School of Computing
National University of Singapore

Thesis Proposal

Reconstruction of Defective Skulls by Non-Rigid Registration with Interpolating Surface

by

Xie Shudong

under guidance of

A/Prof. Leow Wee Kheng

1^{st} November, 2016

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

Human skulls are very important target objects in craniofacial surgery planning, forensic investigation and physical anthropological study. In craniofacial surgery planning, a patient's skull is defective in one of three ways: impact injury, congenital deformity and tumor deformity. Impact injury can be caused by traffic accidents, work accidents or violence. It results in fractures and displacements of bone fragments (Figure 1.1(a)). Congenital deformity is caused by the abnormal development of infant skull, and tumor deformity is caused by the growth of tumor. They both result in the deformation of the patients' skulls (Figure 1.1(b,c)).

Given a patient's defective skull, a craniofacial surgeon's task is to operate on it to restore the patient's normal look. During pre-operative planning, a craniofacial surgeon loads the CT images of a patient's head into a software such as Brainlab [GWL99] that generates the normal shape of the defective part of a patient's skull. During intra-operative surgery, the surgeon moves the patient's fractured bones to their desired positions or replaces the deformed parts with implants using the generated normal shape as a guide.

In forensic investigation, one of the tasks is to identify a victim based on his remains. If the victim's face is unrecognizable due to decay or severe damage, then the forensic investigator has to reconstruct his face from his skull. If the victim's skull is also defective,

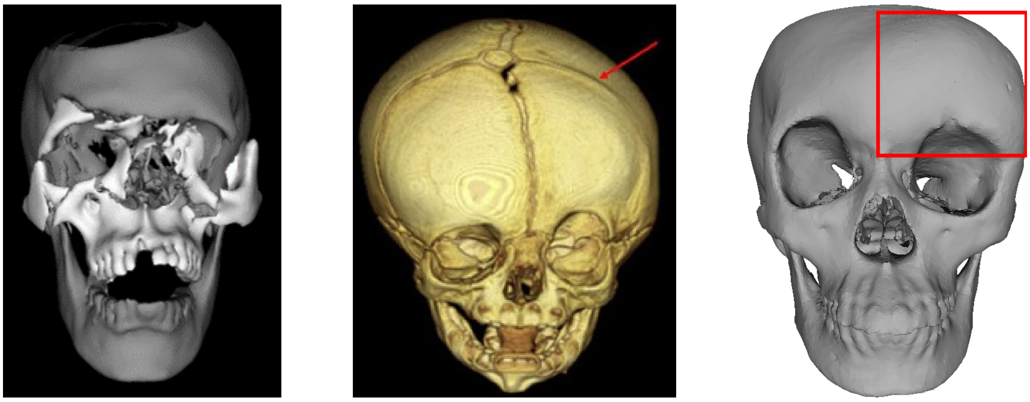

(a) (b) (c)

Figure 1.1: Skull deformities in surgery planning. (a) Fractured skull due to traffic accidents. (b) Congenitally deformed skull. (c) Deformed skull due to tumor.

which is typically caused by violence (Figure 1.2(a)), then the forensic investigator needs to reconstruct a non-defective model of the victim's skull first. He can either reconstruct the skull manually or using software tools [BSMG09]. He then reconstructs the victim's face for identification from the reconstructed skull. The reconstructed skull should be close to the normal shape of the victim's skull before injury so that the reconstructed face correctly identifies the victim.

In physical anthropology, an anthropologist studies the skulls of hominids which include modern human and ancient hominids to understand the major differences in skull shape that characterize various hominid species. The skulls studied could be actual skulls or fossilized skulls. The skulls can be defective (Figure 1.2(b)) due to violence, religious rituals, congenital deformities, and natural processes such as erosion and animal trampling. Given a defective hominid skull, an anthropologist first identifies the species it belongs to. He then reconstructs the hominid skull manually or using software tools [FdCP$^+$08] for anatomical study. Notice that in order to reconstruct an ancient hominid skull, the anthropologist must already know the normal shape of the skulls of the species.

For manual reconstruction, a reconstruction expert obtains a cast of the subject's defective skull and replaces the defective parts of the cast by clay models with normal shape. In forensic investigation, this process is based on the forensic investigator's assessment of

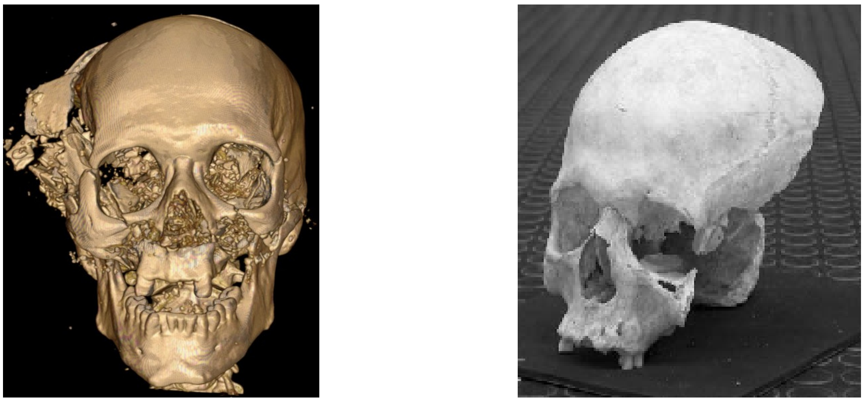

(a) (b)

Figure 1.2: Skull deformities in forensic investigation and physical anthropological study. (a) Fractured skull damaged by a bullet in forensic investigation. (b) Incomplete medieval skull excavated by archaeologists [FdCP+08].

how a complete skull should look like, whereas in physical anthropology, it is based on the shape of existing skull specimens of the same species. In other words, the forensic reconstruction is subject-specific whereas anthropological reconstruction is usually species-specific instead of subject-specific. Since the skull reconstruction process is subjective, different reconstruction experts could produce different reconstructed skulls. The accuracy of reconstruction varies with the location of the damage and decreases as the size of the defect increases [TA98].

For computer-aided reconstruction, a reconstruction expert uses software tools to generate the normal shape of the subject's defective skull. The **skull reconstruction** methods used in software tools can be grouped into four categories: symmetry-based, geometric, statistical, and bone repositioning. **Symmetry-based** methods [CSP+07, dMCP+06, LCL+03, CTS+10, FdCP+08, GWL99, LYW+11, YWML11] regard the reflection of the non-defective parts on one side of a skull about the mid-plane as an estimate of the normal shape of the defective parts on the other side. Due to the natural asymmetry of human skulls [Woo31, RRS03], the reflected model's surface would not flush with the normal parts on the defective side of the skull, producing surface discontinuities. These methods are not applicable when both sides of a skull are defective. **Geometric** methods [BS11, LYW$^+$11, YWML11, WYLL11, BSMG09, GMBW04, FdCP$^+$08, LP00] perform non-rigid registration of a reference model to fit the non-defective parts of the target model, and regard the registered reference model as the reconstructed model. The accuracy of geometric methods depends critically on the amount of corresponding points used in non-rigid registration. **Statistical** methods [Gun05, LAV09, ZLES05, Zha14, ZLC15] build a statistical model from a set of normal training skulls. Given a target skull, they compute the model parameters that best fit the non-defective parts of the target, and generate the reconstructed skull from the best-fitting model parameters. To capture all the essential variations in normal human skulls across age, race, and gender, a large number (probably hundreds) of training skulls is required. The lack of such a large training set has hindered their applications to skull reconstruction. **Bone repositioning** methods [WYLL11, YWML11, YLL12, Che13] reposition fractured bone fragments of a defective skull at their correct positions. These methods are applicable only to fractured skulls with large pieces of fractured bone fragments rather than incomplete, deformed skulls. If the fractured fragments are missing or pulverised, then the defective skull is incomplete and these methods cannot be applied. Moreover, they require the individual fractured bone fragment to be segmented. Without an automatic way to segment the bone fragments, it is very tedious and time-consuming to obtain the bone fragments manually.

1.2 Research Objectives

This thesis focuses on two research objectives. The first objective is to develop an accurate and robust method for reconstructing a normal skull given a defective target skull. Based on the goals of craniofacial surgery planning and forensic investigation, the criteria for skull reconstruction are summarized as follows:

1. **Normality**

   The reconstructed skull model should look like a normal human skull.

2. **Accuracy**

   The reconstruction should be as accurate as possible. The non-defective parts should be preserved in the reconstructed model. The defective parts should be reconstructed

(a) (b) (c) (d)

Figure 1.3: Normal skull variation. The shape of human skulls vary greatly across different identity, age, gender and racial groups.

as close to the unknown actual shape as possible. Moreover, surfaces of the reconstructed parts should flush with those of the adjacent normal parts.

3. **Robustness**

   The reconstruction process should not be misled by outliers such as metal artifacts (Figure 1.4), fractured bone fragments (Figure 1.1(a)) and deformed shape (Figure 1.1(b,c)) in a defective skull.

4. **Symmetry**

   The reconstructed skull model should be laterally symmetric. Although anthropological studies have shown that human skulls have lateral asymmetry [Woo31, RRS03], making the reconstructed model symmetric is reasonable in the absence of relevant shape information such as the actual shape of a subject's fractured skull. This principle is often used by craniofacial surgeons, forensic investigators and anthropologists in skull reconstruction.

The second objective is to develop two software tools for skull reconstruction. The first tool is the **skull segmentation tool** which is used to generate a 3D skull model from a set of CT images. The second tool is the **skull reconstruction tool** which is used to generate a normal reconstructed skull from a defective skull. These tools should satisfy the following criteria:

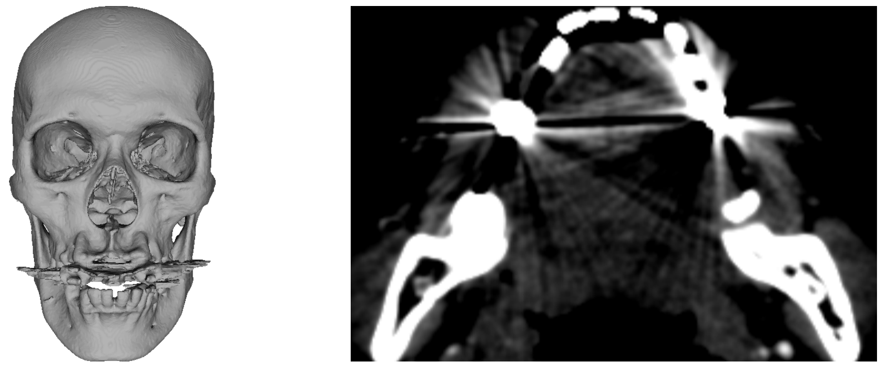

(a) (b)

Figure 1.4: Metal artifacts (a) on a 3D skull model due to (b) the scattering of X-ray on patients' metal implants.

1. **Wide Applicability**

   This thesis proposal focuses on modern human skulls which are involved in craniofacial surgery planning and forensic investigation. Thus the skull reconstruction tool should be able to deal with fractured, incomplete and deformed skulls. In addition to craniofacial surgery planning and forensic investigation, the skull reconstruction tool has the potential of reconstructing defective skulls in physical anthropology. For reconstructing modern human skulls, the tool can be used directly. For reconstructing non-human hominid skulls, the tool requires a normal skull of the same species as reference.

2. **Ease of Use**

   Both tools should be intuitive and easy to use. The logic flow should be clear without ambiguity. The tools should require as few user inputs as possible.

3. **Efficiency**

   The skull reconstruction tool should be able to generate a reconstructed model as quickly as possible (e.g., within 10 minutes) after user inputs. In this way, a user can explore different parameter settings to generate a satisfactory reconstructed model within a short amount of time.

With the software tools equipped with the reconstruction algorithm, the user should be able to generate an accurate reconstructed skull model conveniently.

1.3 Organization of Thesis Proposal

This thesis proposal is organized as follows. Chapter 2 first describes skull anatomy and characteristics. Chapter 3 then reviews existing work related to this research. After that, Chapter 4 proposes the guidelines for designing software tools and formulates the skull reconstruction problem. Subsequently, Chapter 5 and Chapter 6 present the preliminary work on the software tools and the reconstruction algorithm, respectively. Finally, Chapter 7 concludes this thesis proposal.

Chapter 2

Human Skull

2.1 Skull Anatomy

The human skull is composed of two parts: the cranium and the mandible (lower jaw bone) (Figure 2.1). The cranium consists of 8 cranial bones, 13 facial bones and 3 pairs of auditory bones that are fused together. The cranial bones include the occipital bone and other bones. The facial bones include two maxilla (upper jaw bones), two nasal bones, and other bones (Figure 2.1). The maxilla and the mandible form the jaw structure that provides the chewing function. Most of the bones are fixed except the mandible which can be rotated with respect to the jaw joints. The skull bones are bound together by bone fibres, and the connection between adjacent bones is called a suture (Figure 2.2). The cranial sutures include the sagittal suture, the lambdoidal suture, the coronal suture and other sutures.

The human orbit is a complex bony structure. It is made up of seven bones that are fused together (Figure 2.3(a)). The medial orbital wall is thin whereas the lateral wall is thick. The shape of the orbit is an irregular cone (Figure 2.3(b)).

The skulls of different people have subtle differences in local shape details. For example, the skull in Figure 2.4(a) has narrow nasal bones and large eye orbits, whereas the skull in Figure 2.4(b) has wide nasal bones and small eye orbits.

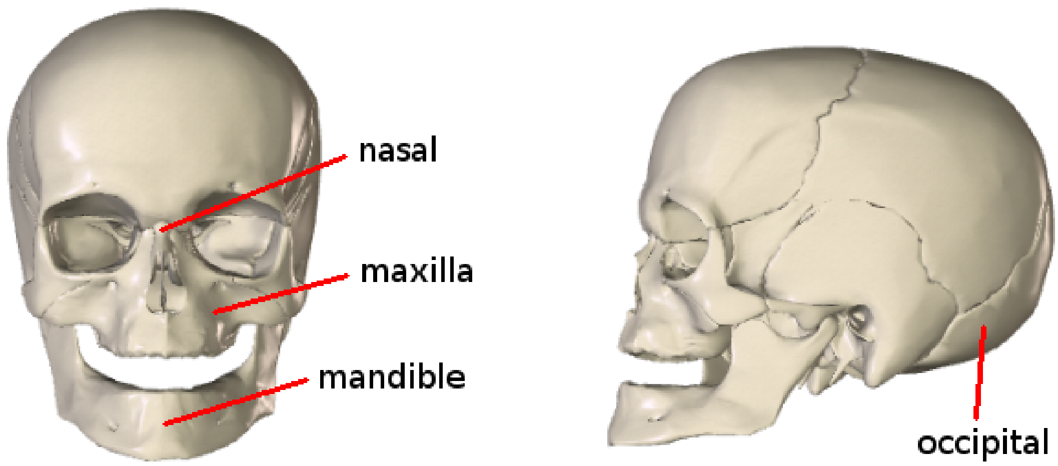

Figure 2.1: Human skull. The human skull is composed of two parts: the cranium and the mandible. The cranium consists of the occipital bone, maxilla, nasal bones and other bones.

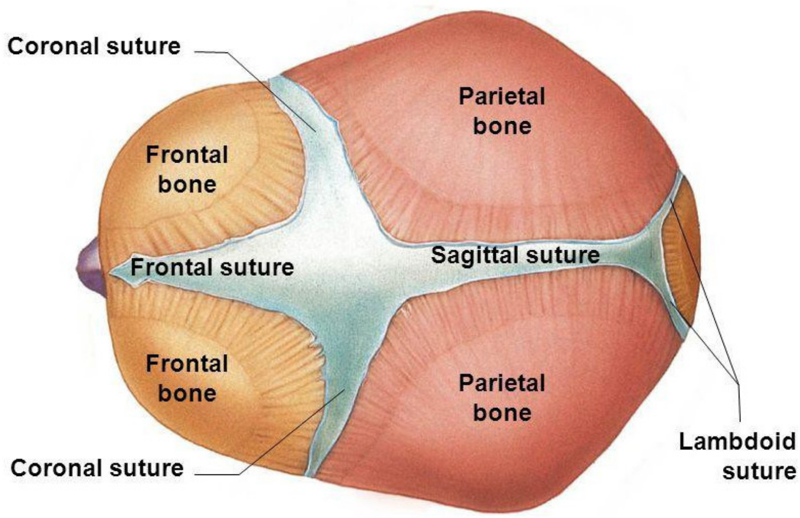

Figure 2.2: Skull sutures. The skull sutures include the sagittal suture, the lambdoidal suture, the coronal suture and other sutures.

Skull landmarks are distinctive points on the surface of a skull. For example, [İH93] defines a set of landmarks (Figure 2.5) that is used for forensic analysis of skull and cranio-facial reconstruction. Some of these landmarks are used to define two important anatomical planes called the Frankfurt plane and the mid-sagittal plane. The Frankfurt plane

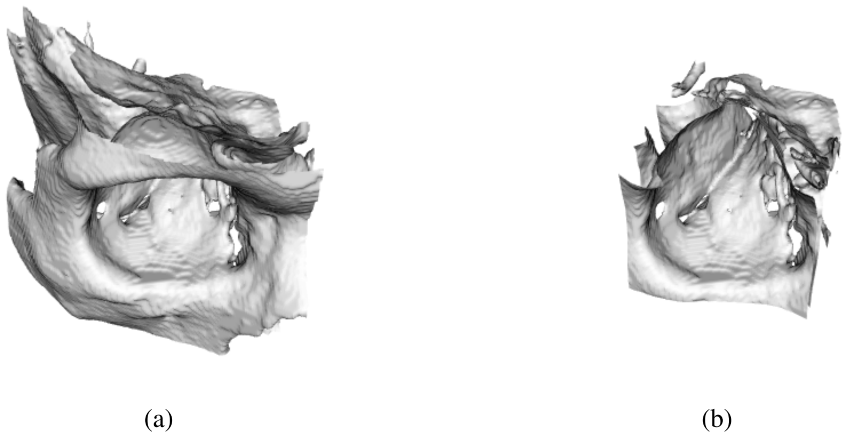

Figure 2.3: Orbit of human skull. (a) Right orbit of human skull. (b) The shape of human orbit is an irregular cone.

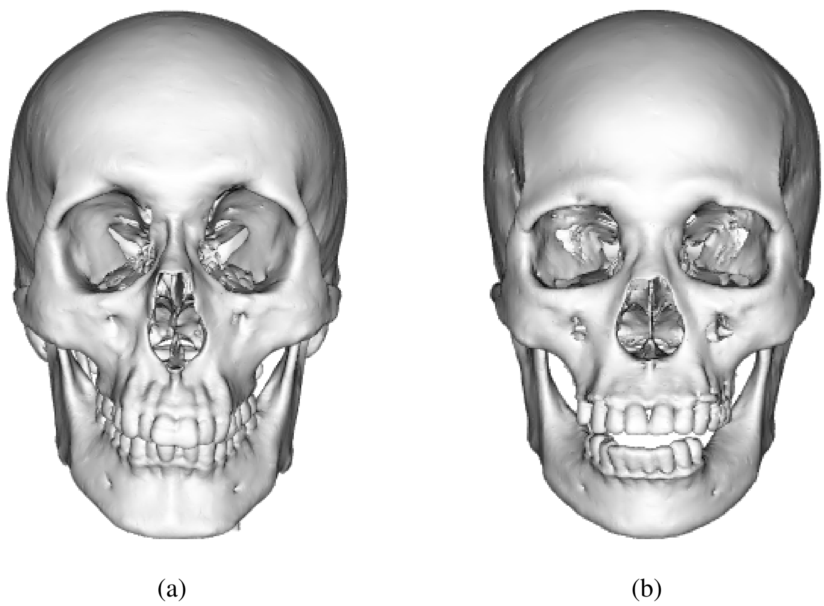

Figure 2.4: Local difference of skulls. The skull in (a) has narrow nasal bones and large eye orbits, whereas the skull in (b) has wide nasal bones and small eye orbits.

(Figure 2.6) is defined as the horizontal plane that passes through two pairs of landmarks [Mos16]:

- **Orbitale (Or)**: The left (right) orbitale is the lowest point on the margin of the left (right) orbit.

- **Porion**: The left (right) porion is the highest point of the left (right) ear canals.

The mid-sagittal plane (Figure 2.7), on the other hand, is defined as the vertical plane that passes through the midline of a skull [Mos16]. It should be orthogonal to the Frankfurt plane. There are many landmarks on the midline of a skull that can be used to define the mid-sagittal plane, and eight of them are chosen to define it:

- **Nasion (N)**: The midpoint of the suture between the frontal and the two nasal bones.
- **Nasale (Na)**: The front tip of the nasal bones at their junction with the lateral nasal cartilages.
- **Pogonion (Pog)**: The most frontal point in the midline on the mandible.
- **Menton (Me)**: The lowest point on the mandible.
- **Bregma (B)**: The point of intersection of the sagittal and coronal sutures.
- **Lambda (L)**: The point of intersection of the sagittal and lambdoidal sutures.
- **Inion (I)**: The most prominent point that sticks out on the occipital bone at the base of the skull.
- **External Occipital Crest (EOC)**: The ridge structure along the midline at the bottom of the skull.

