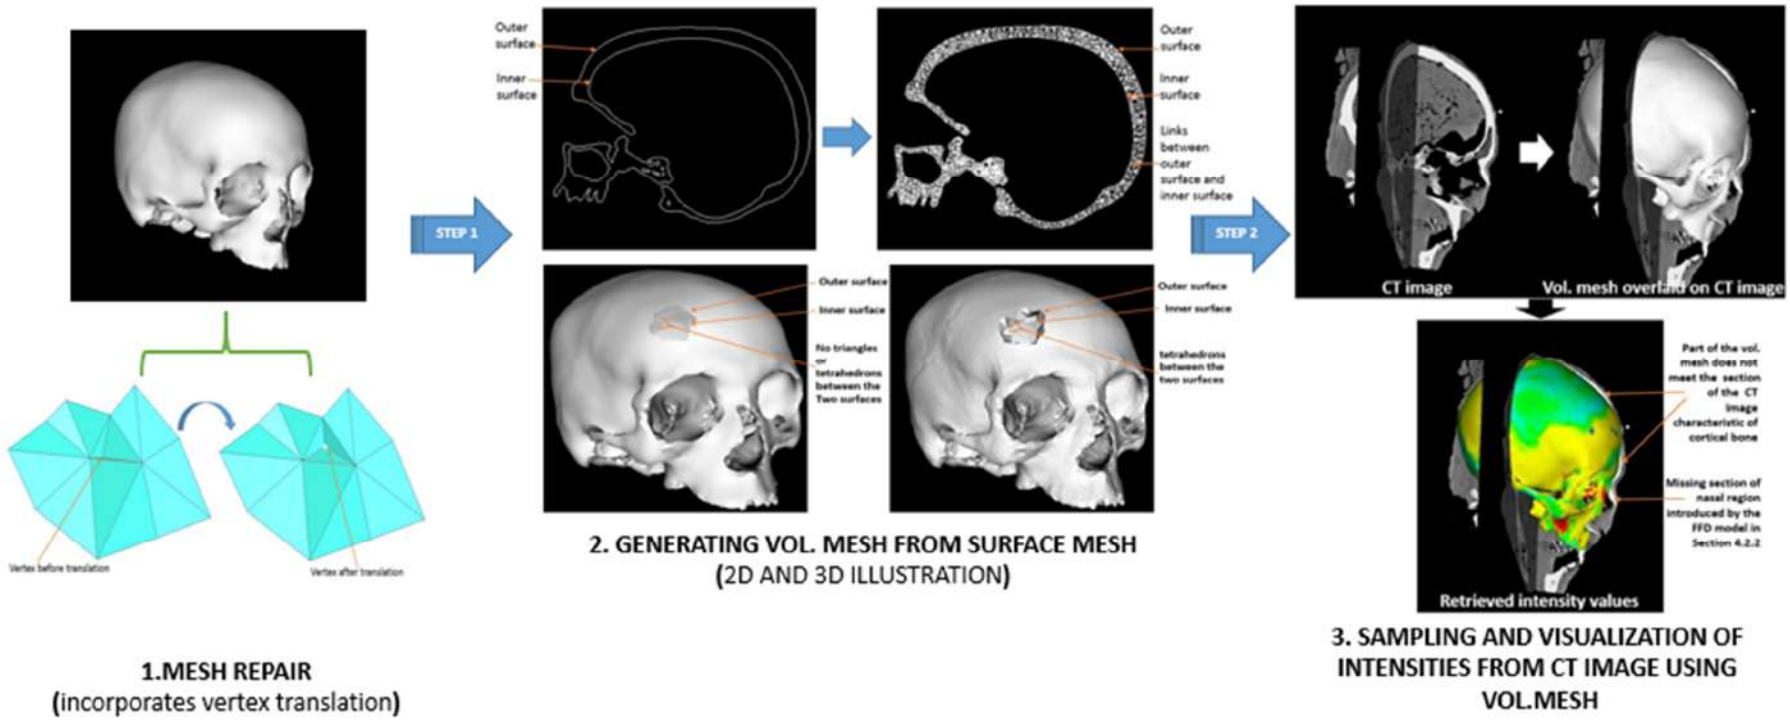

Figure 38: *A summary of the process of SAM generation from in-correspondence volumetric meshes.*