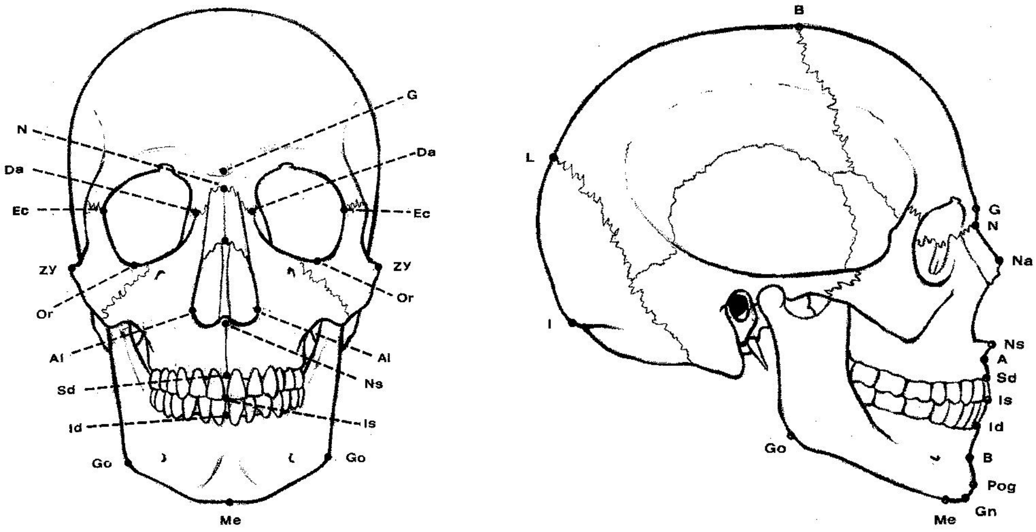

(a) (b)

Figure 2.5: Skull landmarks defined in [İH93]. (a) Frontal view. (b) Side view.

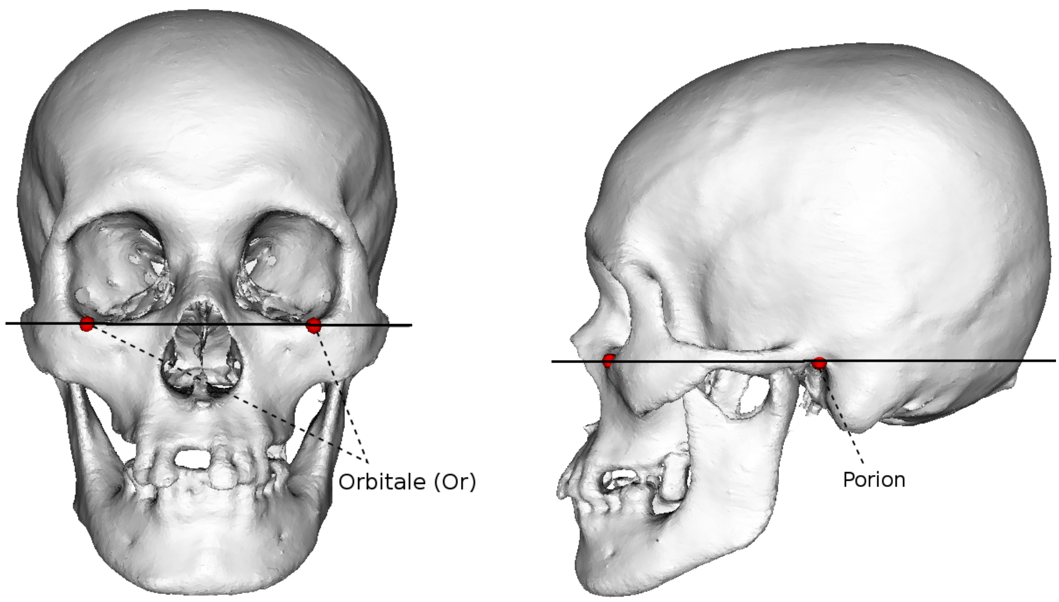

(a) (b)

Figure 2.6: The Frankfurt plane. (a) Frontal view. (b) Side view.

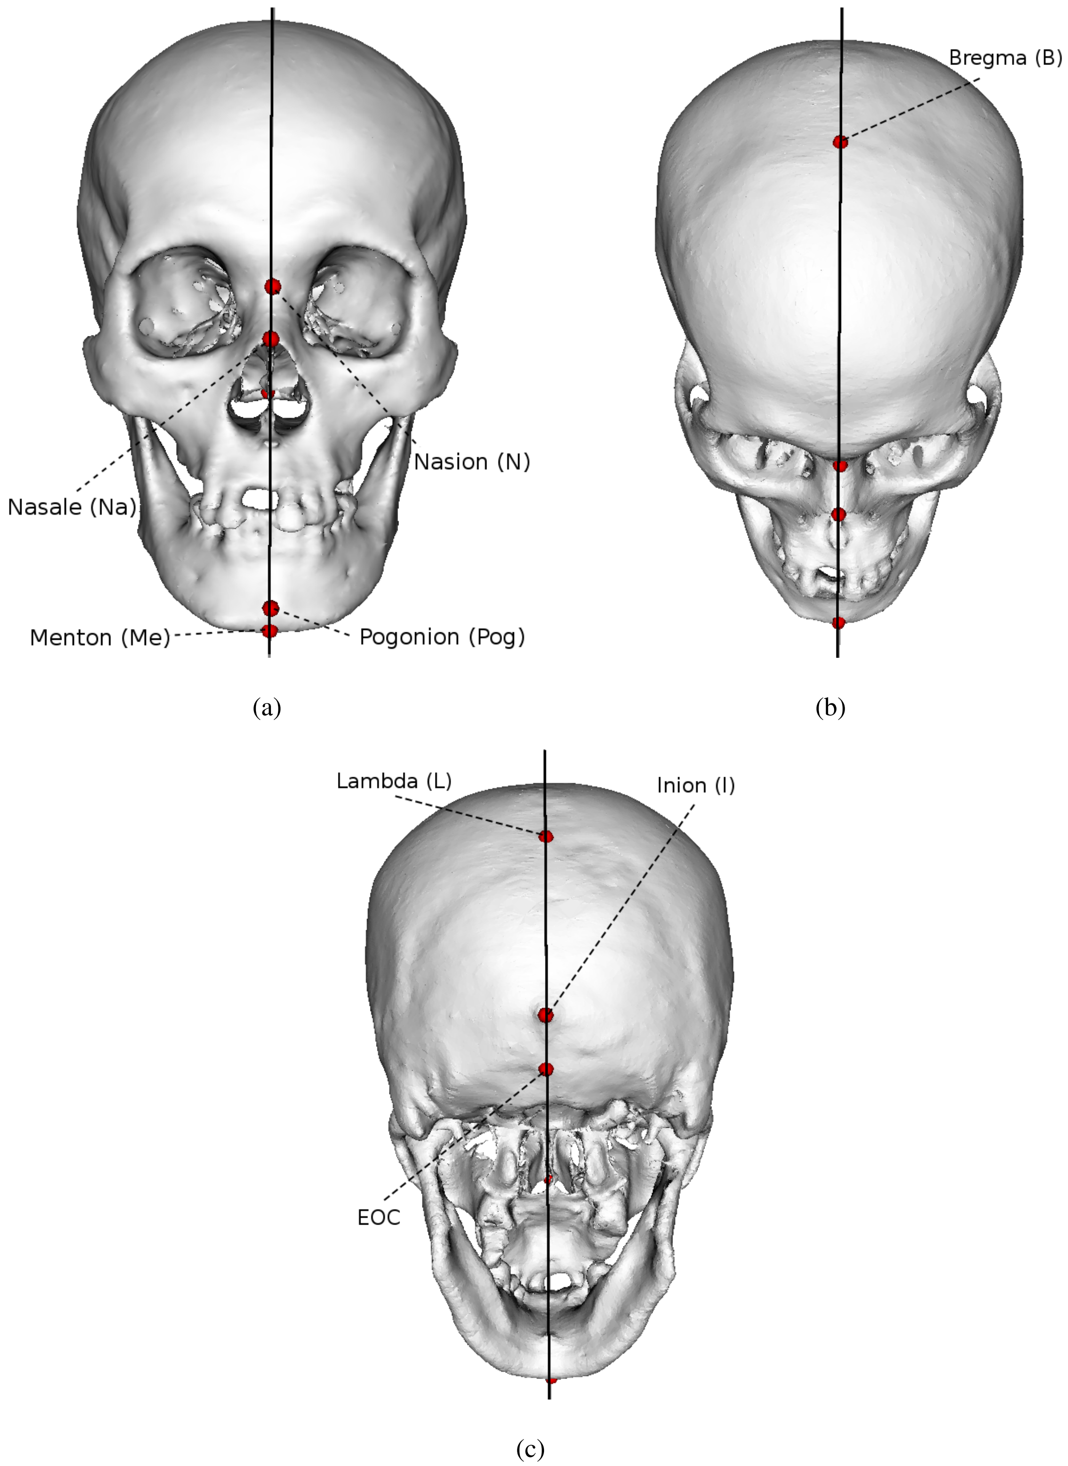

Figure 2.7: The mid-sagittal plane. (a) Frontal view. (b) Top-frontal view. (c) Bottom-rear view.

2.2 Defective Skulls and Other Complications

A skull can be defective due to fractures, congenital deformity and tumor deformity. Skull fractures can be caused by traffic accidents, work accidents or violence. Fractured bone fragments are often displaced from their original positions (Figure 2.8). A defective skull with congenital deformity is due to the abnormal development of infant skull (Figure 1.1(b)), whereas a defective skull with tumor deformity is caused by the growth of tumor (Figure 1.1(c)). Congenital deformity and tumor deformity result in deformed skull shape.

The 3D mesh model of a skull can have a spine, tubes and metal artifacts (Figure 2.8). Metal artifacts are outliers caused by the scattering of radiation from the patient's dental implants. In addition, the top and rear of the skull may be incomplete due to incomplete CT scanning. Incompleteness due to CT scanning is not a skull defect.

Figure 2.9 illustrates the orbits of a patient with skull fractures. The orbits of a skull are very thin and not clearly shown in CT images. As a result, the orbits of the 3D skull model generated from CT images may be incomplete (Figure 2.9(b,c,e)). Moreover, the orbital floor and orbital wall may be fractured (Figure 2.9(d)). Therefore, skull fractures can involve fractures of the facial bones, jaw bones as well as the orbits.

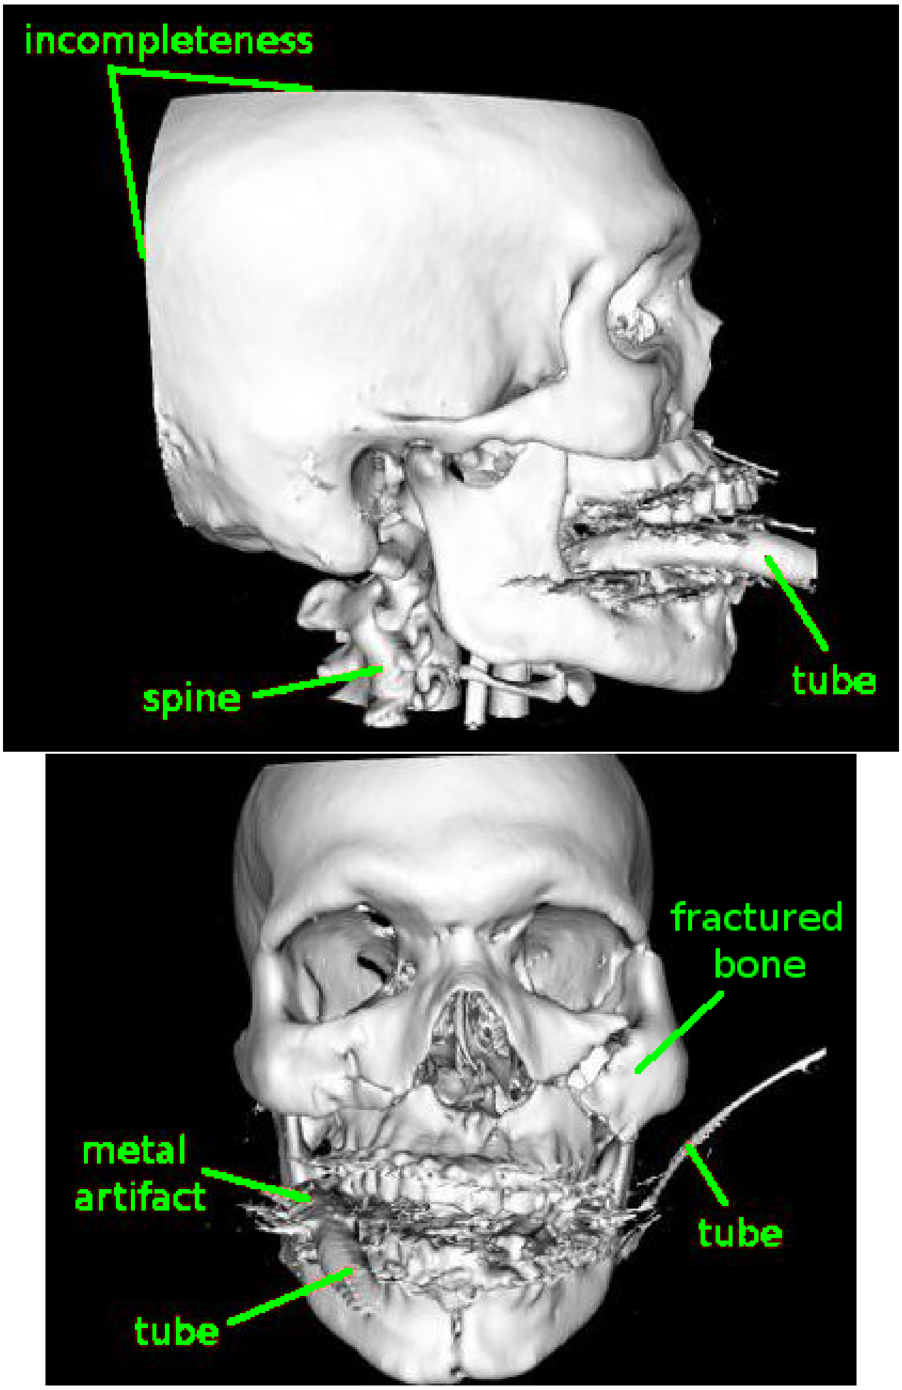

Figure 2.8: Fractures and other complications of skull model. This skull has fractured orbit and facial bone. Moreover, it has a spine, tubes, and metal artifacts, and it is not completely scanned in CT.

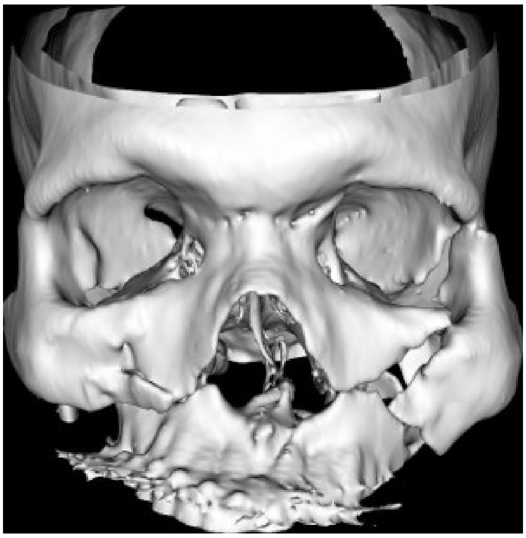

(a)

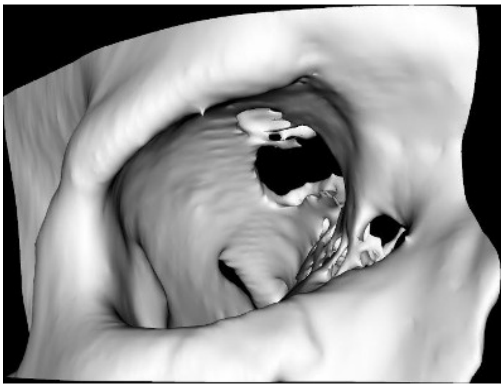

(b)

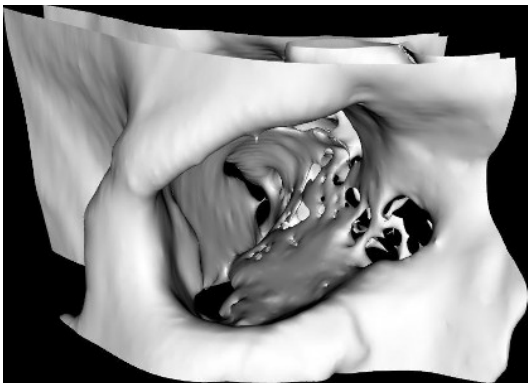

(c)

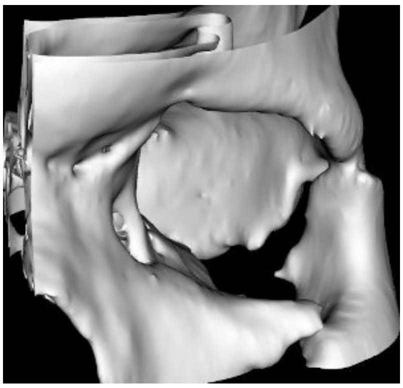

(d)

(e)

Figure 2.9: Orbital complications. (a) The orbits of a fractured skull. The right orbit is intact, whereas the left orbit is fractured. The holes in (b) the right orbital ceiling, (c) the right orbital floor and (e) the left orbital wall are due to insufficient resolution in CT scan. (d) The left orbital floor is fractured.

Chapter 3

Existing Methods

As discussed in Section 1.2, one of the objectives is to develop an accurate and robust method for skull reconstruction. Thus, this chapter first reviews existing reconstruction approaches (Section 3.1). Among them, the geometric reconstruction approach uses non-rigid registration methods. Hence, this chapter also reviews non-rigid registration methods (Section 3.2). In general, non-rigid registration methods can cause self-interference of registered surfaces. Therefore, this chapter finally reviews techniques for handling self-interference (Section 3.3).

3.1 Skull Reconstruction Methods

Existing approaches for generating the normal shape of a defective skull can be grouped into four broad categories: symmetry-based reconstruction, geometric reconstruction, statistical reconstruction, and bone repositioning. The following subsections review existing methods in each of these categories.

3.1.1 Symmetry-Based Reconstruction

Symmetry-based reconstruction methods [CSP$^+$07, dMCP$^+$06, LCL$^+$03, CTS$^+$10, GWL99, LYW$^+$11] are based on the fact that human skulls are approximately left-right symmetric.

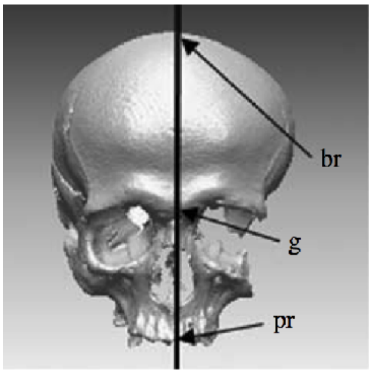

(a)

(b)

Figure 3.1: Symmetry-based reconstruction (figures collected from [FdCP$^+$08]). (a) Incomplete skull model. (b) Reconstructed skull model.

These methods use the reflection of the non-defective part on one side of the skull about the mid-sagittal plane to replace the defective part on the other side (Figure 3.1). They require the non-defective parts to be specified by the user. The mid-sagittal plane can be identified semi-automatically [CSP$^+$07, CTS$^+$10] or automatically [dMCP$^+$06, LYW$^+$11, CLL12].

Symmetry-based methods are easy to apply. They are commonly used for implant design [LCL$^+$03] and in surgery planning systems [CSP$^+$07, CTS$^+$10, GWL99]. However, due to the natural asymmetry of human skulls [RRS03, Woo31], the reflected model's surface produced by symmetry-based methods may not flush with the normal parts on the defective side of the skull, causing surface discontinuities. Moreover, symmetry-based methods cannot be applied directly to a defective skull with bilateral defects, i.e., defects on both sides of a skull.

3.1.2 Geometric Reconstruction

Geometric reconstruction methods [BS11, LYW$^+$11, YWML11, WYLL11, BSMG09, GMBW04, FdCP$^+$08, LP00] non-rigidly register a reference model to the non-defective parts of the target model, and generate the reconstructed model from the registered reference model. They do not require symmetry information and thus can be applied to target skulls with bilateral defects. These methods use various non-rigid registration algorithms such as thin-plate spline [BS11, BSMG09, GMBW04], free-form deformation [FdCP^{+}08], and shape-preserving deformation [LYW^{+}11, YWML11, WYLL11]. They change the shape of the reference model appropriately to match the shape of the non-defective parts of the target as closely as possible. The accuracy of these methods depends critically on the amount of corresponding points used in non-rigid registration. Methods that use a small set of manually marked landmarks [BSMG09, DZS^{+}11, LP00, RB02] cannot achieve high accuracy. Thus, some methods automatically detect more corresponding points [TBK^{+}05, ZCL13].

Among the geometric reconstruction methods, the methods of [BSMG09, LP00] simply regard the non-rigidly registered reference model as an approximation of the reconstruction of the target model. In general, the registered reference does not match the non-defective parts of the target perfectly. Thus, regarding the registered reference as the reconstruction introduces undesirable errors in the non-defective parts. On the other hand, the methods in [BS11, LYW^{+}11, YWML11, WYLL11, GMBW04, FdCP^{+}08] generate a reconstructed model by replacing the defective parts of the target skull by their corresponding parts on the non-rigidly registered reference model. These methods keep the non-defective parts of the target skull unchanged. However, surfaces of the replaced parts may not flush with the non-defective parts of the target because the non-rigidly registered reference does not match the non-defective parts of the target exactly.

Instead of using a generic reference model, some methods [BS11, FdCP^{+}08] use the reflection of the non-defective parts of the target as the reference model for geometric reconstruction. These methods further preserve the shape information because the replaced parts are indigenous to the target. However, they are applicable only when one side of the target is defective.

3.1.3 Statistical Reconstruction

Statistical reconstruction methods [Gun05, LAV09, ZLES05, Zha14, ZLC15] typically apply active shape model (ASM) to build a statistical model from a set of normal training skulls. To construct the active shape model using PCA [Zha14, ZLC15, ZLES05], they need to resample the training skulls to obtain training models with the same number of vertices and mesh connectivity. This is accomplished by non-rigidly registering a reference model to each of the normal skulls, and regarding each registered reference model as one of the training models. Then, given a defective target model, they compute the model parameters that best fit the non-defective parts of the target, and generate the reconstructed skull from the best-fitting model parameters.

Standard ASM requires the training models to have the same number of vertices and mesh connectivity. Normal skulls obtained in medical practices may be incomplete because they are not always completely scanned (Section 2.2). Therefore, they cannot be used to build ASM, which is a waste of valuable information. To overcome this issue, Lüthi et al. [LAV09] propose a statistical shape model based on probabilistic PCA (PPCA) [Row98, TB99]. This statistical shape model can be built from incomplete training skulls. To begin with, a complete and normal reference model is segmented into parts based on anatomical structure. The reference model is then non-rigidly registered to fit the shape of each model in the training set to detect the outliers and missing data. The surface of the registered reference model that corresponds to the non-missing parts of the target is regarded as a training model, and the training models are used to build the statistical shape model.

To capture all the essential variations in normal human skulls across age, race, and gender, a large number of (for example, $\gg 100$) training skulls are required. The lack of such a large training set has hindered the applications of statistical methods to skull reconstruction. Moreover, generating the reconstructed model by adjusting shape parameters is an indirect method of changing the reconstructed model's shape. It cannot guarantee perfect match of the reconstructed surface and the non-defective parts of the target. Therefore, statistical reconstruction methods typically introduce undesirable errors in the non-defective parts of the reconstructed model.

3.1.4 Bone Repositioning

Bone repositioning methods [WYLL11, YWML11, YLL12, Che13] reconstruct a skull by repositioning fractured bone fragments of a defective skull at their correct positions. The bone fragments of a defective skull can be rotated and displaced. Thus, it is necessary to determine the spatial relationships among the bone fragments and then join adjacent bone fragments together, like solving the jigsaw puzzle. To determine the rough spatial relationships among the bone fragments, some methods [WYLL11, YWML11, YLL12] establish correspondence between each bone fragment and a reference model by matching their surface features, and rigidly register the bone fragments to the reference based on their correspondence. They then join adjacent bone fragments together by matching their fractured surfaces. This matching process requires features of the fractured surfaces to be well preserved. For cases where the target skull is defective due to impact injuries, the fractured surfaces of bone fragments may abrade each other, damaging the fractured surfaces. In this case, these methods cannot accurately match the fractured surfaces of bone fragments.

Alternatively, the method of Cheng [Che13] does not need to match the fractured surfaces of fragments. Instead, it repositions the bone fragments by rigidly registering them to an estimate of the ideal surface of the target skull. The ideal surface is estimated by non-rigidly registering a reference model to the non-defective parts of the target in the same manner as geometric reconstruction.

Bone repositioning approach reconstructs a skull by repositioning its bone fragments. Thus, it is applicable only to fractured skulls rather than incomplete and deformed ones. It requires every individual fractured bone fragment to be segmented from CT images. Without an automatic way to segment the bone fragments, it is very tedious and time-consuming to obtain the bone fragments manually.

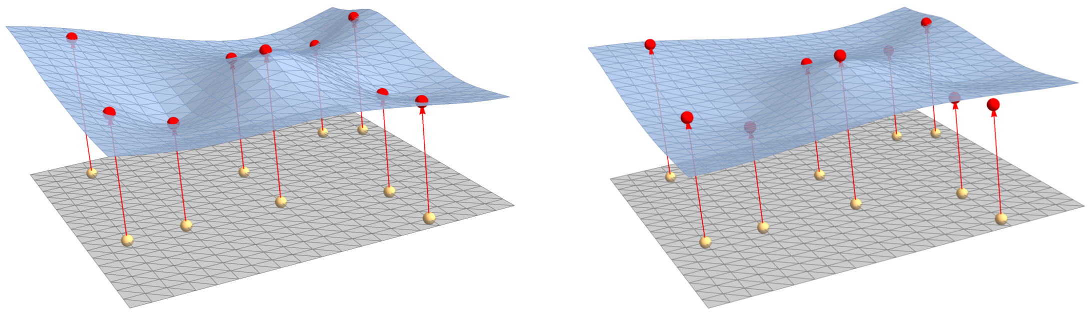

(a) (b)

Figure 3.2: Interpolating surfaces vs. approximating surfaces. (a) An interpolating surface passes through landmark points exactly, whereas (b) an approximating surface minimizes its average distance to the landmarks.

3.2 Non-Rigid Registration

Non-rigid registration methods can be grouped into two broad categories based on the goal of registration: interpolation and approximation. The fundamental difference between these two approaches is that the system of linear equations of interpolation yields a unique exact solution, whereas that of approximation yields a least-square solution.

3.2.1 Interpolating Surfaces

Non-rigid registration methods that produce an **interpolating surface** fit the reference surface to pass through the corresponding target points. They regard the positional correspondence as **hard constraints**, and thus, their registered surfaces have zero error with respect to the corresponding target points (Figure 3.2(a)).

There are three approaches for performing non-rigid registration with interpolating surfaces: free-form deformation, thin-plate spline and Laplacian deformation. Free-form deformation (FFD) [SP86] deforms a mesh by transforming the underlying 3D space enclosing it. It encloses the mesh within a 3D space represented by a rectangular grid. The 3D space is divided into parallelepiped regions whose vertices are the control points (Figure 3.3). When the control points are moved to new positions, FFD transforms the points

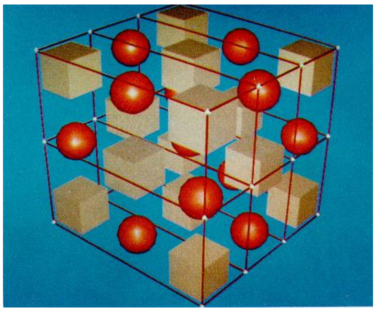

Figure 3.3: Free-form deformation [SP86]. FFD encloses 3D objects within a rectangular grid. The grid is divided into parallelepiped regions whose vertices are the control points.

within the underlying 3D space by interpolating with Bernstein polynomial. The control points used by FFD are not necessarily on the mesh, which makes it difficult to manipulate the control points to ensure correct interpolation of surface points.

To overcome the shortcoming of FFD, Hsu et al. [HHK92] present direct manipulation free-form deformation that allows the user to specify the target interpolating points instead of the control points. The method represents the displacements of control points by the displacements of mesh vertices, allowing a 3D mesh to be deformed by moving its mesh vertices to their desired locations directly.

Among the interpolating methods, thin-plate spline (TPS) [Boo89] is the most popular for reconstruction of skulls [DZS+11, LP00, RB02, TBK+05, ZCL13]. The process of TPS warping is analogous to the bending of a thin metal plate. TPS transforms the surface points of a mesh using a smooth mapping function that minimizes the bending energy. The mapping function consists of an affine transformation and a nonlinear warping defined as the weighted sum of thin-plate spline functions. The parameters of the mapping function are solved through a system of linear equations that yields a unique solution. The size of the data matrix of this linear system is proportional to the square of the number of landmarks. Therefore, the time and space complexity of TPS increase rapidly with increasing number of corresponding landmarks.

Laplacian deformation [SCOL+04, MYF06] deforms a reference model to fit target points by preserving the curvatures and surface normals of mesh vertices. It moves the surface points with correspondence to their desired positions exactly, and transforms the other surface points in such a way that the surface of the mesh after deformation preserves the curvatures and surface normals of mesh vertices. Unlike TPS, the system of linear equations of Laplacian deformation captures only the mesh vertices without correspondence. The size of its data matrix is proportional to the square of the number of mesh vertices without correspondence. Therefore, the size of data matrix as well as the time and space complexity decreases with increasing number of correspondence. Laplacian deformation has not been used for skull reconstruction.

3.2.2 Approximating Surfaces

Non-rigid registration methods that produce an **approximating surface** fit a reference surface to the target by minimizing the average distance between corresponding reference and target surfaces. They regard the positional correspondence as **soft constraints**, and their registered surfaces have non-zero distance or error to the target surfaces (Figure 3.2(b)).

There are two general approaches for performing non-rigid registration with approximating surfaces: non-rigid ICP [ARV07, HMS12, BSB14] and approximating TPS [RSS$^+$01]. Amberg et al. [ARV07] extends iterative closest point (ICP) to non-rigid ICP by applying **piecewise affine transformation** between the reference and the target models. Similar to ICP, this method first establishes correspondence between two models by associating every reference vertex with its closest point on the target. Based on the correspondence, it computes an affine transformation per reference vertex such that every reference vertex is moved towards its corresponding target point as much as possible. The difference of transformation matrices between adjacent vertices is kept as small as possible. The process of establishing correspondence and applying affine transformation is iterated until convergence is achieved. Non-rigid ICP is applicable to incomplete target models, but it is not robust to outliers.

There are several variants of non-rigid ICP. Hontani et al. [HMS12] proposes a robust non-rigid ICP by incorporating a statistical model into the non-rigid ICP. The statistical shape model is trained by a set of normal models and can be used to discriminate between inliers and outliers. Thus, their proposed non-rigid ICP can achieve robustness by weakening the transformations of outliers. Bonarrigo et al. [BSB14] presents a non-rigid ICP method that applies **piece-wise rigid transformation** between the reference and the target models. Unlike the non-rigid ICP of [ARV07], their method does not establish correspondence for every reference vertex. Instead, it discretizes the reference into a set of overlapping patches, and establishes correspondence between each patch and the target model. They compute a rigid transformation per patch, and the rigid transformation matrix of every vertex is interpolated from the transformation matrices of its adjacent patches. Since the rigid transformations are computed for each patch instead of each vertex, the size of data matrix of this method is smaller than that of [ARV07], resulting in more efficient computation.

Rohr et al. [RSS$^+$01] modifies the interpolating thin-plate spline into an approximating thin-plate spline. They add additional terms into the original system of linear equations of TPS. As a result, the function parameters obtained from this modified system of linear equations are different from those of interpolating thin-plate spline, leading to an approximating mapping function.

3.3 Handling Surface Self-Intersection and Flipping

Interpolating methods can produce **flipped surfaces** when there are conflicts in the hard constraints. Flipped surfaces cause severe distortion of surface shape (Fig. 3.4(e, f)), and are very difficult to remove. Note that surface flipping is a direct consequence of surface interpolation with conflicting hard constraints. Imposing surface smoothness constraint by energy minimization, such as TPS, cannot remove surface flipping (Fig. 3.4(f)). In contrast, approximating surfaces can avoid surface flipping because they regard the correspondence as soft constraints and are allowed to ignore conflicting constraints. Their shortcoming is the non-zero reconstruction error of the non-defective parts.

There are two general approaches for handling surface flipping and self-intersection: (1) detection and resolution, and (2) avoidance. The method of McInerney and Terzopoulos [MT99] detects self-intersections by examining the deformation result, and resolves

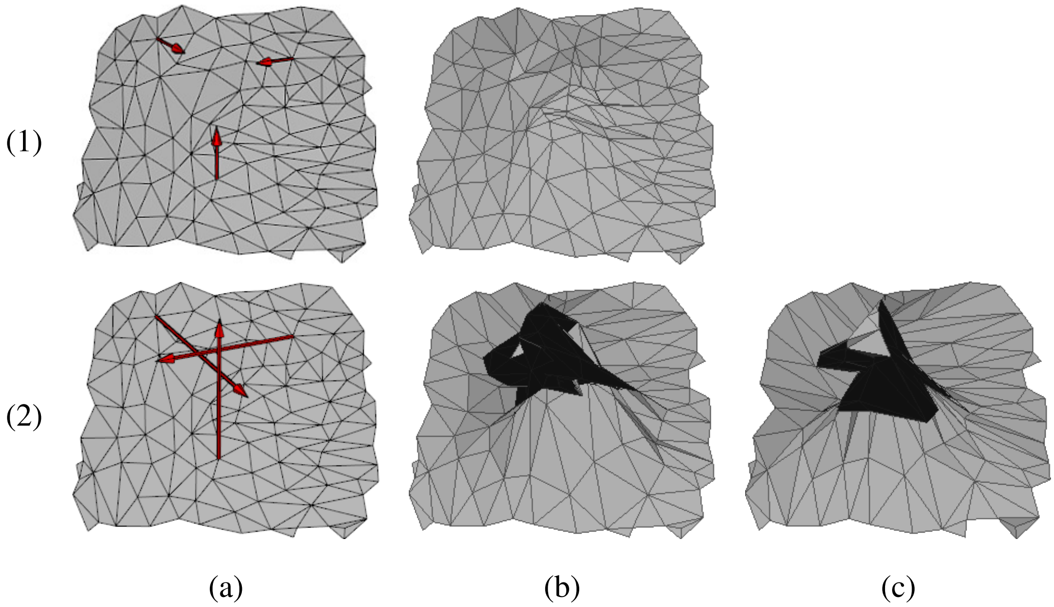

Figure 3.4: Surface flipping. Red arrows indicate the target locations of selected mesh vertices. (1a) Non-crossing correspondence vectors (arrows) produce (1b) no surface flipping. (2a) Crossing correspondence vectors cause (2b) Laplacian deformation and (2c) TPS to produce flipped and distorted surfaces even when they do not intersect. Black regions are surface patches that have flipped.

self-intersections by rolling back the mesh to the state before deformation and imposing repulsive forces to keep the potentially intersecting surfaces apart. The method of Lachaud and Montanvert [LM98] imposes proximity conditions between mesh vertices and detect violations of proximity conditions, whereas the methods in [DM01, JSC04, ZBH07] detect self-intersections through collision detection, and resolve self-intersections by remeshing.

The methods in [CL00, HF98] avoid self-intersection by imposing injectivity (one-to-one) condition on free-form deformation function. The injectivity condition confines the free-form deformation of mesh to regions that do not have self-intersection. Khan et al. [KAB$^+$05] and Zhuang et al. [ZRA$^+$08] apply diffeomorphic deformation function. A diffeomorphic function and its inverse are both one-to-one and smooth, and it preserves the topology of the mesh after deformation, thus avoiding self-intersection. The computation of a diffeomorphic function is very expensive.

Ding et al. [DYLV09] devise an ingenious quadrilateral mesh (Figure 3.5) that permits

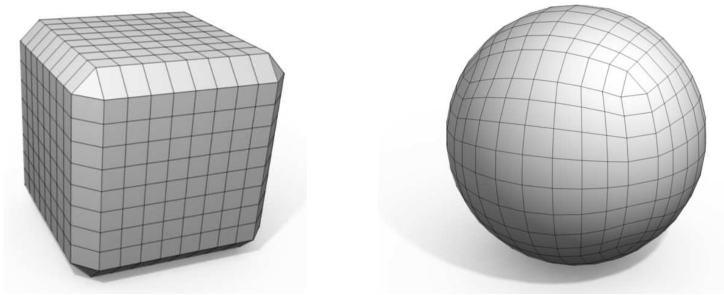

(a) (b)

Figure 3.5: Quadrilateral mesh of [Din10]. (a) The cubical quadrilateral mesh is defined by three groups of closed contours where each contour of a group is orthogonal to the other contours of a different group. The cubical quadrilateral mesh is warped into (b) a spherical quadrilateral mesh.

easy detection of possible flippings of mesh edges. The quadrilateral mesh is defined by three groups of closed contours where each contour of a group is orthogonal to the other contours of a different group (Figure 3.5). Each mesh vertex is an intersection of two contours from different groups, and the mesh vertices along any closed contour are linearly ordered. The linear ordering of vertices on a closed contour makes it simple to detect possible flippings. These flippings are removed from the constraint set for mesh deformation, thus avoiding flippings. Unfortunately, it is non-trivial to convert a triangular mesh to the special quadrilateral mesh, limiting the applicability of this method.

3.4 Summary

Existing skull reconstruction methods are summarized in Table 3.1. Among them, geometric reconstruction is the most promising because it is fast, it needs just a single reference model, and it is applicable to any defective skull. Thus, the skull reconstruction algorithm proposed in this thesis proposal (Chapter 6) adopts the geometric reconstruction approach. To achieve high reconstruction accuracy, the proposed method will use non-rigid registration with interpolating surfaces instead of approximating surfaces because it can achieve zero distance or error with respect to the reference surface points with correspondence. Thus, it can achieve highly accurate reconstruction with a very dense set of correspondence. Among the two common interpolating algorithms, Laplacian deformation is chosen over the more popular TPS because Laplacian deformation runs faster as the number of correspondence increases whereas TPS runs slower. To handle possible surface flipping of interpolating algorithms, flip avoidance method is selected because it is more efficient than the detection and resolution method.

Table 3.1: Comparison of existing skull reconstruction methods.

In the third column, RI, RA, and RR refer to local replacement with interpolating surface, approximating surfaces, and reflection, respectively.

Approach	References	Type	No. of refs	Bilateral defects	Local replacement	Non-defective parts unchanged	Flushed surfaces	Non-rigid registration	Correspon-dence	Efficiency
Symmetry-based reconstruction	[CSP$^+$07, dMCP$^+$06, LCL$^+$03, CTS$^+$10, GWL99, LYW$^+$11]	–	1	no	yes	yes	no	–	–	fast
Geometric reconstruction	[BSMG09, LP00]	generic	1	yes	no	no	yes	interpolating	sparse	fast
	[GMBW04]	RI	1	yes	yes	yes	no	interpolating	sparse	fast
	[LYW$^+$11, YWML11, WYLL11]	RA	1	yes	yes	yes	no	approximating	dense	fast
	[BS11, FdCP$^+$08]	RR	1	no	yes	yes	no	interpolating	sparse	fast
Statistical reconstruction	[Gun05, LAV09, ZLES05, Zha14, ZLC15]	–	many	yes	no	no	yes	–	–	fast
Bone Repositioning	[WYLL11, YWML11, YLL12, Che13]	–	1	yes	yes	yes	may be	–	–	slow
Proposed		–	1	yes	effectively	yes	yes	interpolating	very dense	fast

Chapter 4

Proposed Thesis Project

The proposed research project consists of the development of software tools and a skull reconstruction algorithm. The software tools include a skull segmentation tool and a skull reconstruction tool. The skull segmentation tool (Section 4.1) is used to generate a 3D mesh model from a set of CT images. It is needed when the input is a set of CT images instead of a mesh model of skull. The skull reconstruction tool (Section 4.1) is used to generate a normal reconstructed skull model from a defective skull model. Its core is the skull reconstruction algorithm (Section 4.2) that produces the reconstructed model given the defective model and other information provided by the tool.

4.1 Requirements of Software Tools

The skull segmentation tool is used to generate a 3D mesh model from a set of CT images. The input of the skull segmentation tool is a set of CT images of a defective skull. The output of the tool is a 3D mesh model of the skull with possibly irrelevant parts such as metal artifacts, a spine and tubes. The resolution of the output should not be too high or too low. If the resolution is too high, it would need more time and memory for the skull reconstruction tool to produce the output model. If the resolution is too low, the output model will lose important local shape information. Moreover, the skull segmentation tool should be intuitive to the users and should involve as few user inputs as possible.

The skull reconstruction tool is used to generate a normal reconstructed skull model from a defective skull model. It should be intuitive, easy to use and involve as few user inputs as possible. The inputs of the skull reconstruction tool include:

1. **Target Skull Model**

   The target skull model is the object to reconstruct from. It can be either fractured, deformed or incomplete. It may contain possibly irrelevant parts such as a spine and tubes.

2. **Normal Reference Model**

   A normal reference skull model is needed to provide normal shape information that can be used to guide the reconstruction. For example, the shape information of the reference can be used to estimate the missing parts of the target. The reference model should be aligned with the target model so that its shape information can be utilized.

3. **Corresponding Landmarks**

   Skull landmarks on the reference model (Figure 4.1) and their corresponding landmarks on the target model are important and reliable information for skull reconstruction. They can be used for spatial alignment and building dense correspondence between two models which are involved in the skull reconstruction process. The reference landmarks are placed on the reference model in advance. The target landmarks are placed by the user on the non-defective parts of the target model. Defective parts can be displaced severely, and landmarks placed on defective parts may not be reliable. Therefore, the set of target landmarks is a subset of the reference landmarks.