7 DISCUSION AND CONCLUSION

This project involved the development of a SSM and SAM of the skull. In the development, a total of 16 cadaveric heads were used. This number was limited because of the availability of cadaveric samples. The effect of the limited number of training dataset was a model that had a limited range of variation defined by a mean and only fifteen modes of variation. However, the primary modes of variation carry most of the variation in a model. In the project, out of 16 training dataset, the first eight principal components accounted for 92.13% of the variation of the SSM.

7.1 DATA COLLECTION AND PRE-PROCESSING

During the segmentation protocol, the thresholding technique was used. The technique was effective at separating two regions with different intensity values. However, where segmentation was required for regions with similar intensity values such as the segmentation of the mandibular bone from the cranium or segmenting out sections of the nasal sinuses, the technique was not effective. Manual segmentation was utilised for the mentioned cases. The manual segmentation was time consuming. However, because of the limited number of training dataset, the process was manageable. Therefore, there is need for a consistent segmentation protocol and segmentation quality assessment technique. Additionally, if a larger sample size is utilised, the manual segmentation can increase the time required to develop the shape and appearance model pipelines. An automatic segmentation technique not only based on intensity values but also prior information of the bone being segmented, would therefore be useful. Model-based segmentation is an application of SSM (Heimann & Meinzer, 2009); by further including the intensity characteristic to a SSM, the model-based segmentation technique can be developed for the skull based on the pipeline presented. The development of an automated model-based segmentation technique was not part of the scope of this dissertation, but is recommended for future work.

7.2 REGISTRATION OF 3D SURFACE MESHES

For the identification of reliable landmarks to be used in the project, a total of 33 anatomical landmarks were used. These landmarks were chosen based on the ease of an observer to identify their location on a 3D surface. The landmark identification process is therefore subjective. Additionally, the 3D surfaces are meshes defined by nodes. The location of a landmark identified is a vertex position near the mark placed by an observer. Some meshes are composed of tens to hundreds of thousands of points and it is impossible to distinguish one vertex from another with the naked eye. It was therefore not possible to precisely select the same vertex position more than once even for the

same observer. The landmark will always be the vertex position closest to the mark placed by the observer on the 3D surface. A better land-marking protocol that can reliably and repeatability select the vertex position of a landmark to ensure consistency is required, especially when more than one observer is identifying the landmarks.

The establishment of a dense correspondence protocol was based on the free-form deformation (FFD) registration technique. A FFD model is defined by a reference and a kernel. A squared exponential kernel was used in this project where its optimal values for the span and magnitude of the kernel were determined experimentally. Future work should improve the process to determine the optimal values of the span and magnitude of the kernel simultaneously. The FFD model was registered onto target skulls and its fit used as a representation of the target. However, the registered surfaces with respect to the target skulls had large Hausdorff distances in the tens of millimetres. In future, the individual registered surfaces can be used to develop new FFD models that can be registered to the target skull. Because the new FFD is almost identical to the target skull, the registration would be further improved thus reducing the Hausdorff distance between the registered surface and the target mesh.