4. **Defective Parts of Target**

   The defective parts of the target model, such as deformed bones, fractured bone fragments and missing parts, should be used cautiously during the reconstruction process. They may provide wrong shape information that can mislead the reconstruction algorithm. The user needs to mark the defective parts of the target.

The output of the tool is a normal skull model reconstructed from the input. The reconstructed model contains outer surfaces only because in applications such as craniofacial

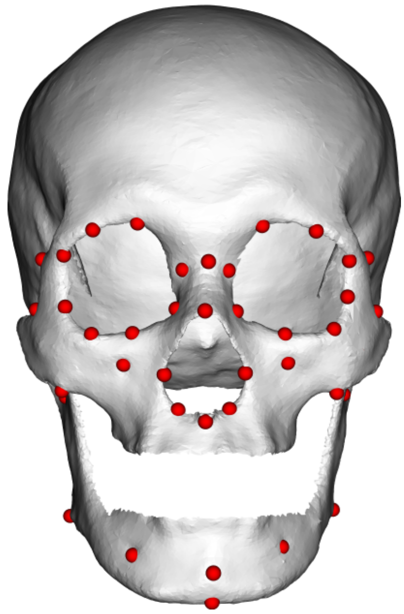

Figure 4.1: Reference landmarks.

surgery and forensic investigation, only outer surfaces are important as they define the shape of the human face. It is noted that in craniofacial surgery, the orbits can be fractured and need to be reconstructed. Due to the complexity of orbits (Section 2.2), the reconstruction of orbits will be studied in the continuing work (Section 6.7). More requirements on the output are discussed in Section 4.2.

4.2 Problem Formulation for Skull Reconstruction

The inputs of the skull reconstruction problem include a normal reference model F, a defective target model T, and their landmark sets L_F (Figure 4.1) and L_T. The input mesh models are the outer surfaces of the respective skull models. Let D denote the set of surface points on the defective parts of T. Its complement $\bar{D}$ is the set of surface points on the non-defective parts of T.

The skull reconstruction problem is to find an appropriate, possibly non-linear, transformation g such that the output skull model $R = g(T)$ is non-defective. The function g maps the target T to the reconstructed model R. It also maps the target landmarks in L_T

to the landmarks in L_R of the reconstructed model. Let π_M define the mid-sagittal plane of R that passes through the MSP landmarks in L_R. The mapping function g and, thus, the reconstruction should satisfy the following constraints:

1. **Accuracy of Non-Defective Parts**

   The non-defective parts of the target should be preserved in the reconstructed model, which means that for any surface point $\mathbf{p} \in \bar{D}$, $g(\mathbf{p}) = \mathbf{p}$.

2. **Matching of Reference Shape**

   The reconstruction of defective parts should be close to their corresponding parts on the reference model. For a surface point $\mathbf{p}$ on T, let $f(\mathbf{p})$ denote the anatomically corresponding point of $\mathbf{p}$ on the reference F. Then for any surface point $\mathbf{p} \in D$, $\|g(\mathbf{p}) - f(\mathbf{p})\|$ should be small. The mapping function f is generally unknown and needs to be determined.

3. **Surface Continuity**

   The surface of the reconstructed model should flush with the reconstruction of non-defective parts in R. Let $s(\mathbf{p})$ denote the surface characteristics such as curvature and surface normal at a surface point $\mathbf{p}$. Then for any surface point $\mathbf{p}$ on the reconstruction $g(D)$ and any surface point $\mathbf{q}$ on the boundary of the non-defective part $\bar{D}$ that is adjacent to D,

   $$\lim_{\mathbf{p} \to \mathbf{q}} s(\mathbf{p}) = s(\mathbf{q}). \tag{4.1}$$

4. **Symmetry**

   The reconstructed model R should be laterally symmetric. Let $\mathbf{r}$ and $\mathbf{r}'$ denote any pair of anatomically corresponding points on the left and right side of the reconstructed skull R. Symmetry means the distance from $\mathbf{r}$ and $\mathbf{r}'$ to the mid-sagittal plane π_M of R are the same and the vector $\mathbf{r} - \mathbf{r}'$ is parallel to the normal $\mathbf{n}$ of π_M:

   $$d(\pi_M, \mathbf{r}) = d(\pi_M, \mathbf{r}') \quad \text{and} \quad \frac{|\mathbf{n}^\top(\mathbf{r} - \mathbf{r}')|}{\|\mathbf{r} - \mathbf{r}'\|} = 1. \tag{4.2}$$

   It is noted that human skulls are not exactly laterally symmetric. Nevertheless, this condition can be useful when the other conditions are insufficient to guide the reconstruction algorithm.

5. **Normality**

   The reconstructed model R should look like a normal human skull. This condition is automatically implied by the first three conditions provided that the reference F is close to the target T. The first two conditions indicate that the reconstruction of both the non-defective parts and defective parts are normal. The third condition implies that the surfaces of the reconstruction of the defective parts and the adjacent non-defective parts are flushed. Thus the reconstructed model R is normal. Consider the ideal case where the reference model F is exactly the same as the ground truth T^*. Then by setting $g(\mathbf{p}) = f(\mathbf{p})$ for all surface point p on the target T, the first two conditions will be satisfied, giving $R = T^*$. And thus, the third condition about surface continuity is satisfied as well. That is, the problem formulation includes the ideal case as a possible solution. In practice, the ground truth T^* is not available. Therefore, the reference F should be chosen such that it is as close to the target T as possible.

Chapter 5

Software Tools

Two software tools have been developed based on the guidelines proposed in Section 4.1: a skull segmentation tool (Section 5.1) and a skull reconstruction tool (Section 5.2). The software packages chosen to develop the tools are Qt and visualization toolkit (VTK) because they are free (open source) and platform independent. Qt is used to design the graphical user interface, whereas VTK is mainly used for visualization and simple processing of DICOM images and 3D mesh models.

5.1 Skull Segmentation Tool

The skull segmentation tool is used to generate a 3D mesh model from a set of CT images which are usually stored in DICOM format. There are 5 main stages in the segmentation tool:

1. **Import CT Images**
   Use the import function to import CT images. The imported CT images are displayed in three standard anatomical views that are familiar to surgeons. These three views are the transverse view, the coronal view and the sagittal view (Figure 5.1).

2. **Determine Volume of Interest (VOI)**
   Adjust the VOI lines in each of the three views to define the VOI (Figure 5.2(a)). Form as small a volume as possible that includes the skull but excludes the neck and

spine. Then, apply the VOI cropping tool to remove the volumes outside the VOI (Figure 5.2(b)).

3. **Adjust Intensity Threshold**

   Skull segmentation is achieved by filtering out voxels with intensity values no greater than a given threshold. This method is simple, efficient, and skull segmentation can be performed within a short amount of time. Use the thresholding function of the tool to adjust the intensity threshold so that most of soft tissues disappear in three views (Figure 5.3). It is noted that the intensity threshold should not be too low; otherwise the skull model to be constructed will contain much soft tissues. Also, the intensity threshold should not be too high; otherwise the skull model to be constructed will be incomplete.

4. **Generate Mesh Model**

   Use the mesh generating function of the tool to generate the mesh model from the intensity values greater than the threshold (Figure 5.4). Then, simplify and smooth the mesh if necessary.

5. **Save Mesh Model**

   Use the export function to save the generated mesh in either PLY or STL format. A mesh model stored in the PLY format is usually smaller in file size than the STL format. The STL format is used in computer-aided surgery planning systems such as Brainlab.

The skull segmentation tool has default settings for VOI, intensity threshold and parameters for mesh simplification and smoothing. Most of the time, the default settings can be used, and the user does not need to manually adjust the settings.

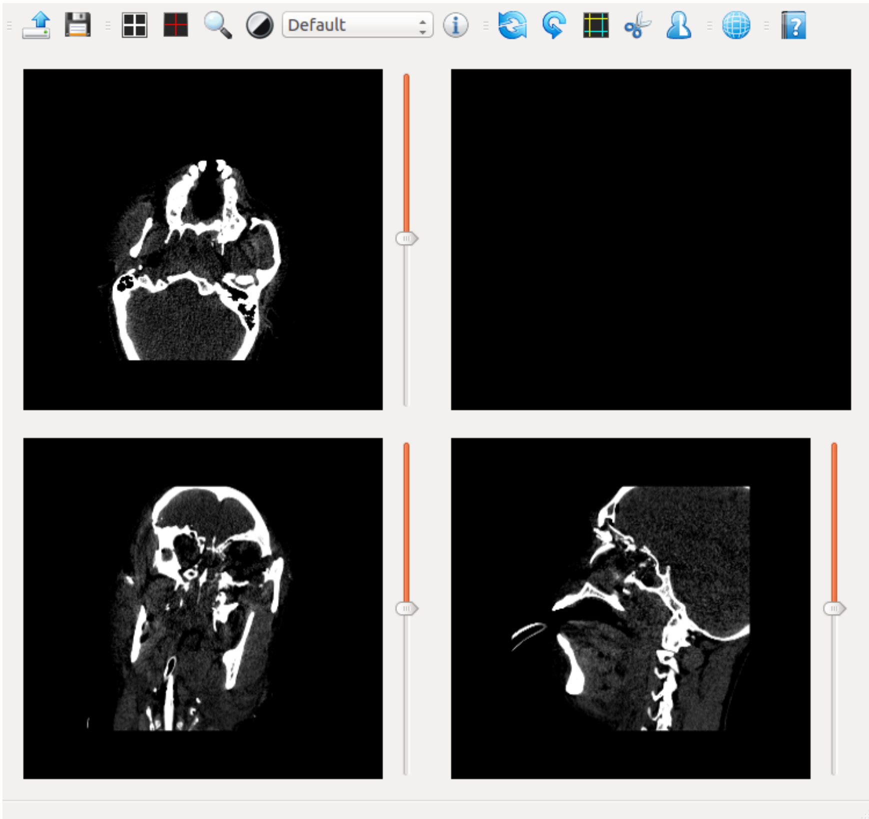

Figure 5.1: Imported CT images. The imported CT images are displayed in three standard anatomical views that are familiar to surgeons.

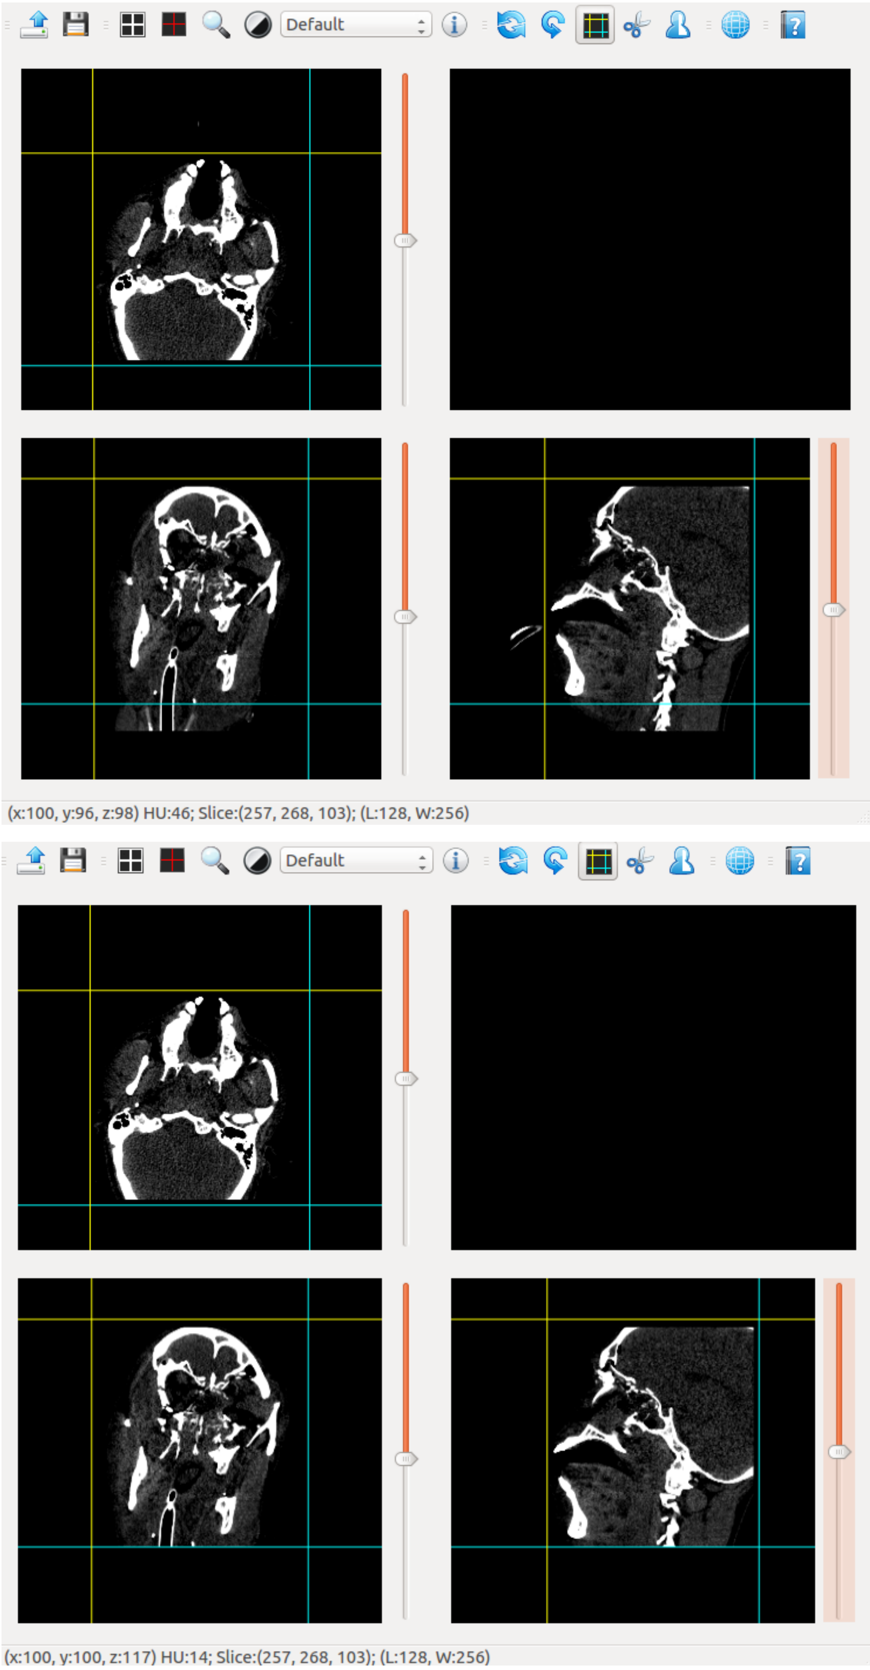

(a)

(b)

Figure 5.2: Determine VOI. (a) VOI lines in each of the three views define the VOI. (b) Volume outside VOI is cropped.

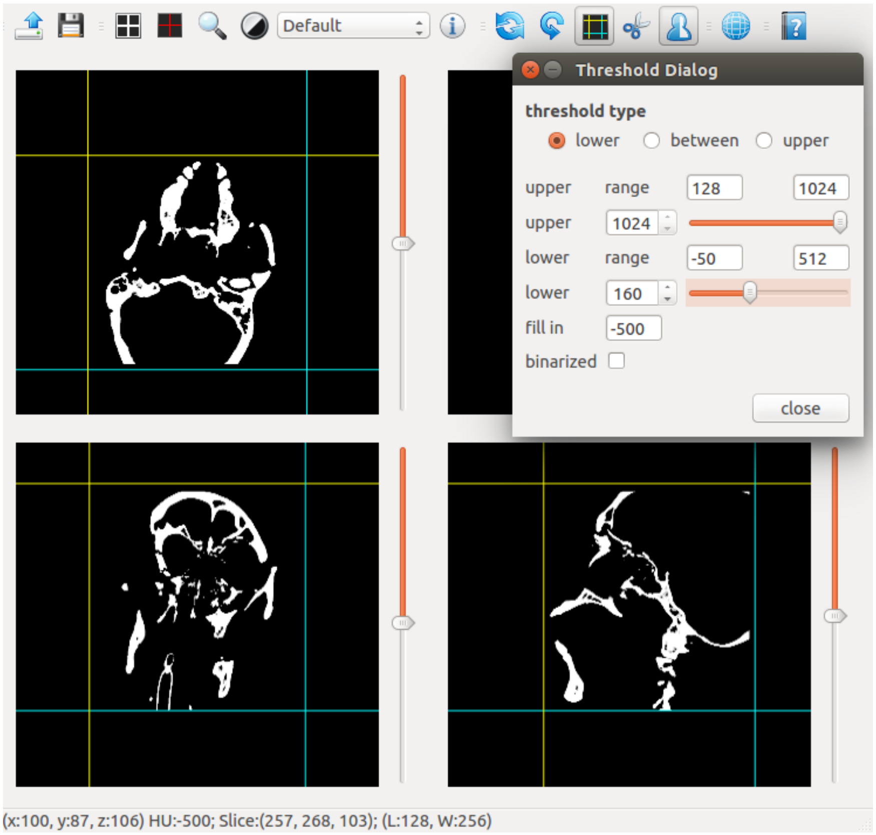

Figure 5.3: Adjust intensity threshold. Voxels with intensity no greater than the threshold are filtered out.

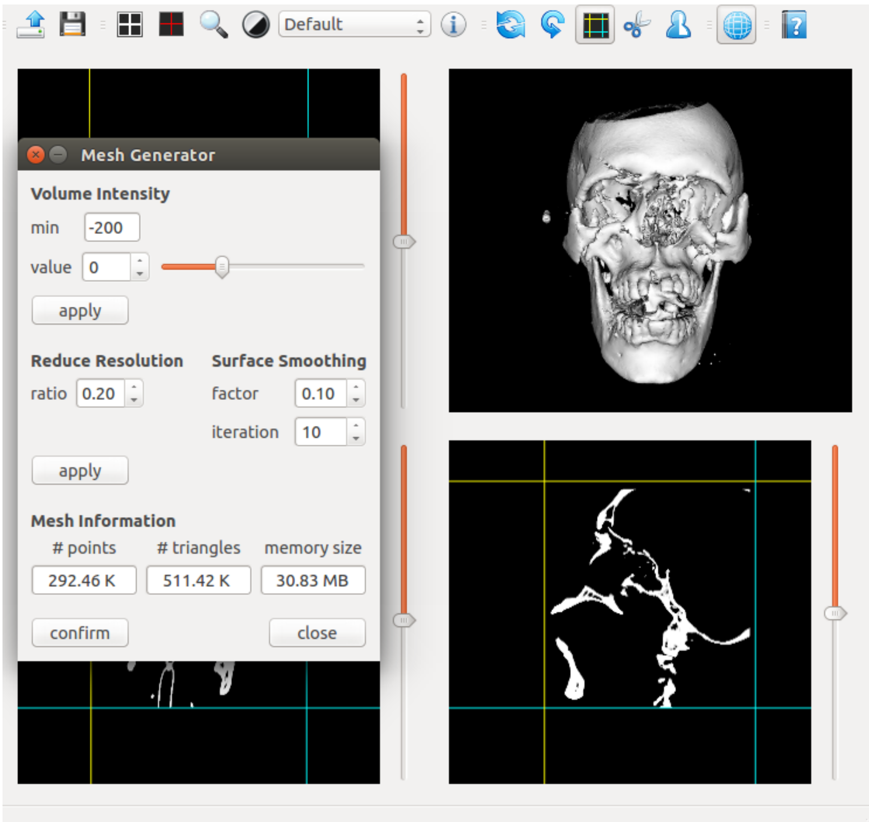

Figure 5.4: Generate mesh model. Mesh generated from voxels with intensity greater than the threshold.

5.2 Skull Reconstruction Tool

The skull reconstruction tool is used to generate a normal reconstructed model from a defective skull model. There are 7 main stages in the workflow. Follow the guide on the left panel of the tool to execute these stages sequentially:

1. **Import Mesh Model**

   Use the import function of the tool to import a mesh model. The imported mesh model is rendered in the central region of the tool (Figure 5.5).

2. **Remove Irrelevant Parts**

   The mesh model may contain a spine, tubes and things that are irrelevant to skull reconstruction. Use the cutting function of the tool to remove these irrelevant parts (Figure 5.6). Cutting is achieved by fitting a cutting plane that separates the irrelevant part from the main skull.

3. **Place Landmarks**

   Use the landmarking function of the tool to place landmarks on the target (Figure 5.7). Each of the target landmarks corresponds to a landmark on a reference model. It is noted that the landmarks should not be placed on the defective parts of the target. The landmarks can be used to align the reference to the target later.

4. **Mark Defective Parts**

   Use the painting function of the tool to mark defective parts (Figure 5.8). The marking of defective parts is performed by locating a VOI that encloses the defective parts.

5. **Select Reference Model**

   Use the reference selection function of the tool to select a reference model (Figure 5.9). The reference model should be as close to the target as possible.

6. **Restore Skull**

   This stage consists of several processes that are executed automatically by the tool. To begin with, the tool extracts the outer surfaces of the target model. Then, it spatially aligns the selected reference to the outer surfaces of the target model using the

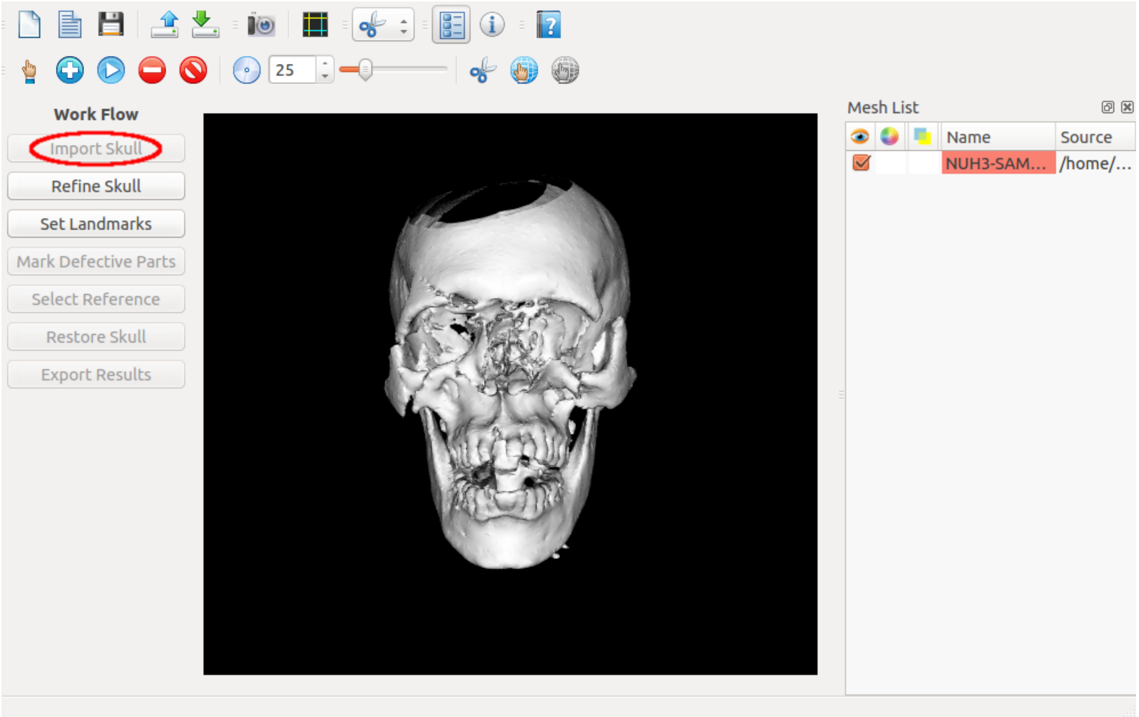

Figure 5.5: Imported mesh model. The imported mesh model is rendered in the central region of the tool.

shape and landmark information of two models. Spatial alignment of the reference and the target provides a good initialization for the skull reconstruction algorithm. In the current implementation, FICP [PLT07] is used to spatially align two models. It uses only the shape information of two models. Finally, it executes the skull reconstruction algorithm (Chapter 6) to generate a normal skull model for the target.

7. **Export Result**

   Use the export function to save the reconstructed model. Choose a file format to store the result. A mesh stored in PLY format is smaller in size than that in STL format. The STL format is used in surgery planning tools such as Brainlab.

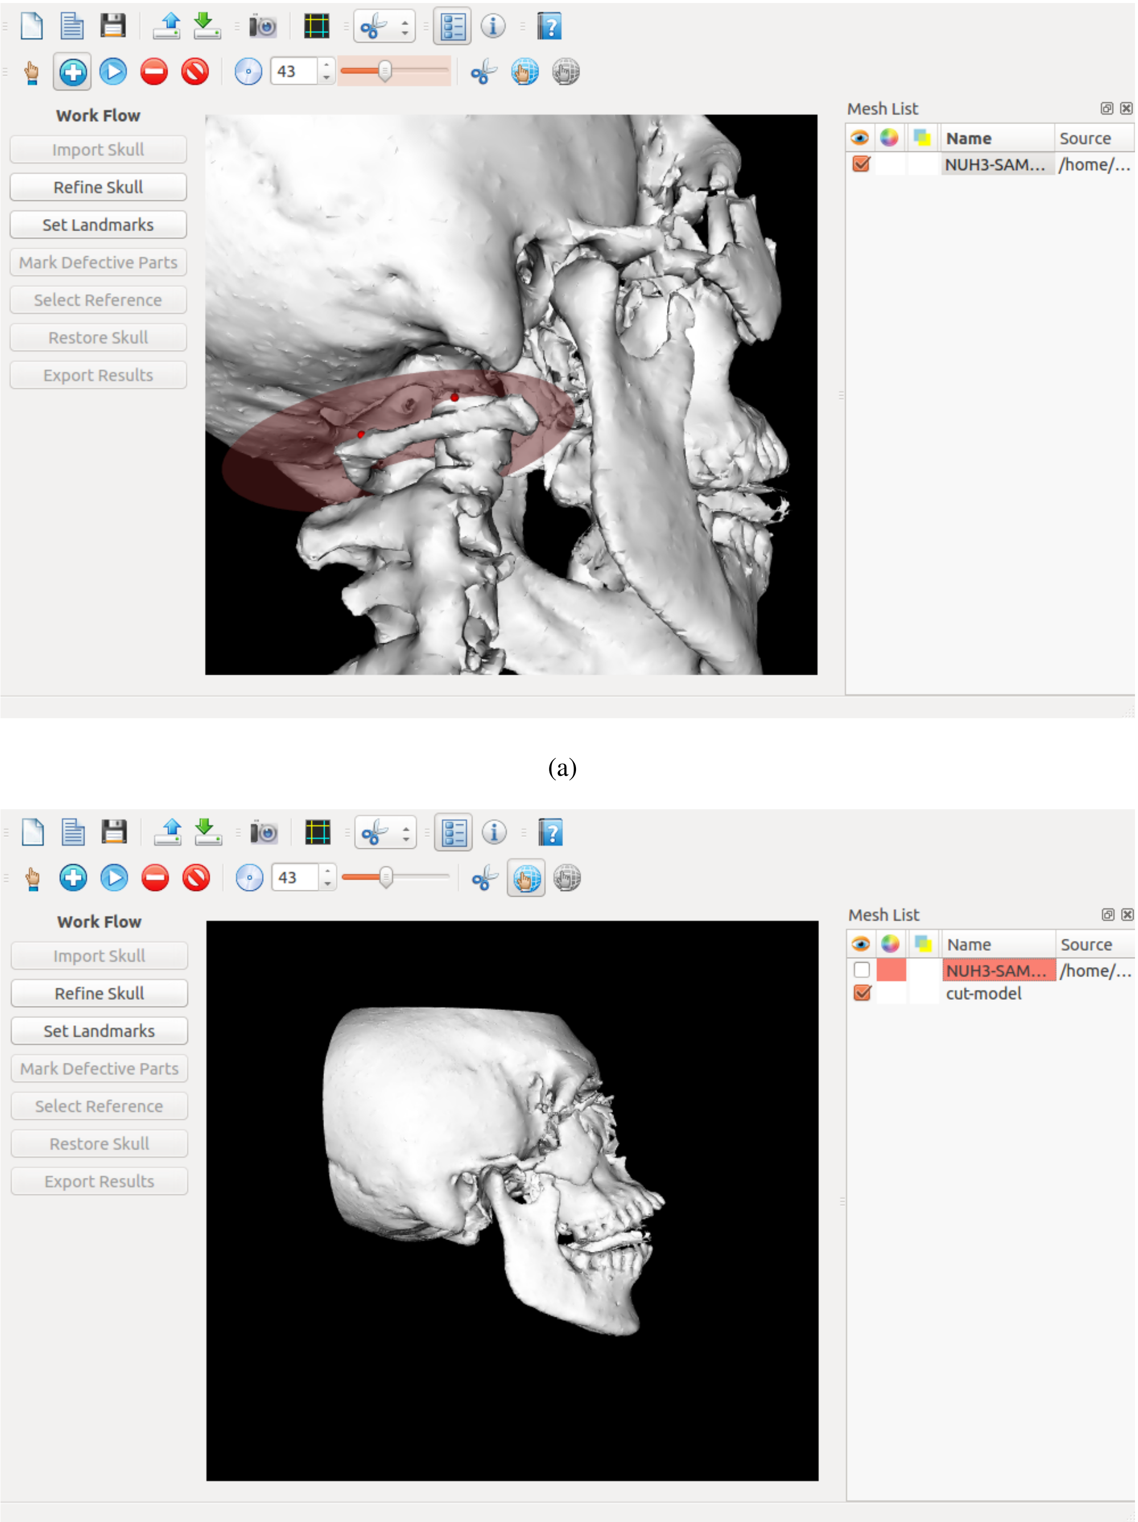

(a)

(b)

Figure 5.6: Remove irrelevant parts. (a) A cutting plane separates the spine from the main skull. (b) The main skull after the spine is removed.

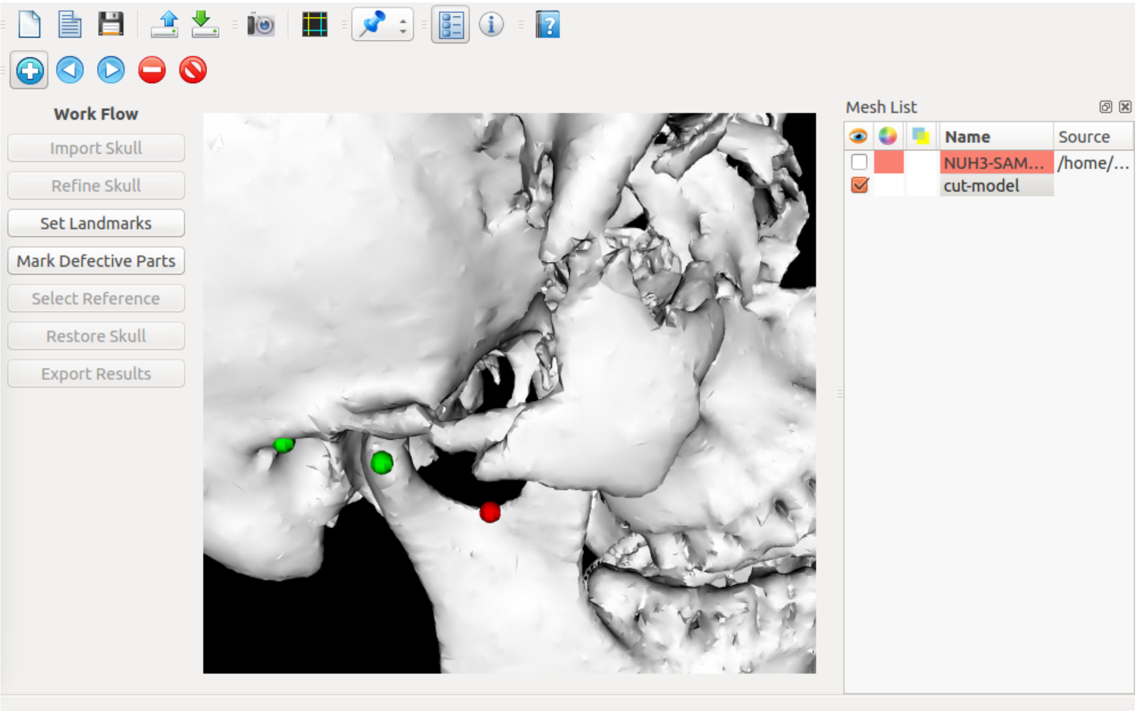

(a)

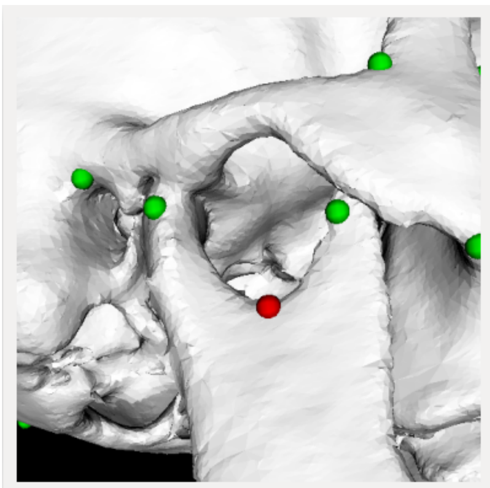

(b)

Figure 5.7: Place landmarks. (a) Each of the target landmarks corresponds to (b) a landmark on a reference model.

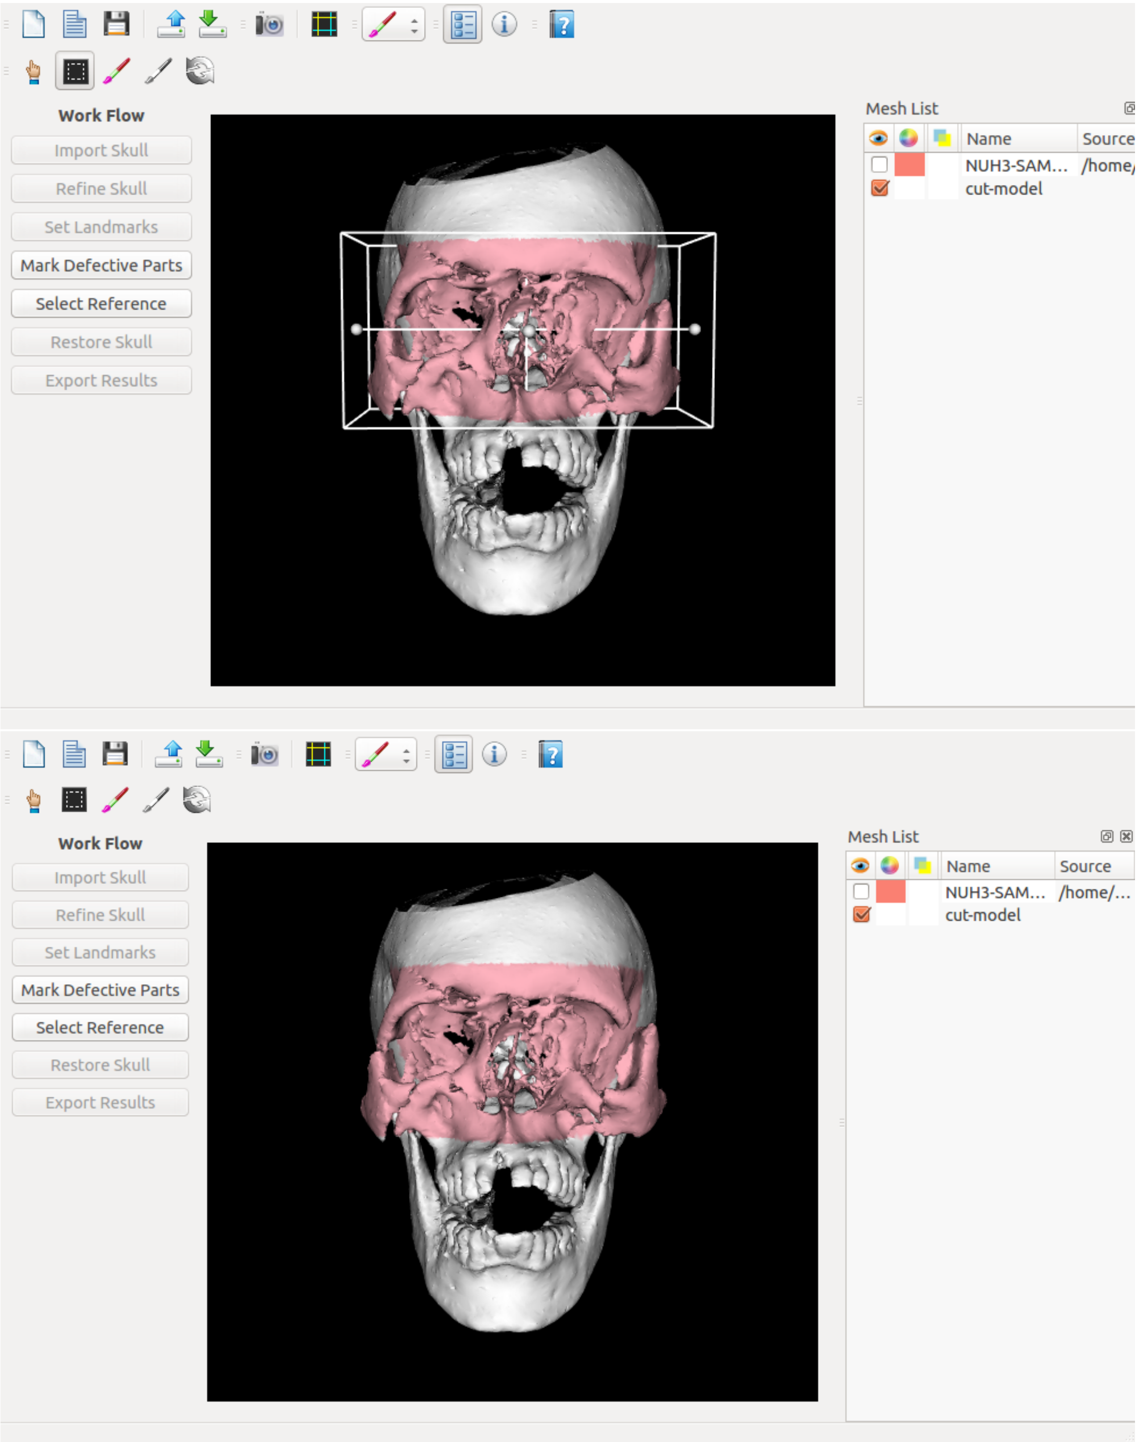

Figure 5.8: Mark defective parts. (a) VOI enclosing defective parts. (b) The defective parts on the target skull are painted pink.

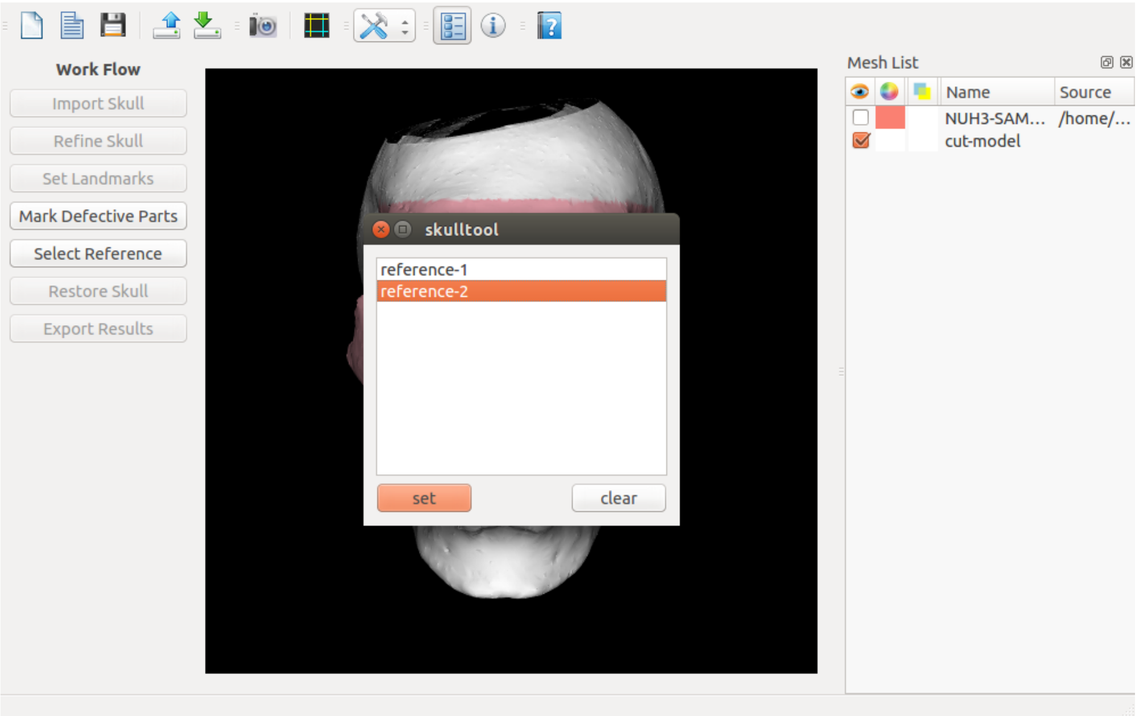

(a)

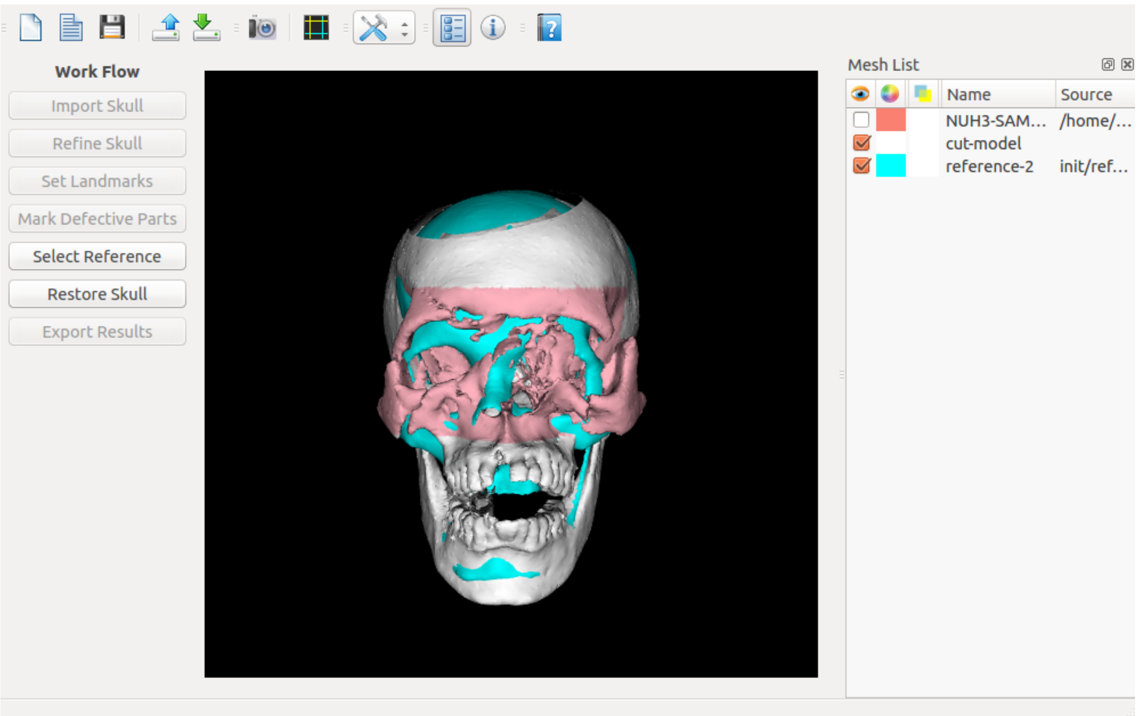

(b)

Figure 5.9: Select reference model. (a) Select a reference model from the reference candidates that is close to the target. (b) The reference model (green) is rendered in the central region of the tool.