7.3 DEVELOPMENT AND VALIDATION OF A SKULL SSM AND SAM

Lack of literature on the development of a skull SSM meant that there were no results with which to compare the SSM validation. The compactness, specificity and generality of the model were as follows: 8 principal components represented 92.13% of the variance; the specificity of the model had a steady state value of 2.04mm and; the generality of the model had a steady state value of 1.48mm. The SAM was developed from five training datasets. The number reduced from sixteen to five because ten of the CT scans from which the volumetric meshes were to sample intensities, were incomplete and therefore could not be used. The SAM was characterised by scalar values of the mean intensity values across the various nodes of the volumetric meshes and their respective standards of deviation. To validate the SAM, the SSM and SAM were combined to form a combined shape and appearance model. However, the combined model could not be visualised, neither could it be used in the fitting process of an incomplete surface mesh of the skull. The GP regression module in Scalismo could not be constrained by a section of the model (mesh points) to generate a posterior with both mesh points and their respective intensities. In future, the GP regression module in Scalismo should be improved so that it can used to develop a combined model for the generation of posterior models and fitting exercises. Because of the mentioned limitation, the combined model and in extension the SAM could not be validated.

7.4 CONCLUSION

All objectives set out at the beginning of the project were achieved except a section of the final objective which was to develop and validate a SAM of the skull. Completing this objective would have required a significant modification of the software libraries in Scalismo, with the associated time investment. It was therefore decided to recommend this as future work. All other objectives produced results that enabled further analysis of the SSM and SAM development pipeline. The project produced academic outputs that were presented at conferences in 2017:

- B Lugadilu, C Richards, C Reyneke, T Mutsvangwa, TS Douglas. A statistical shape model of the skull developed from a South African population. Proceedings, Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Jeju, South Korea, July 2017.
- B Lugadilu, C Richards, C Reyneke, TS Douglas, T Mutsvangwa. A statistical shape and appearance model of the skull from a South African population. Human Biology, Integrative Biomedical Sciences, and Pathology Research Day, University of Cape Town, South Africa, October 2017.
- B Lugadilu, C Reyneke, TS Douglas, T Mutsvangwa. A statistical appearance model of the skull from a South African population. Proceedings, South African Institute of Mechanical Engineering Western Cape Postgraduate Conference, Stellenbosch, South Africa, November 2017.

Where the SAM developed using this pipeline is to be applied in the determination of bone mineral density, calibrated CT is needed to ensure the pixel values across the different CT images represent the same bone mineral density values.

REFERENCES

Abdelmunim, H., & Farag, A. A. (2011). *Elastic shape registration using an incremental free form deformation approach with the icp algorithm.* Paper presented at the Computer and Robot Vision (CRV), 2011 Canadian Conference on.

Albrecht, T., Lüthi, M., Gerig, T., & Vetter, T. (2013). Posterior shape models. *Medical image analysis, 17*(8), 959-973.

Altbeker, A. (2008). Murder and robbery in South Africa: A tale of two trends. *Crime, violence and injury prevention in South Africa: data to action. Tygerberg: MRC-UNISA Crime, Violence and Injury Lead Programme*, 122-149.

Anon (Producer). (2017a, 06 18). Futurelearn. *statistical shape modelling, University of Basel*. Retrieved from https://www.futurelearn.com/courses/statistical-shape-modelling/2/todo/7755

Anon (Producer). (2017b, 06 18). Scalismo, scalable image analysis and shape modelling. *Scalismo, scalable image analysis and shape modelling*. Retrieved from https://github.com/unibas-gravis/scalismo

Babaud, J., Witkin, A. P., Baudin, M., & Duda, R. O. (1986). Uniqueness of the Gaussian kernel for scale-space filtering. *IEEE Transactions on Pattern Analysis & Machine Intelligence*(1), 26-33.

Baka, N., Kaptein, B. L., de Bruijne, M., van Walsum, T., Giphart, J., Niessen, W. J., & Lelieveldt, B. P. (2011). 2D–3D shape reconstruction of the distal femur from stereo X-ray imaging using statistical shape models. *Medical image analysis, 15*(6), 840-850.

Barr, A. H. (1987). Global and local deformations of solid primitives *Readings in Computer Vision* (pp. 661-670): Elsevier.

Belkin, M., & Niyogi, P. (2002). *Laplacian eigenmaps and spectral techniques for embedding and clustering.* Paper presented at the Advances in neural information processing systems.

Bennis, C., Vézien, J.-M., & Iglésias, G. (1991). *Piecewise surface flattening for non-distorted texture mapping.* Paper presented at the ACM SIGGRAPH computer graphics.

Berar, M., Desvignes, M., Bailly, G., & Payan, Y. (2006). Statistical skull models from 3D X-ray images. *arXiv preprint physics/0610182*.

Besl, P. J., & McKay, N. D. (1992). *Method for registration of 3-D shapes.* Paper presented at the Sensor Fusion IV: Control Paradigms and Data Structures.

Boas, F. E., & Fleischmann, D. (2011). Evaluation of two iterative techniques for reducing metal artifacts in computed tomography. *Radiology, 259*(3), 894-902.

Boehler, T., & Peitgen, H.-O. (2009). *Evaluation of Image Registration using a Mass-spring Model of the Breast.* Paper presented at the World Congress on Medical Physics and Biomedical Engineering, September 7-12, 2009, Munich, Germany.

Bonaretti, S., Seiler, C., Boichon, C., Reyes, M., & Büchler, P. (2014). Image-based vs. mesh-based statistical appearance models of the human femur: implications

for finite element simulations. *Medical engineering & physics, 36*(12), 1626-1635.

Bookstein, F. L. (1989). Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 11*(6), 567-585.

Bouabene, G. L., M;. (2017). Statistical Shape Modelling. Retrieved from https://www.futurelearn.com/courses/statistical-shape-modelling/2/todo/7748

Brown. (1992). A survey of image registration techniques. *ACM computing surveys (CSUR), 24*(4), 325-376.

Brown, & Rusinkiewicz, S. (2004). *Non-rigid range-scan alignment using thin-plate splines.* Paper presented at the 3D Data Processing, Visualization and Transmission, 2004. 3DPVT 2004. Proceedings. 2nd International Symposium on.

Cahill, N. D., Noble, J. A., & Hawkes, D. J. (2009). *A demons algorithm for image registration with locally adaptive regularization.* Paper presented at the International Conference on Medical Image Computing and Computer-Assisted Intervention.

Chen, C.-S., Hung, Y.-P., & Cheng, J.-B. (1999). RANSAC-based DARCES: A new approach to fast automatic registration of partially overlapping range images. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 21*(11), 1229-1234.

Cootes, & Taylor, C. J. (2001). *Statistical models of appearance for medical image analysis and computer vision.* Paper presented at the Medical Imaging 2001: Image Processing.

Cootes, Taylor, C. J., Cooper, D. H., & Graham, J. (1995). Active shape models-their training and application. *Computer vision and image understanding, 61*(1), 38-59.

Coovadia, H., Jewkes, R., Barron, P., Sanders, D., & McIntyre, D. (2009). The health and health system of South Africa: historical roots of current public health challenges. *The Lancet, 374*(9692), 817-834.

Crum, W. R., Hartkens, T., & Hill, D. (2004). Non-rigid image registration: theory and practice. *The British journal of radiology, 77*(suppl_2), S140-S153.

David, D. J., Hemmy, D. C., & Cooter, R. D. (2012). *Craniofacial deformities: atlas of three-dimensional reconstruction from computed tomography*: Springer Science & Business Media.

Davies, R. H., Twining, C. J., Cootes, T. F., Waterton, J. C., & Taylor, C. J. (2002). A minimum description length approach to statistical shape modeling. *IEEE transactions on medical imaging, 21*(5), 525-537.

Davies, R. H., Twining, C. J., & Taylor, C. (2008). Groupwise surface correspondence by optimization: Representation and regularization. *Medical image analysis, 12*(6), 787-796.