5.3 Summary and Remaining Work

This chapter illustrates two software tools for skull reconstruction. The skull segmentation tool is completely developed, whereas the skull reconstruction tool is mostly complete. More work can be done to make the skull reconstruction tool more comprehensive and convenient. The work to be done includes:

1. **Reconstruction of Orbits**
   The tool is currently not applicable to the reconstruction of orbits. To reconstruct orbits, a complete and normal orbital model is needed as a reference. It is noted that orbits constructed from CT images are often incomplete due to segmentation error. A complete orbital model can be obtained by fitting a cone model to the orbit of a normal skull. In addition to a complete orbital model, the tool may need an algorithm for extracting the orbits of the target skull.

2. **Automatic Reference Selection**
   In the current reconstruction tool, the reference model is selected by the user. To choose a reference that is close to the target, the user needs to compare the target skull with each of the reference candidates, which is troublesome and time-consuming. The tool can automate this process by comparing the shape of the target and those of the reference candidates.

Chapter 6

Skull Reconstruction Algorithm

The skull reconstruction problem formulated in Section 4.2 is very complex. So, this thesis proposal first solves a simplified version of the problem. For this simplified problem, the defective parts D and the non-defective parts $\bar{D}$ are not known. In other words, the whole target skull is regarded as defective. Moreover, the symmetry constraint is not considered, and the orbit is not reconstructed.

6.1 Overview

Based on the simplified problem definition, an iterative optimization method called **FAIS** (**F**lip-**A**voiding **I**nterpolating **S**urface) is developed for skull reconstruction. It exploits the strength of the interpolating algorithm to produce an exact match between the reconstructed model and the non-defective part of the target model. It overcomes the shortcoming of the interpolating algorithm by detecting and excluding conflicting hard constraints that can cause surface flipping. Such exclusion is affordable when a very dense set of corresponding points is available. Thus, FAIS can reconstruct a skull without flipped surfaces and achieve practically zero error for the non-defective parts. It uses Laplacian deformation instead of the more popular thin-plate spline (TPS) because Laplacian deformation runs faster with increasing amount of hard constraints, whereas TPS runs slower.

For handling surface interference, FAIS is similar in spirit to [DYLV09], except FAIS detects possible flippings of triangular faces before deformation, which are removed from the constraint set for mesh deformation. FAIS's advantage is that it can be applied to triangular meshes, and it is conceptually simpler than [DYLV09].

6.2 Flip-Avoiding Registration

FAIS performs non-rigid interpolating registration of a reference model to a defective target model. It is assumed that the two models are already spatially aligned by rigid transformation. To achieve the goals discussed in Section 1.2 and 6.1, FAIS applies the following principles:

1. FAIS uses a small set of correspondence provided by the user to ensure anatomically correct registration of the reference model to the target model.

2. FAIS applies automatic correspondence search methods to obtain dense correspondence. It matches the surface characteristics of the reference and the target (Section 6.3), which allows FAIS to ignore outliers. Similar techniques are commonly used in existing methods.

3. FAIS detects and removes correspondence that may cause surface flipping (Section 6.4), thus achieving flip-avoiding reconstruction with interpolating surfaces.

4. Correspondence search is a local operation that is not guaranteed to be anatomically accurate. To reduce the risks of wrong correspondence, FAIS adopts an *iterative incremental* approach that deforms the reference model very slightly in the early iterations (Section 6.5). As the reference registers closer to the target in subsequent iterations, the risk of finding wrong correspondence is reduced, and the reference is allowed to deform more.

5. FAIS registers an interpolating surface to the non-defective parts of the target model exactly, resulting in zero error for the non-defective parts. In particular, Laplacian deformation is used for non-rigid registration.

6.3 Correspondence Search

FAIS applies two correspondence search methods. The first method is applied in the early iterations of FAIS. It searches for a corresponding mesh vertex $\mathbf{p}'$ on the target T for each mesh vertex $\mathbf{p}$ on the reference F that satisfies the conditions:

- $\mathbf{p}'$ is near enough to $\mathbf{p}$: $\|\mathbf{p} - \mathbf{p}'\| \leqslant D_1$, where D_1 is a constant parameter for search range; and
- $\mathbf{p}'$ and $\mathbf{p}$ have similar surface normals that differ by no greater than $10°$.

In the current implementation, D_1 is empirically set at 0.5mm.

The second method is applied in the final step. It searches for a corresponding point $\mathbf{p}'$ on the target T for each mesh vertex $\mathbf{p}$ on the reference F, such that

- $\mathbf{p}'$ is $\mathbf{p}$'s nearest surface point on T, i.e., the nearest intersection of the surface normal at $\mathbf{p}$ with T, and
- $\|\mathbf{p} - \mathbf{p}'\| \leqslant D_2$, where D_2 is a constant parameter.

D_2 is larger than D_1 but not so large that wrong correspondence is found. In the current implementation, $D_2 = 3$mm. The second method can find more corresponding points but is less efficient than the first. So, it is used only in the final step.

If a corresponding point $\mathbf{p}'$ is found for $\mathbf{p}$, then the vector $\mathbf{v}(\mathbf{p}) = \mathbf{p}' - \mathbf{p}$ is the *correspondence vector* of $\mathbf{p}$. Otherwise, $\mathbf{p}$ has no correspondence vector. The set C of correspondence contains tuples of the form $(\mathbf{p}, \mathbf{p}')$.

6.4 Flip Avoidance

Surface flipping is caused by the crossing of correspondence vectors that results in the flipping of a surface patch relative to its neighbouring surfaces (Figure 6.1). There is no surface flipping if the correspondence vectors do not cross. To derive the condition for flip avoidance, consider two points $\mathbf{p}$ and $\mathbf{q}$ on the surface of a mesh. If their correspondence vectors $\mathbf{v}(\mathbf{p})$ and $\mathbf{v}(\mathbf{q})$ meet at the same point, then they form a triangle with the vector

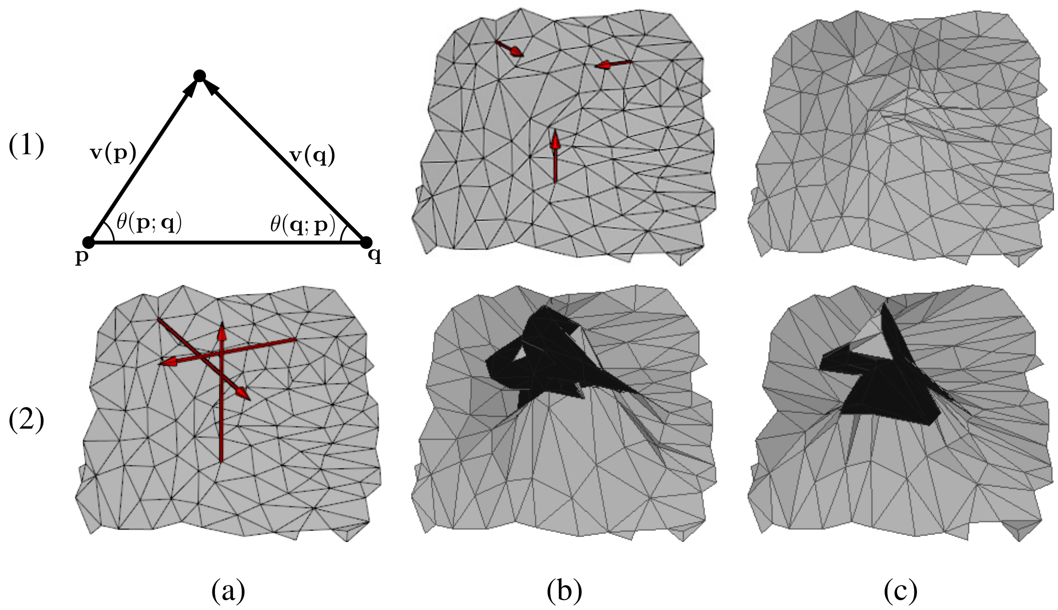

Figure 6.1: Surface flipping. Red arrows indicate the target locations of selected mesh vertices. (1a) Correspondence vectors $\mathbf{v}(\mathbf{p})$ and $\mathbf{v}(\mathbf{q})$ form a triangle with the line joining $\mathbf{p}$ and $\mathbf{q}$ when they meet at the same point. (1b) Non-crossing correspondence vectors (arrows) produce (1c) no surface flipping. (2a) Crossing correspondence vectors cause (2b) Laplacian deformation and (2c) TPS to produce flipped and distorted surfaces even when they do not intersect. Black regions are surface patches that have flipped.

$\mathbf{q} - \mathbf{p}$ from $\mathbf{p}$ to $\mathbf{q}$ (Figure 6.1(a)). Let $\theta(\mathbf{p}; \mathbf{q})$ denote the angle made by $\mathbf{v}(\mathbf{p})$ and $\mathbf{q} - \mathbf{p}$, and similarly for $\theta(\mathbf{q}; \mathbf{p})$. Then, basic trigonometry states that

$$\|\mathbf{v}(\mathbf{p})\| \cos\theta(\mathbf{p}; \mathbf{q}) + \|\mathbf{v}(\mathbf{q})\| \cos\theta(\mathbf{q}; \mathbf{p}) = \|\mathbf{p} - \mathbf{q}\|. \tag{6.1}$$

In general, $\mathbf{p}$ and $\mathbf{q}$ do not meet or intersect at a point in 3D space. Then, the left-hand side of Eq. 6.1 is the sum of the projections of $\mathbf{p}$ and $\mathbf{q}$ on the vector $\mathbf{q} - \mathbf{p}$. If $\|\mathbf{p} - \mathbf{q}\|$ is less than the left-hand side of Eq. 6.1, $\mathbf{v}(\mathbf{p})$ and $\mathbf{v}(\mathbf{q})$ will cross in 3D space, causing surface flipping (Figure 6.1(2)). If $\|\mathbf{p} - \mathbf{q}\|$ is greater than the left-hand side, $\mathbf{v}(\mathbf{p})$ and $\mathbf{v}(\mathbf{q})$ will not cross, and there is no flipping (Figure 6.1(1b,1c)).

To reduce the risk of obtaining wrong correspondence, the correspondence vectors should not be too long. Let D denote the upper bound on the length of the correspondence vectors:

$$\|\mathbf{v}(\mathbf{p})\| \leqslant D, \ \forall \mathbf{p}. \tag{6.2}$$

Then, $\mathbf{p}$ and $\mathbf{q}$ will not cross if

$$\cos\theta(\mathbf{p};\mathbf{q}) + \cos\theta(\mathbf{q};\mathbf{p}) < \frac{\|\mathbf{p}-\mathbf{q}\|}{D}. \tag{6.3}$$

This condition can be simplified as

$$\cos\theta(\mathbf{p};\mathbf{q}) < \frac{\|\mathbf{p}-\mathbf{q}\|}{2D} \text{ and } \cos\theta(\mathbf{q};\mathbf{p}) < \frac{\|\mathbf{q}-\mathbf{p}\|}{2D} \tag{6.4}$$

since Condition 6.4 implies Condition 6.3.

In order that $\mathbf{v}(\mathbf{p})$ does not cross any vector $\mathbf{v}(\mathbf{q})$, Condition 6.4 must be satisfied for all the points $\mathbf{q}$ on the mesh. Since $\cos\theta(\mathbf{p};\mathbf{q}) \leqslant 1$, Condition 6.4 is trivially satisfied for all points $\mathbf{q}$ at a distance larger than $2D$ from $\mathbf{p}$. Thus, we can state the following conditions for no crossing:

Simple No-Crossing Condition

There is no crossing if, for all pairs $(\mathbf{p},\mathbf{p}')$ and $(\mathbf{q},\mathbf{q}')$ in correspondence set C,

$$\|\mathbf{p}-\mathbf{q}\| > 2D. \tag{6.5}$$

General No-Crossing Condition

There is no crossing if, for each $(\mathbf{p},\mathbf{p}') \in C$,

$$\begin{aligned} &\cos\theta(\mathbf{p};\mathbf{q}) < \frac{\|\mathbf{p}-\mathbf{q}\|}{2D}, \\ &\forall\mathbf{q} \in N(\mathbf{p}) = \{\mathbf{q} \mid \|\mathbf{p}-\mathbf{q}\| \leqslant 2D\} \text{ and } (\mathbf{q},\mathbf{q}') \in C. \end{aligned} \tag{6.6}$$

The simple condition is a special case of the general condition.

6.5 Reconstruction Algorithm

FAIS reconstructs the resultant model R given a reference model F, a target model T, and known correspondence C^*. C^* is obtained from manual marking of significant anatomical landmarks on F and T that are adequately separated to ensure no crossing. In addition, F and T are assumed to be already spatially aligned by rigid transformation. FAIS is summarized in Algorithm 1.

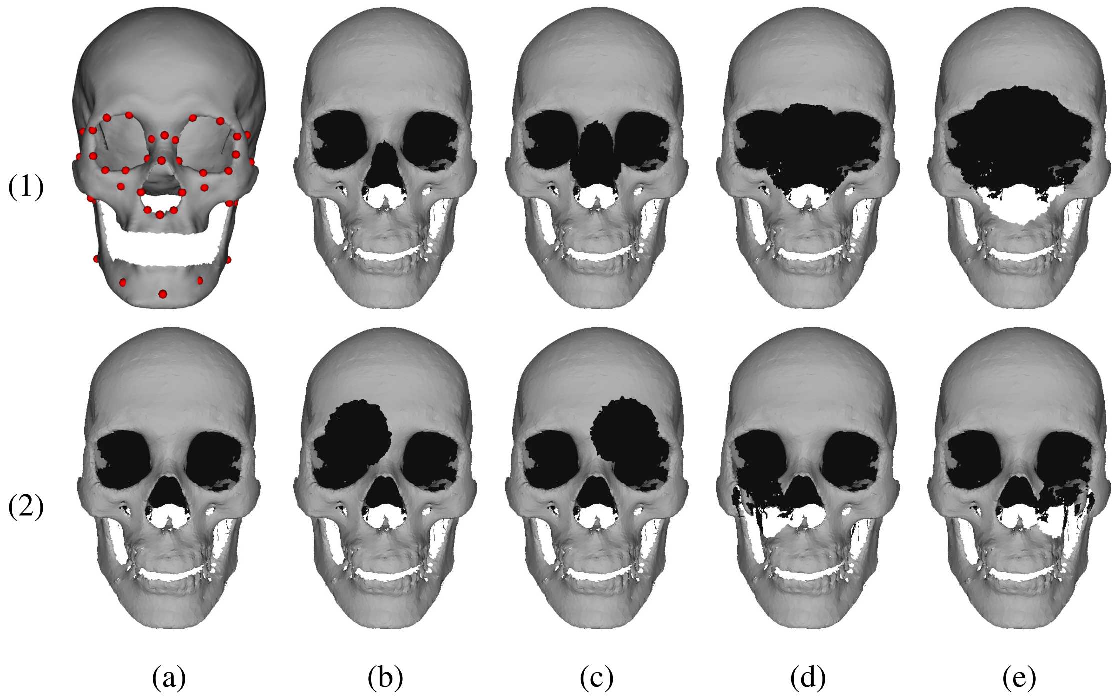

Figure 6.2: Sample skull models. (1a) Reference model with manual landmarks (red dots). (2a) Normal testing skull. (1b–e) Synthetic testing skulls with defective parts of different sizes, ranging from radius of 10mm to 40mm. (2b–e) Synthetic testing skulls with defective parts at various locations.

Algorithm 1: FAIS Flip-Avoiding Interpolating Surface

Input: Reference F, target T, known correspondence C^*.

1 Non-rigidly register F to T with positional constraints C^*, and set R as the registered F.

2 **for** k *from* 1 *to* K **do**

3   Find correspondence C from R to T using first correspondence search method.

4   Choose a sparse subset C^+ from $C^* \cup C$.

5   Non-rigidly register R to T with constraints C^+ incrementally.

6 **end**

7 Find correspondence C from R to T using second correspondence search method.

8 Remove crossings in $C^* \cup C$ giving C^+.

9 Non-rigidly register R to T with constraints C^+.

Output: Resultant R.

Step 1 non-rigidly registers reference F to target T with known correspondence C^* as the positional constraints, and sets the result R as the registered F. This step matches the overall anatomical shape of R to that of T in order to improve correspondence search in subsequent steps.

Steps 2 to 6 perform K iterations of non-rigid registration in small steps. First, Step 3 finds correspondence C from R to T using the first correspondence search method, which restricts all $\|\mathbf{v}(\mathbf{p})\|$ to be no longer than D_1 (Section 6.3). Step 4 chooses a sparse subset C^+ as follows: First, the upper bound D is set to the longest $\|\mathbf{v}(\mathbf{p})\|$ in $C^* \cup C$, thus, $D \leqslant D_1$. C^+ is initialized with known correspondence C^*. Then, each tuple $(\mathbf{p}, \mathbf{p}')$ in C is checked for sparse distribution: If there is a tuple $(\mathbf{q}, \mathbf{q}')$ in C^+ such that $\|\mathbf{p} - \mathbf{q}\| \leqslant 2D$, the tuple $(\mathbf{p}, \mathbf{p}')$ is discarded. Otherwise, it is added to C^+. This step ensures that all the reference points in C^+ are separated by a distance greater than $2D$, thereby satisfying the Simple No-Crossing Condition. Step 5 non-rigidly registers R to T, with each $\mathbf{p}$ in C^+ moved by an amount $(k/K)\|\mathbf{v}(\mathbf{p})\|$ along $\mathbf{v}(\mathbf{p})$. Thus, $\mathbf{p}$ is moved toward $\mathbf{p}'$ incrementally, allowing FAIS to recover from possible wrong correspondence in subsequent iterations.

Step 7 finds correspondence C from R to T using the second correspondence search method (Section 6.3). Step 8 removes crossings in $C^* \cup C$ as follows: First, the upper bound D is set to the longest $\|\mathbf{v}(\mathbf{p})\|$ in $C^* \cup C$, and the correspondence set C^+ is initialized to C^*. Next, each tuple $(\mathbf{p}, \mathbf{p}')$ in C is checked according to the General No-Crossing Condition. If the condition is satisfied, the tuple is added to C^+; otherwise, it is discarded. This step obtains a much denser set of correspondence than the sparse set in Step 4. Finally, Step 9 performs the final registration of R to T with C^+ as the positional constraints.

FAIS differs significantly from non-rigid ICP [ARV07, HMS12], although they have similar iterative structure. Non-rigid ICP performs locally affine registration of approximating surface, which has no surface flipping problem. On the other hand, FAIS performs non-rigid registration of interpolating surface, and needs to avoid surface flipping.

6.6 Experiments and Discussions

6.6.1 Data Preparation and PC Configuration

3D mesh models of skulls were constructed from patients' CT images. A normal skull was randomly selected as the reference model (Figure 6.2(a)), which had 92550 mesh vertices and 42 manual landmarks. Skull models with much fewer vertices cannot model the surfaces accurately. Teeth in the reference model were omitted because they were less important than facial bones in defining facial appearance, and many patients had missing teeth. Moreover, the resolution of the head CT was insufficient for constructing skull models with sufficient resolution for modelling each tooth accurately. Twenty normal skulls were used for testing. Five of the normal testing skulls were each used to synthesize 8 defective testing skulls with missing parts, giving a total of 40 defective testing skulls. For the examples in Figure 6.2(1b) to 6.2(1e), the defective parts were produced by removing mesh vertices within a spherical region of radius 10mm to 40mm. For the examples in Figure 6.2(2b) to 6.2(2e), the spherical regions were all 20mm but located at different locations. The 5 normal testing skulls serve as the ground truth.