Dedner, A., Lüthi, M., Albrecht, T., & Vetter, T. (2007). *Curvature guided level set registration using adaptive finite elements.* Paper presented at the Joint Pattern Recognition Symposium.
Drake, R., Vogl, A. W., Mitchell, A. W., Tibbitts, R., & Richardson, P. (2014). *Gray's Atlas of Anatomy E-Book*: Elsevier Health Sciences.
Dryden, I. L., Faghihi, M. R., & Taylor, C. C. (1997). Procrustes shape analysis of planar point subsets. *Journal of the Royal Statistical Society: Series B (Statistical Methodology), 59*(2), 353-374.
Dryden, I. L., & Mardia, K. V. (2016). *Statistical shape analysis: with applications in R*: John Wiley & Sons.
Floater, M. S., & Hormann, K. (2005). Surface parameterization: a tutorial and survey *Advances in multiresolution for geometric modelling* (pp. 157-186): Springer.
Galeotti, J. (2016). Lecture 20: deformable/ non-rigid registration. *powerpoint presentation*.
Gateno, J., Xia, J. J., Teichgraeber, J. F., Christensen, A. M., Lemoine, J. J., Liebschner, M. A., . . . Briggs, M. E. (2007). Clinical feasibility of computer-aided surgical simulation (CASS) in the treatment of complex cranio-maxillofacial deformities. *Journal of Oral and Maxillofacial Surgery, 65*(4), 728-734.
Gavin, H. P. (2013). The Nelder-Mead algorithm in two dimensions. *CEE 201L. Duke U*.
George, P.-L., & Borouchaki, H. (1998). Delaunay triangulation and meshing.
Gerig, T., Shahim, K., Reyes, M., Vetter, T., & Lüthi, M. (2014). *Spatially varying registration using gaussian processes.* Paper presented at the International Conference on Medical Image Computing and Computer-Assisted Intervention.
Gobindchandra, K., & Santhosh, K. (2015). Analysis of image segmentation techniques. *International Research Journal of Computer Science, 6*(2).
Goodall, C. (1991). Procrustes methods in the statistical analysis of shape. *Journal of the Royal Statistical Society. Series B (Methodological)*, 285-339.
Goshtasby, A. A., Sonka, M., & Udupa, J. (2000). Analysis of Volumetric Images. *Computer vision and image understanding, 77*(2), 79-83.
Guide, A. U. s. (2009). Amira 5. *San Diego, CA: Visage Imaging*, 40-44.
Hajnal, J. V., & Hill, D. L. (2001). *Medical image registration*: CRC press.
Härdle, W. K., Müller, M., Sperlich, S., & Werwatz, A. (2012). *Nonparametric and semiparametric models*: Springer Science & Business Media.
Hedley, M., & Yan, H. (1992). Motion artifact suppression: a review of post-processing techniques. *Magnetic resonance imaging, 10*(4), 627-635.
Heimann, T., & Meinzer, H.-P. (2009). Statistical shape models for 3D medical image segmentation: a review. *Medical image analysis, 13*(4), 543-563.
Hiemann, T., & Hans-Peter, M. (2009). Statistical shape models for 3D medical image segmentation: a review. *medical image analysis*, 543-563.
Hochfeld, M., Lamecker, H., Thomale, U.-W., Schulz, M., Zachow, S., & Haberl, H. (2014). Frame-based cranial reconstruction. *Journal of Neurosurgery: Pediatrics, 13*(3), 319-323.

Hormann, K., Lévy, B., & Sheffer, A. (2007). Mesh parameterization: Theory and practice.

Hotta, S. (2008). Manifold Matching for High-Dimensional Pattern Recognition *Pattern Recognition Techniques, Technology and Applications*: InTech.

Hristidis, V. (2009). *Information discovery on electronic health records*: CRC Press.

Hu, Y., Duan, F., Zhou, M., Sun, Y., & Yin, B. (2012). Craniofacial reconstruction based on a hierarchical dense deformable model. *EURASIP Journal on Advances in Signal Processing, 2012*(1), 217.

Huang, X., & Tsechpenakis, G. (2009). Medical image segmentation. *Information Discovery on Electronic Health Records, 10*, 251-289.