The programs were implemented in Mathematica which used Intel MKL to solve linear systems. All tests were run on a PC with Intel i7-2600 CPU at 3.4GHz and 8GB RAM.

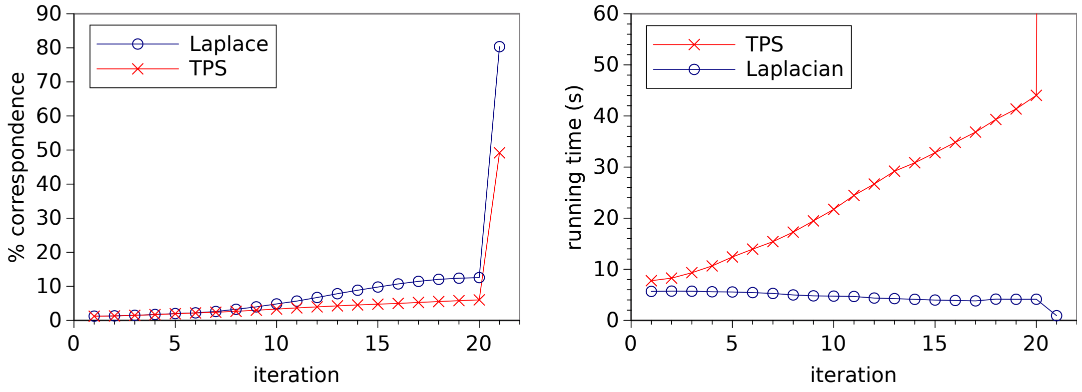

(a) (b)

Figure 6.3: Correspondence search. (a) Amount of correspondence at various iteration. (b) Running time (sec) is influenced by the amount of correspondence.

6.6.2 Dense Correspondence

In this experiment, FAIS was tested in turn with Laplacian deformation and TPS as the non-rigid registration method on a normal testing skull. Figure 6.3(a) shows that FAIS with Laplacian deformation finds more corresponding points than FAIS with TPS. This is because FAIS with Laplacian deformation is more accurate than FAIS with TPS (Section 6.6.3).

During the iterative stage from Step 2 to 6, up to 90% of the corresponding points are rejected by the Simple No-Crossing condition. At Step 8, FAIS's General No-Crossing Condition accepts 80% of the mesh vertices, amounting to about 74,000 corresponding points. In comparison, existing methods such as [DZS$^+$11, LP00, RB02, ZCL13] use several tens to hundreds of corresponding points, which are two orders of magnitude smaller than that of FAIS. With comparatively sparser correspondence sets, existing methods cannot achieve reconstruction accuracy as high as FAIS.

Figure 6.3(b) shows that Laplacian deformation runs faster and TPS runs slower with increasing amount of hard constraints. Moreover, TPS cannot run in Step 9 because its memory requirement exceeds available memory. FAIS's execution time is roughly proportional to the number of iterations K.

6.6.3 Reconstruction Accuracy

Reconstruction accuracy is highly dependent on the amount of correspondence available. To ensure that the methods tested have comparable amounts of correspondence, we compare FAIS with existing interpolating surface methods as follows:

- FAIS-1: FAIS with K = 10.
- FAIS-2: FAIS with K = 20.
- Laplace-1: As Step 1 of FAIS.
- Laplace-2: As Steps 1–6 of FAIS with K = 1.
- TPS-1: As Step 1 of FAIS except TPS is used; similar to [DZS$^+$11, LP00, RB02].
- TPS-2: As Steps 1–6 of FAIS with K = 1 except TPS is used; similar to [ZCL13].

The last four methods are equivalent to FAIS with different steps omitted. TPS-1 and TPS-2 are similar to existing methods, except that existing methods implicitly avoid surface flipping by choosing sparse correspondence sets (Section 6.6.2).

This test was performed on 20 normal skulls, 40 synthetic defective skulls and 2 real defective skulls. Reconstruction error was measured as the average distance from a mesh vertex on the reconstructed result R to the surface of the ground truth. Besides the overall error E averaged over the testing skulls, errors of the defective parts E_D, non-defective parts with and without positional constraints, respectively E_C and E_N, were measured separately. Some non-defective parts were excluded from the correspondence sets due to flip avoidance, and thus had no positional constraints.

Table 6.1 summarizes the reconstruction errors and Figure 6.5 details the reconstruction errors for various defective conditions. All methods have zero error E_C for the non-defective parts with positional constraints because they adopt interpolating surfaces and regard positional constraints as hard constraints. E_D is affected by the severity of defects as expected. Since the defective testing skulls were synthesized from only 5 normal skulls, they have a smaller variance compared to the 20 normal testing skulls. Therefore, their error E_N for non-defective parts without correspondence is smaller than that of the normal

Table 6.1: Average reconstruction errors of various algorithms.
E_C, E_N, E_D, and E are, respectively, error (in mm) of non-defective parts with positional constraints, non-defective parts without positional constraints, defective parts, and whole skull.

	Normal Skulls			Synthetic Skulls			
	E_C	E_N	E	E_C	E_N	E_D	E
FAIS-1	0.00	1.13	0.32	0.00	1.05	**1.00**	0.29
FAIS-2	0.00	**0.92**	**0.20**	0.00	**0.94**	1.01	**0.22**
Laplace-1	0.00	2.97	2.97	0.00	2.97	1.69	2.95
Laplace-2	0.00	2.43	2.39	0.00	2.44	1.60	2.39
TPS-1	0.00	5.46	5.46	0.00	5.63	1.42	5.56
TPS-2	0.00	5.24	5.17	0.00	5.39	1.24	5.26

skulls when tested with FAIS-1. The converse is true when tested on the other algorithms, as expected.

Compared to Laplace and TPS, FAIS has significantly smaller errors E_N for the non-defective parts without positional constraints, and slightly smaller errors E_D for the defective parts. Laplace and TPS have larger E_N than E_D because most of the manual landmarks are placed on the facial bones. Although FAIS-1 runs only half as many iterations as FAIS-2, its error is only slightly larger than that of FAIS-2 and significantly smaller than those of Laplace and TPS. Therefore, FAIS's performance is not significantly affected by the number of iterations K. With FAIS, up to 80% of the reference's vertices have corresponding points that serve as hard constraints, whereas only 10% are available to Laplace and TPS (Section 6.6.2). Therefore, FAIS has the lowest overall error of $E \leqslant 0.3$mm, whereas Laplace and TPS have higher errors of, respectively, $E = 2$–3mm and $E > 5$mm.

Visual inspection of the reconstructed results for synthetic skulls (Figure 6.4(1, 2)) confirms that FAIS's reconstructions are close to the ground truth. On the other hand, the reconstructions of Laplace and TPS have visible errors. Moreover, TPS reconstructions have less accurate aspect ratios.

For the real defective skulls (Figure 6.4(3,4)), FAIS's reconstruction of the jaws are

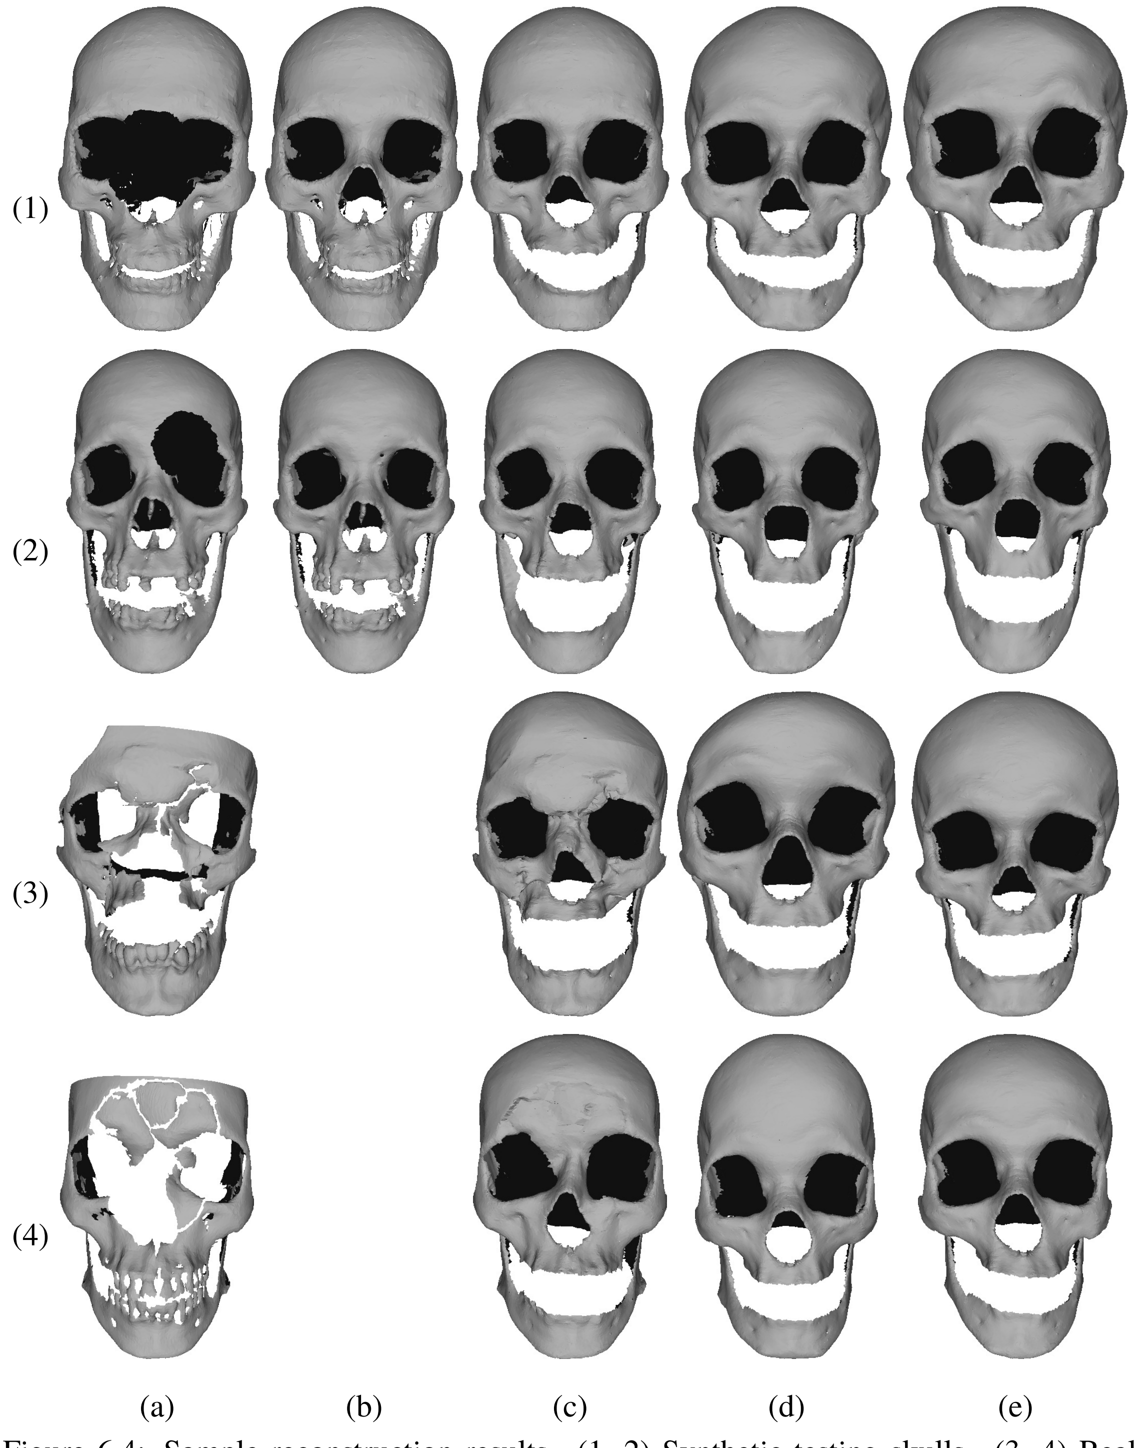

Figure 6.4: Sample reconstruction results. (1, 2) Synthetic testing skulls. (3, 4) Real defective skulls. Ground truth is not available. (a) Target, (b) ground truth, (c) FAIS-2, (d) Laplace-2, (e) TPS-2. Results of FAIS-1, Laplace-1, and TPS-1 are similar, and are omitted.

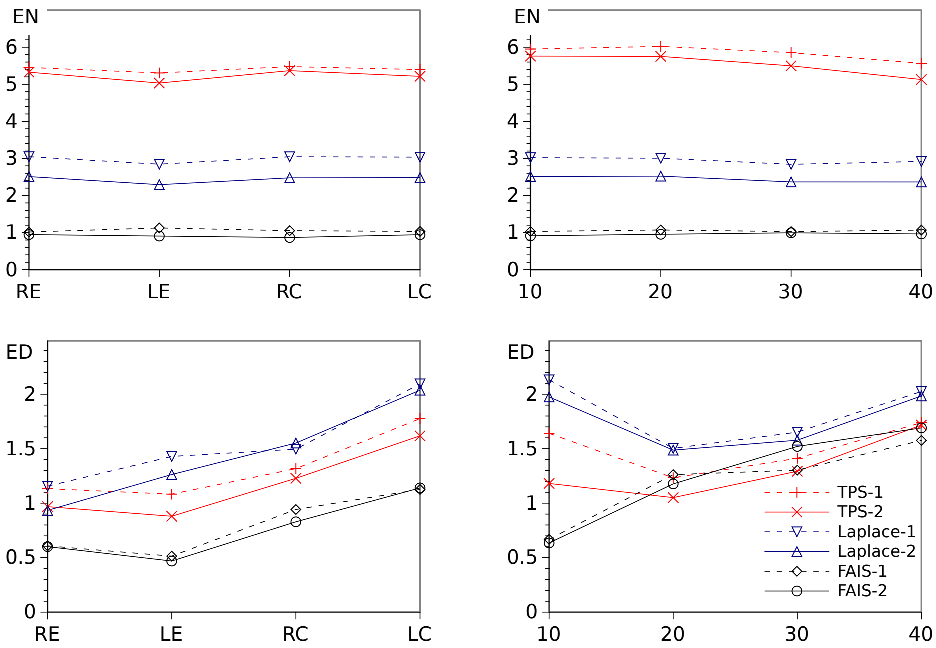

(a) Location (b) Size

Figure 6.5: Reconstruction errors for various defective conditions. E_N and E_D are, respectively, error of non-defective parts without positional constraints and error of defective parts. RE: right eye, LE: left eye, RC: right cheek, LC: left cheek.

close to the targets. Its reconstruction of the missing facial bones are structurally correct but distorted due to the fractured and deformed bones in the targets. The reconstructions of Laplace and TPS appear less distorted because they do not fit the targets closely and are thus closer to the reference (Figure 6.2(a)) than the targets.

6.6.4 Robustness Against Outliers

CT images that are used to construct 3D skull models can contain radiation artifacts caused by metallic dental implants [MMvS$^+$13] that are very difficult to remove. Thus, 3D skull models segmented and constructed from CT images often contain metal artifacts (Figure 6.6). This test evaluates the robustness of FAIS against outliers such as metal artifacts. The test was performed on 5 actual skulls with metal artifacts. FAIS-2 was tested using two reference models, one without teeth and one with teeth. For FAIS-2 with teeth, Step

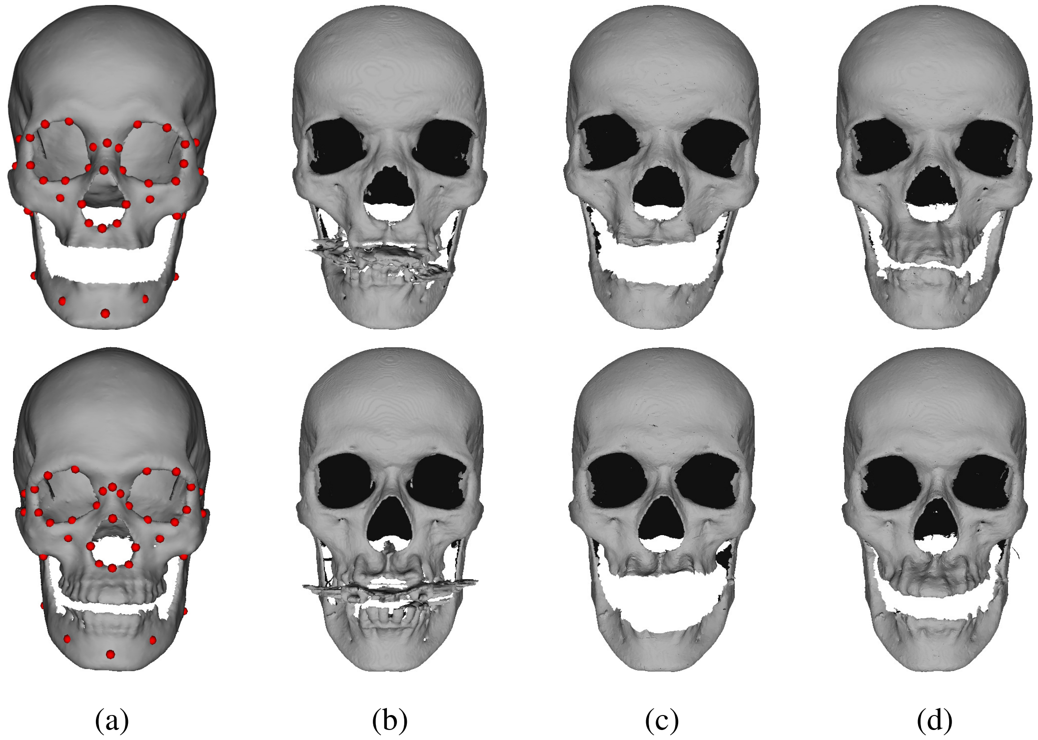

Figure 6.6: Robustness against metal artifacts. (a, row 1) Reference without teeth, (a, row 2) reference with teeth, (b) target with metal artifacts, (c) FAIS-2 using reference without teeth, (d) FAIS-2 using reference with teeth.

7 differed slightly such that the first, instead of the second, correspondence search method was applied on the mesh vertices in the teeth region.

Test results show that FAIS-2's correspondence search is robust enough to exclude metal artifacts as possible corresponding points. Consequently, it does not register the reference to the metal artifacts, and the reconstruction is free of metal artifacts. Using reference with teeth, some of the teeth reconstructed by FAIS-2 are slightly distorted because the skull mesh has insufficient resolution to model each tooth accurately. Test results also show that FAIS's reconstruction of the non-defective parts is independent of the reference used, which is expected when an interpolating surface is registered to a large amount of target points that serve as hard constraints.

Figure 6.7: References and targets. (1) Reference models (R1–R5). (2) Target models with teeth (T1–T4). (3) Target models without teeth (T5, T6).

6.6.5 Effect of Reference Model

A test was carried out to investigate how the choice of reference model affects FAIS's reconstruction accuracy. There are two specific goals in the test. The first goal is to investigate whether a reference model can be used to reconstruct every target accurately. The second goal is to see which reference is the best for reconstructing a specific target.

The test was performed on 5 reference models and 6 complete target models (Figure 6.7). All the reference models had no teeth. Four target models had teeth whereas the other two had no teeth. The error measurement used in the previous sections is asymmetric because correspondence between the reference and the target is not a one-to-one mapping. Thus two error measures, E and E', were adopted. The error E measured the average distance from the reconstructed model to the ground truth (i.e., the target model), whereas the error E' measured the average distance from the ground truth to the reconstructed model.

Table 6.2 summarizes the average reconstruction errors E of various targets given a reference model. Test results show that every reference model can reconstruct at least one target model accurately with error $E < 0.2$mm. In addition, R1–R4 can reconstruct all the targets with error $E < 0.5$mm. This suggests that a few references may be enough for reconstructing all the targets. Moreover, the more similar is the reference to a target, the more accurate is the reconstruction (Figure 6.8(1,2)). In all cases, the largest reconstruction error is still smaller than 1mm, which is very accurate.

Table 6.3 summarizes the average reconstruction error E' of various reference models given a target model. Test results show that each target can be reconstructed by a suitable reference model with error E' between 0.5mm and 1.5mm. E' is larger than E for the same pair of reference and target. In particular, the errors E' for targets T1–T4 are much larger than E because the targets have teeth whereas the references have no teeth. The additional error is due to the wrong correspondence of the teeth in the targets. In all cases, the largest E' is still smaller than 1.7mm, which is quite accurate.

Figure 6.8 illustrates reconstructed results with the smallest and largest errors. For cases with smallest reconstruction errors (Figure 6.8(1,3)), the references are close to the targets, and the reconstructed models are very close to the targets. For cases with largest reconstruction errors, the references differ significantly in shape from the targets. These test results imply that, for best reconstruction, it is advisable to choose a reference that is as similar to the target as possible.

Table 6.2: Average reconstruction errors E of various targets given a reference. The error E is the average error (mm) measured from the reconstructed model to the ground truth (i.e., the target).