Huber, P., & Ronchetti, E. (1981). Robust statistics, ser. *Wiley Series in Probability and Mathematical Statistics. New York, NY, USA: Wiley-IEEE, 52*, 54.

Huttenlocher, D. P., Klanderman, G. A., & Rucklidge, W. J. (1993). Comparing images using the Hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 15*(9), 850-863.

Jones, T., & Forrest, S. (1995). Genetic algorithms and heuristic search. *Santa Fe Institute Tech. Report 95-02, 21*.

Kim, S.-G., Yi, W.-J., Hwang, S.-J., Choi, S.-C., Lee, S.-S., Heo, M.-S., . . . Yoo, J. H. (2012). Development of 3D statistical mandible models for cephalometric measurements. *Imaging science in dentistry, 42*(3), 175-182.

Kittle, J., & Young, P. (1973). A new application to feature selection based on the karhunenloeve expansion. *Pattern recognition, 5*.

Kohavi, R. (1995). *A study of cross-validation and bootstrap for accuracy estimation and model selection.* Paper presented at the Ijcai.

Kotcheff, A. C., & Taylor, C. J. (1998). Automatic construction of eigenshape models by direct optimization. *Medical image analysis, 2*(4), 303-314.

Kun, Z. (2014). *Dense Correspondence and Statistical Shape Reconstruction of Fractured, Incomplete Skulls.*

Li, H., Sumner, R. W., & Pauly, M. (2008). *Global correspondence optimization for non-rigid registration of depth scans.* Paper presented at the Computer graphics forum.

Lorensen, W. E., & Cline, H. E. (1987). *Marching cubes: A high resolution 3D surface construction algorithm.* Paper presented at the ACM siggraph computer graphics.

Loubele, M., Jacobs, R., Maes, F., Schutyser, F., Debaveye, D., Bogaerts, R., . . . Marchal, G. (2005). Radiation dose vs. image quality for low-dose CT protocols of the head for maxillofacial surgery and oral implant planning. *Radiation protection dosimetry, 117*(1-3), 211-216.

Luo, L., Wang, M., Tian, Y., Duan, F., Wu, Z., Zhou, M., & Rozenholc, Y. (2013). Automatic sex determination of skulls based on a statistical shape model. *Computational and mathematical methods in medicine, 2013*.

Masters, S. J., McClean, P. M., Arcarese, J. S., Brown, R. F., Campbell, J. A., Freed, H. A., . . . Koziol, D. F. (1987). Skull x-ray examinations after head trauma. *New England Journal of Medicine, 316*(2), 84-91.

Mather, R. (2005). Patient Focused Imaging--Aquilion's Low Dose Vision. *Diagnostic Imaging, 89*, 4-8.

Mohiuddin, S. (2013). Change in bone density as a function of water content. *World J Med Sci, 8*(1), 48-51.

Myronenko, A., Song, X., & Carreira-Perpinán, M. A. (2007). *Free-form nonrigid image registration using generalized elastic nets.* Paper presented at the Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on.

Neuert, M. A. (2016). The Development And Application Of A Statistical Shape Model Of The Human Craniofacial Skeleton.

Nutcher, A., Lingemann, K., & Hertzberg, J. (2007). Cached KD tree search for ICP algorithms. *In 3D digital imaging and modelling. 3DIM'07. Sixth international conference on*, 419-426.

Offut, C. V., M.; Gilula, L.; marsh, J.; Sutherland, C.;. (1990). Volumetric 3D imaging of computed tomography scans. *Radiological technology, 61*.

Ono, I., Ohura, T., Narumi, E., Kawashima, K., Matsuno, I., Nakamura, S., . . . Tanaka, F. (1992). Three-dimensional analysis of craniofacial bones using three-dimensional computer tomography. *Journal of cranio-maxillo-facial surgery, 20*(2), 49-60.

Otsu, N. (1979). A threshold selection method from gray-level histograms. *IEEE transactions on systems, man, and cybernetics, 9*(1), 62-66.

Pal, N. R., & Pal, S. K. (1993). A review on image segmentation techniques. *Pattern recognition, 26*(9), 1277-1294.

Pardhi, Shruti, & Wanjale, K. H. (2016). Survey on techniques involved in image segmentation. *International journal of computer science trends and technology (IJCST)*, 275-280.

Park, H., Martin, G. R., & Bhalerao, A. (2010). Local affine image matching and synthesis based on structural patterns. *IEEE Transactions on Image Processing, 19*(8), 1968-1977.

Paulsen, R. R., Baerentzen, J. A., & Larsen, R. (2009). Markov random field surface reconstruction. *IEEE Transactions on Visualization & Computer Graphics*(4), 636-646.

Paysan, P., Lüthi, M., Albrecht, T., Lerch, A., Amberg, B., Santini, F., & Vetter, T. (2009). *Face reconstruction from skull shapes and physical attributes.* Paper presented at the Joint Pattern Recognition Symposium.

Phillips, J. M., Liu, R., & Tomasi, C. (2007). *Outlier robust ICP for minimizing fractional RMSD.* Paper presented at the 3-D Digital Imaging and Modeling, 2007. 3DIM'07. Sixth International Conference on.

Rangarajan, A., Chui, H., & Bookstein, F. L. (1997). *The softassign procrustes matching algorithm.* Paper presented at the Biennial International Conference on Information Processing in Medical Imaging.

Rooppakhun, S., Chantarapanich, N., & Sitthiseripratip, K. (2011). Advanced medical imaging and reverse engineering technologies in craniometric study *Forensic Medicine-From Old Problems to New Challenges*: InTech.

Rueckert, D., & Aljabar, P. (2010). Nonrigid registration of medical images: Theory, methods, and applications [applications corner]. *IEEE Signal Processing Magazine, 27*(4), 113-119.

Ruppert, J. (1995). A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *Journal of algorithms, 18*(3), 548-585.

Saboret, L., Alliez, P., & Levy, B. (2006). Planar parameterization of triangulated surface meshes. *CGAL User and Reference Manual. CGAL Editorial Board, 3*.

Salfati, C. G., Horning, A. M., Sorochinski, M., & Labuschagne, G. N. (2015). South African serial homicide: Consistency in victim types and crime scene actions across series. *Journal of Investigative Psychology and Offender Profiling, 12*(1), 83-106.

Sapsford, R., & Jupp, V. (2006). *Data collection and analysis*: Sage.

Sarkalkan, N., Weinans, H., & Zadpoor, A. A. (2014). Statistical shape and appearance models of bones. *Bone, 60*, 129-140.

Schwarz, L. A. (2007). Non-rigid registration using free-form deformations. *Technische Universität München*.

Sclaroff, S., & Pentland, A. (1995). *Modal matching for correspondence and recognition*. Retrieved from

Semper-Hogg, W., Fuessinger, M. A., Schwarz, S., Ellis III, E., Cornelius, C.-P., Probst, F., . . . Schlager, S. (2017). Virtual reconstruction of midface defects using statistical shape models. *Journal of Cranio-Maxillofacial Surgery, 45*(4), 461-466.

Sheffer, A., Praun, E., & Rose, K. (2007). Mesh parameterization methods and their applications. *Foundations and Trends® in Computer Graphics and Vision, 2*(2), 105-171.

Shen, J.-K., Matuszewski, B. J., Shark, L.-K., & Moore, C. J. (2006). *Deformable Image Registration using Spring Mass System.* Paper presented at the BMVC.

Sinha, A., Leonard, S., Reiter, A., Ishii, M., Taylor, R. H., & Hager, G. D. (2016). *Automatic segmentation and statistical shape modeling of the paranasal sinuses to estimate natural variations.* Paper presented at the Medical Imaging 2016: Image Processing.