	T1	T2	T3	T4	T5	T6
R1	**0.13**	0.23	0.23	**0.13**	0.22	0.33
R2	**0.10**	0.23	0.21	0.21	**0.11**	0.28
R3	**0.16**	0.22	0.26	**0.11**	0.21	0.32
R4	0.21	0.34	0.36	**0.18**	0.30	0.43
R5	**0.15**	0.66	0.77	0.40	0.39	0.58

Table 6.3: Average reconstruction errors E' of various references given a target. The error E' is the average error (mm) measured from the ground truth (i.e., the target) to the reconstructed model.

	T1	T2	T3	T4	T5	T6
R1	1.14	1.25	1.50	1.39	0.58	0.87
R2	1.19	**1.21**	**1.41**	1.48	**0.46**	0.78
R3	1.22	1.27	1.42	1.38	0.59	0.84
R4	1.65	1.60	1.56	1.42	0.63	**0.51**
R5	**0.78**	1.26	1.56	**1.00**	0.49	0.64

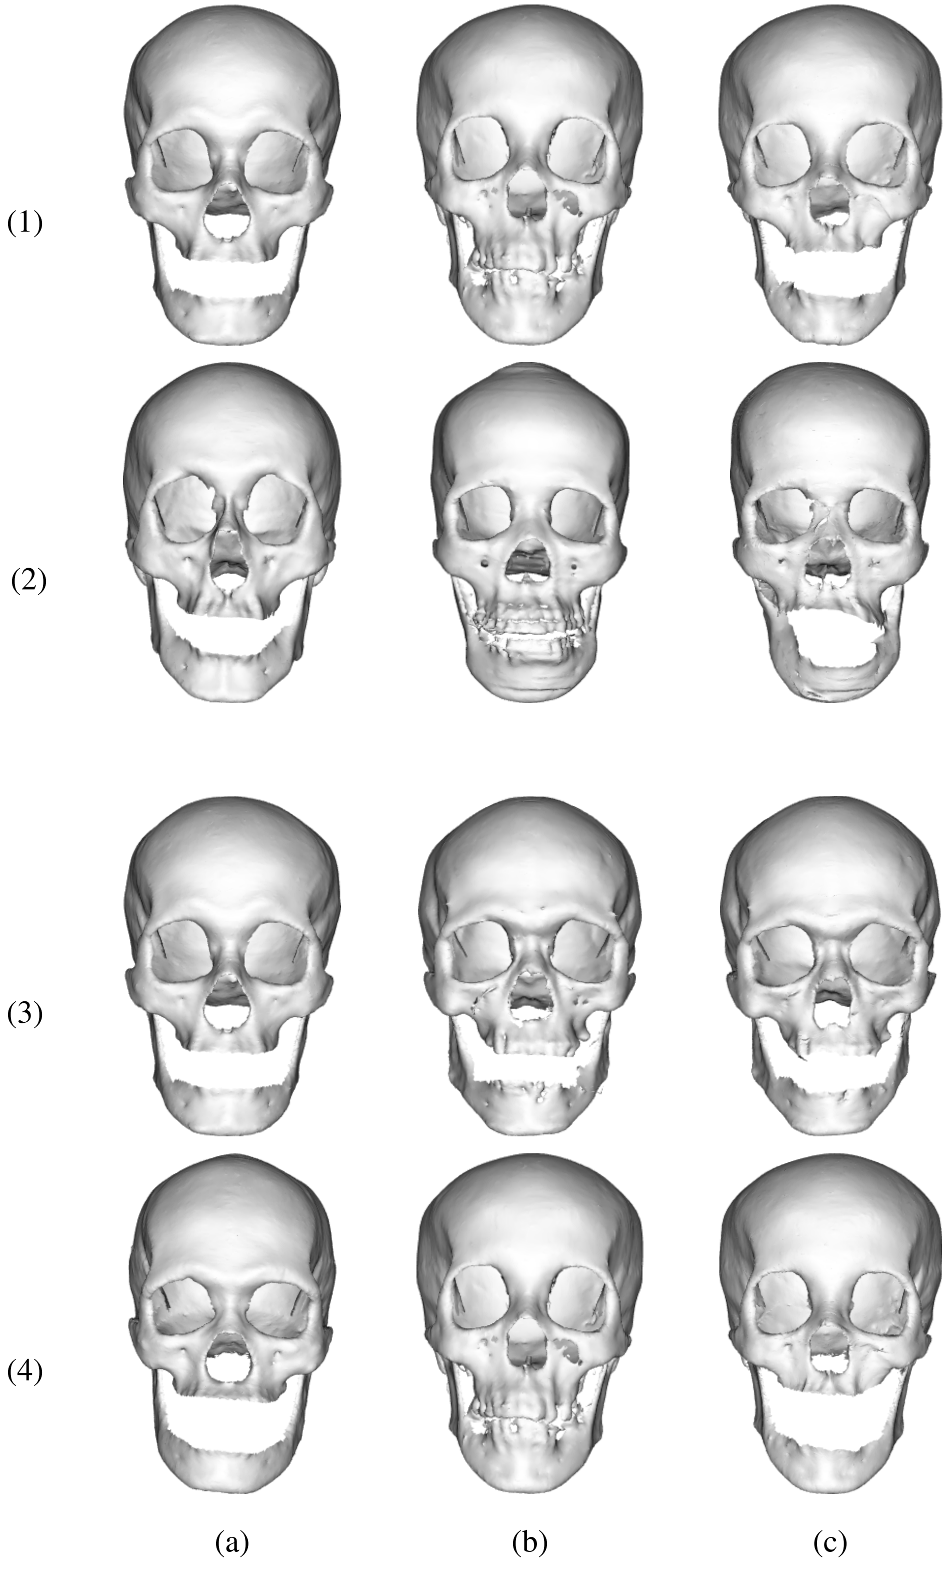

Figure 6.8: Sample reconstruction results. (a) Reference models. (b) Target models. (c) Reconstructed models. (1) Reconstruction of T1 by R2 has smallest error $E = 0.10$mm. (2) Reconstruction of T3 by R5 has largest error $E = 0.77$mm. (3) Reconstruction of T5 by R2 has smallest error $E' = 0.46$mm. (4) Reconstruction of T1 by R4 has largest error $E' = 1.65$mm.

6.7 Summary and Remaining Work

This chapter proposed a skull reconstruction algorithm, namely FAIS, that solves a simplified skull reconstruction problem. FAIS has the following nice properties:

- FAIS can accurately generate normal reconstructed models not only from fractured skulls, but also from incomplete and deformed ones. Thus FAIS is applicable to a wider range of applications including craniofacial surgery planning, forensic investigation and anthropological studies.
- FAIS can handle bilaterally defective skulls because of the nature of geometric reconstruction. This overcomes the limitation of symmetry-based methods which are applicable only to laterally defective skulls.
- FAIS can achieve practically zero errors with respect to the non-defective parts without surface distortion. This is because FAIS adopts an interpolating non-rigid registration method with a very dense set of correspondence. As a result, the reconstructed shape of non-defective parts is almost independent of the reference model used. Test results also show that a few reference models may be enough for accurately reconstructing all targets, which overcomes the limitation of statistical methods that require a large number of reference models for training.
- FAIS's reconstruction of missing parts are very close to the ground truth (Figure 6.4(1,2)), and FAIS is robust to outliers such as dental artifacts (Figure 6.6). Thus, FAIS is capable to recover the complete shape of the incomplete models, and it can ignore dental artifacts. Moreover, FAIS can be used to establish dense correspondence between skull models, which is useful for shape analysis and statistical modelling.

Further tests and possible enhancements of FAIS are needed:

1. **Effect of Initial Alignment**

   FAIS requires the reference model to be spatially aligned to the target model. In the current implementation of skull reconstruction tool, the alignment is performed using FICP [PLT07]. Since the skull reconstruction tool captures landmarks on the target, it is also possible to perform spatial alignment by computing the best similarity transformation or applying plane-fitting registration [CXLZ15]. It is worth investigating which rigid registration algorithm provides the best initial alignment.

2. **Symmetry Constraint**

   The current FAIS does not include symmetry constraint. Figure 6.4(3c,4c) show that the nasal bones of the reconstructed models are distorted. This problem may be overcome by introducing symmetry constraint into the Laplacian deformation algorithm used in FAIS.

3. **Handling Defective Parts**

   The current FAIS does not know which part of the target skull is defective. Consequently, the reconstructed results can be misled by wrong correspondence due to fractured bones (Figure 6.4(3c,4c)). Since the skull reconstruction tool provides functionality for marking defective parts, the tool can pass the information about defective parts to FAIS. Then, FAIS can ignore the correspondence on the defective parts and perform the reconstruction as though they are missing.

4. **Reconstruction of Orbits**

   The current FAIS has been tested on the reconstruction of the outer surfaces of the skull but not the orbits (eye sockets). As landmarks on the rims of the orbits are captured by the skull reconstruction tool, it may be possible for FAIS to reconstruct the orbits in a similar way as skull surface reconstruction. Further investigation is needed to check whether FAIS or the skull reconstruction tool needs to be enhanced for orbit reconstruction.

Chapter 7

Conclusion

This thesis proposal describes the proposed research work of skull reconstruction that includes two research objectives. The first objective is to develop an accurate and robust algorithm for reconstructing a normal skull given a defective target skull. The second objective is to develop two software tools for skull reconstruction, namely a skull segmentation tool and a skull reconstruction tool. The tools, together with the reconstruction algorithm, are applicable to craniofacial surgery planning and forensic investigation.

The skull segmentation tool is completely developed. The skull reconstruction tool is mostly complete except for orbit reconstruction. To reconstruct the orbits, it requires a complete orbital model and may need an algorithm for extracting the orbits of the target skull. Another possible enhancement is to provide the functionality of automatically selecting the most suitable reference for the target.

A skull reconstruction algorithm has been developed to solve a simplified version of the skull reconstruction problem. To solve the complete problem, FAIS needs to include symmetry constraint, handle defective parts of the target and accurately reconstruct the orbits. Moreover, it is worth investigating which rigid registration algorithm provides the best initial registration. These works will be done in the continuing research.

Bibliography

[ARV07] B. Amberg, S. Romdhani, and T. Vetter. Optimal step nonrigid ICP algorithms for surface registration. In *Proc. IEEE Conference on CVPR*, pages 1–8. IEEE, 2007.

[Boo89] F.L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (6):567–585, 1989.

[BS11] S. Benazzi and S. Senck. Comparing 3-dimensional virtual methods for reconstruction in craniomaxillofacial surgery. *Journal of Oral and Maxillofacial Surgery*, 69(4):1184–1194, 2011.

[BSB14] F. Bonarrigo, A. Signoroni, and M. Botsch. Deformable registration using patch-wise shape matching. *Graphical Models*, 76(5):554–565, 2014.

[BSMG09] S. Benazzi, E. Stansfield, C. Milani, and G. Gruppioni. Geometric morphometric methods for three-dimensional virtual reconstruction of a fragmented cranium: The case of Angelo Poliziano. *International Journal of Legal Medicine*, 123(4):333–344, 2009.

[Che13] Y. Cheng. *Computer-aided craniomaxillofacial surgery planning for fractured skulls*. PhD thesis, National University of Singapore, 2013.

[CL00] Y. Choi and S. Lee. Injectivity conditions of 2D and 3D uniform cubic B-spline functions. *Graphical Models*, 62:411–427, 2000.

[CLL12] Y. Cheng, W.K. Leow, and T.C. Lim. Automatic identification of frankfurt plane and mid-sagittal plane of skull. In *Applications of Computer Vision*

(WACV), 2012 IEEE Workshop on, pages 233–238. IEEE, 2012.

[CSP+07] J. Chapuis, A. Schramm, I. Pappas, W. Hallermann, K. Schwenzer-Zimmerer, F. Langlotz, and M. Caversaccio. A new system for computer-aided preoperative planning and intraoperative navigation during corrective jaw surgery. *Information Technology in Biomedicine, IEEE Transactions on*, 11(3):274–287, 2007.

[CTS+10] L.H.C. Cevidanes, S. Tucker, M. Styner, H. Kim, J. Chapuis, M. Reyes, W. Proffit, T. Turvey, and M. Jaskolka. Three-dimensional surgical simulation. *American Journal of Orthodontics and Dentofacial Orthopedics*, 138(3):361–371, 2010.

[CXLZ15] Y. Cheng, S. Xie, W.K. Leow, and K. Zhang. Plane-fitting robust registration for complex 3D models. In *Proc. CAIP*, pages 640–651, 2015.

[Din10] F. Ding. *3D Segmentation of Soft Tissues by Flipping-free Mesh Deformation*. PhD thesis, National University of Singapore, 2010.

[DM01] H. Delingette and J. Montagnat. Shape and topology constraints on parametric active contours. *Computer Vision and Image Understanding*, 83(2):140–171, 2001.

[dMCP+06] E. de Momi, J. Chapuis, I. Pappas, G. Ferrigno, W. Hallermann, A. Schramm, and M. Caversaccio. Automatic extraction of the mid-facial plane for cranio-maxillofacial surgery planning. *International Journal of Oral and Maxillofacial Surgery*, 35(7):636–642, 2006.

[DYLV09] F. Ding, W. Yang, W. K. Leow, and S. Venkatesh. 3D segmentation of soft organs by flipping-free mesh deformation. In *Proc. WACV*, 2009.

[DZS+11] Q. Deng, M. Zhou, W. Shui, Z. Wu, Y. Ji, and R. Bai. A novel skull registration based on global and local deformations for craniofacial reconstruction. *Forensic Science Int.*, 208:95–102, 2011.

[FdCP+08] M. Fantini, F. de Crescenzio, F. Persiani, S. Benazzi, and G. Gruppioni. 3D restitution, restoration and prototyping of a medieval damaged skull. *Rapid*

Prototyping Journal, 14(5):318–324, 2008.

[GMBW04] P. Gunz, P. Mitteroecker, F.L. Bookstein, and G.W. Weber. Computer aided reconstruction of human crania. *BAR International Series*, 1227:92–95, 2004.

[Gun05] P. Gunz. *Statistical and Geometric Reconstruction of Hominid Crania: Reconstructing Australopithecine Ontogeny*. PhD thesis, Universität Wien, 2005.

[GWL99] H.K. Gumprecht, D.C. Widenka, and C.B. Lumenta. Brainlab vectorvision neuronavigation system: technology and clinical experiences in 131 cases. *Neurosurgery*, 44(1):97–104, 1999.

[HF98] M. Hagenlocker and K. Fujimura. CFFD: A tool for designing flexible shapes. *The Visual Computer*, 14, 1998.

[HHK92] W.M. Hsu, J.F. Hughes, and H. Kaufman. Direct manipulation of free-form deformations. In *ACM Siggraph Computer Graphics*, volume 26, pages 177–184. ACM, 1992.

[HMS12] H. Hontani, T. Matsuno, and Y. Sawada. Robust nonrigid icp using outlier-sparsity regularization. In *Proc. IEEE Conference on CVPR*, pages 174–181. IEEE, 2012.

[İH93] M.Y. İşcan and R.P. Helmer. *Forensic Analysis of the Skull: Craniofacial Analysis, Reconstruction, and Identification*. John Wiley & Sons Inc, 1993.

[JSC04] W. Jung, H. Shin, and B. K. Choi. Self-intersection removal in triangular mesh offsetting. *Computer-Aided Design and Applications*, 1:477–484, 2004.

[KAB+05] A. Khan, E. Aylward, P. Barta, M. Miller, and M. F. Beg. Semi-automated basal ganglia segmentation using large deformation diffeomorphic metric mapping. In *Proc. MICCAI*, pages 238–245, 2005.

[LAV09] M. Lüthi, T. Albrecht, and T. Vetter. Building shape models from lousy data. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2009*, pages 1–8. Springer, 2009.

[LCL+03] M.Y. Lee, C.C. Chang, C.C. Lin, L.J. Lo, and Y.R. Chen. T medical rapid prototyping in custom implant design for craniofacial reconstruction. In *Systems, Man and Cybernetics, 2003. IEEE International Conference on*, volume 3, pages 2903–2908. IEEE, 2003.

[LM98] J.-O. Lachaud and A. Montanvert. Deformable meshes with automated topology changes for coarse-to-fine three-dimensional surface extraction. *Medical Image Analysis*, 3(2):187–207, 1998.

[LP00] R. J. A. Lapeer and R. W. Prager. 3D shape recovery of a newborn skull using thin-plate splines. *Computerized Medical Imaging & Graphics*, 24(3):193–204, 2000.

[LYW+11] X. Li, Z. Yin, L. Wei, S. Wan, W. Yu, and M. Li. Symmetry and template guided completion of damaged skulls. *Computers & Graphics*, 35(4):885–893, 2011.

[MMvS+13] A. Mouton, N. Megherbi, K. van Slambrouk, J. Nuyts, and T. P. Breckon. An experimental survey of metal artefact reduction in computed tomography. *J. X-Ray Science and Tech.*, 21(2):193–226, 2013.

[Mos16] Mosby. *Mosby's Medical Dictionary*. Elsevier, 2016.

[MT99] T. McInerney and D. Terzopoulos. Topology adaptive deformable surfaces for medical image volume segmentation. *IEEE Trans. Medical Imaging*, 18(10):840–850, 1999.

[MYF06] H. Masuda, Y. Yoshioka, and Y. Furukawa. Interactive mesh deformation using equality-constrained least squares. *Computers & Graphics*, 30(6):936–946, 2006.

[PLT07] Jeff M. Phillips, Ran Liu, and Carlo Tomasi. Outlier Robust ICP for Minimizing Fractional RMSD. In *Proc. of Int. Conf. on 3D Digital Imaging and*

Modeling, pages 427–434, 2007.

[RB02] A. Rosas and M. Bastir. Thin-plate spline analysis of allometry and sexual dimorphism in the human craniofacial complex. *American J. Physical Anthropology*, 117:236–245, 2002.

[Row98] S. Roweis. Em algorithms for pca and spca. *Advances in Neural Information Processing Systems*, pages 626–632, 1998.

[RRS03] M. Rossi, E. Ribeiro, and R. Smith. Craniofacial asymmetry in development: an anatomical study. *The Angle Orthodontist*, 73(4):381–385, 2003.

[RSS+01] K. Rohr, H.S. Stiehl, R. Sprengel, T.M. Buzug, J. Weese, and M.H. Kuhn. Landmark-based elastic registration using approximating thin-plate splines. *Medical Imaging, IEEE Transactions on*, 20(6):526–534, 2001.

[SCOL+04] O. Sorkine, D. Cohen-Or, Y. Lipman, M. Alexa, C. Rössl, and H.P. Seidel. Laplacian surface editing. In *Proc. Eurographics/ACM SIGGRAPH Symp. Geometry Processing*, pages 175–184, 2004.

[SP86] T.W. Sederberg and S.R. Parry. Free-form deformation of solid geometric models. *ACM SIGGRAPH computer graphics*, 20(4):151–160, 1986.

[TA98] R.G. Taylor and C. Angel. Facial reconstruction and approximation. In J.G. Clement and D.L. Ranson, editors, *Craniofacial Identification in Forensic Medicine*, chapter 14, pages 177–186. Arnold London, 1998.

[TB99] M.E. Tipping and C.M. Bishop. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(3):611–622, 1999.

[TBK+05] W. D. Turner, R. E. Brown, T. P. Kelliher, P. H. Tu, M. A. Taister, and K. W. Miller. A novel method of automated skull registration for forensic facial approximation. *Forensic Science Int.*, 154:149–158, 2005.

[Woo31] T.L. Woo. On the asymmetry of the human skull. *Biometrika*, 22(3–4):324–352, 1931.

[WYLL11] L. Wei, W. Yu, M.Q. Li, and X. Li. Skull assembly and completion using template-based surface matching. In *3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT), 2011 International Conference on*, pages 413–420. IEEE, 2011.

[YLL12] W. Yu, M. Li, and X. Li. Fragmented skull modeling using heat kernels. *Graphical Models*, 74(4):140–151, 2012.

[YWML11] Z. Yin, L. Wei, M. Manhein, and X. Li. An automatic assembly and completion framework for fragmented skulls. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 2532–2539. IEEE, 2011.

[ZBH07] A. Zaharescu, E. Boyer, and R. Horaud. Transformesh: A topology-adaptive mesh-based approach to surface evolution. In *Proc. ACCV*, pages 166–175, 2007.

[ZCL13] K. Zhang, Y. Cheng, and W.K. Leow. Dense correspondence of skull models by automatic detection of anatomical landmarks. In *Computer Analysis of Images and Patterns*, pages 229–236. Springer, 2013.

[Zha14] K. Zhang. Dense correspondence and statistical shape reconstruction of fractured and incomplete skulls. Master's thesis, National University of Singapore, 2014.

[ZLC15] K. Zhang, W. K. Leow, and Y. Cheng. Performance analysis of active shape reconstruction of fractured, incomplete skulls. In *Proc. CAIP*, 2015.

[ZLES05] S. Zachow, H. Lamecker, B. Elsholtz, and M. Stiller. Reconstruction of mandibular dysplasia using a statistical 3D shape model. In *International Congress Series*, volume 1281, pages 1238–1243. Elsevier, 2005.

[ZRA+08] X. Zhuang, K. Rhode, S. Arridge, R. Razavi, D. Hill, D. Hawkes, and S. Ourselin. An atlas-based segmentation propagation framework using locally affine registration—Application to automatic whole heart segmentation. In *Proc. MICCAI*, pages 425–433, 2008.