Stegmann. (2000). Active appearance models: Theory, extensions and cases. *Informatics and Mathematical Modelling*, 262.

Stegmann, & Gomez, D. D. (2002). A brief introduction to statistical shape analysis. *Informatics and mathematical modelling, Technical University of Denmark, DTU, 15*(11).

Storm, & Alyson, R. (2010). *Human skeleton asymmetry. A study of directional and fluctuating asymmetry in assessing health, environmental conditions, and social status in English populations from the 7th to the 19th centuries.* Retrieved from Bradford, U.K:

Styner, M. A., Rajamani, K. T., Nolte, L.-P., Zsemlye, G., Székely, G., Taylor, C. J., & Davies, R. H. (2003). *Evaluation of 3D correspondence methods for model building.* Paper presented at the Biennial International Conference on Information Processing in Medical Imaging.

Szeliski, R., & Lavallée, S. (1996). Matching 3-D anatomical surfaces with non-rigid deformations using octree-splines. *International journal of computer vision, 18*(2), 171-186.

Szilvi-Nagy, M., & Matyasi, G. (2003). Analysis of STL files. *Mathematical and Computer Modelling: An International Journal, 38*(7-9), 945-960.

Tam, G. K., Cheng, Z.-Q., Lai, Y.-K., Langbein, F. C., Liu, Y., Marshall, D., . . . Rosin, P. L. (2013). Registration of 3D point clouds and meshes: a survey from rigid to nonrigid. *IEEE transactions on visualization and computer graphics, 19*(7), 1199-1217.

Tau, M., & Hassner, T. (2016). Dense correspondences across scenes and scales. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 38*(5), 875-888.

Thirion, J.-P. (1998). Image matching as a diffusion process: an analogy with Maxwell's demons. *Medical image analysis, 2*(3), 243-260.

Vercauteren, T., Pennec, X., Perchant, A., & Ayache, N. (2007). *Non-parametric diffeomorphic image registration with the demons algorithm.* Paper presented at the International Conference on Medical Image Computing and Computer-Assisted Intervention.

Viljoen, D. L., Gossage, J. P., Brooke, L., Adnams, C. M., Jones, K. L., Robinson, L. K., . . . Kodituwakku, P. (2005). Fetal alcohol syndrome epidemiology in a South African community: a second study of a very high prevalence area. *Journal of studies on alcohol, 66*(5), 593-604.

Vishwanathan, S., & Smola, A. J. (2004). *Binet-Cauchy Kernels.* Paper presented at the NIPS.

von Funck, W., Theisel, H., & Seidel, H.-P. (2006). *Vector field based shape deformations.* Paper presented at the ACM Transactions on Graphics (TOG).

Weinberg, S. M., Putz, D. A., Mooney, M. P., & Siegel, M. I. (2005). Evaluation of non-metric variation in the crania of black and white perinates. *Forensic science international, 151*(2-3), 177-185.

Williams, C. K., & Rasmussen, C. E. (2006). Gaussian processes for machine learning. *the MIT Press, 2*(3), 4.

Yogamangalam, R., & Karthikeyan, B. (2013). Segmentation techniques comparison in image processing. *International Journal of Engineering and Technology (IJET), 5*(1), 307-313.

Zachow, S., Lamecker, H., Elsholtz, B., & Stiller, M. (2005). *Reconstruction of mandibular dysplasia using a statistical 3D shape model.* Paper presented at the International congress series.

Zeng, Y., Wang, C., Wang, Y., Gu, X., Samaras, D., & Paragios, N. (2010). Dense non-rigid registration using high-order graph matching. *In computer vision and pattern recognition (CVPR)*, 382-389.

Zhang. (1996). A survey on evaluation methods for image segmentation. *Pattern recognition, 29*(8), 1335-1346.

Zhang, Cheng, Y., & Leow, W. K. (2013). *Dense correspondence of skull models by automatic detection of anatomical landmarks.* Paper presented at the International Conference on Computer Analysis of Images and Patterns.