Received: 15 November 2017 | Revised: 2 February 2018 | Accepted: 2 February 2018
DOI: 10.1002/rcs.1906

WILEY The International Journal of Medical Robotics and Computer Assisted Surgery

ORIGINAL ARTICLE

Flip-avoiding interpolating surface registration for skull reconstruction

Shudong Xie[1] | Wee Kheng Leow[1] | Hanjing Lee[2] | Thiam Chye Lim[2,3]

[1]Department of Computer Science, National University of Singapore,
[2]Division of Plastic, Reconstruction and Aesthetic Surgery, National University Hospital, Singapore
[3]Department of Surgery, National University of Singapore

Correspondence
Shudong Xie, Department of Computer Science, National University of Singapore.
Email: xshudong@comp.nus.edu.sg

Funding information
Ministry of Education, Singapore, Grant/Award Number: MOE2014-T2-1-062

Handling Editor: Loughlin Clive

Abstract

Skull reconstruction is an important and challenging task in craniofacial surgery planning, forensic investigation and anthropological studies. Existing methods typically reconstruct approximating surfaces that regard corresponding points on the target skull as soft constraints, thus incurring non-zero error even for non-defective parts and high overall reconstruction error. This paper proposes a novel geometric reconstruction method that non-rigidly registers an interpolating reference surface that regards corresponding target points as hard constraints, thus achieving low reconstruction error. To overcome the shortcoming of interpolating a surface, a flip-avoiding method is used to detect and exclude conflicting hard constraints that would otherwise cause surface patches to flip and self-intersect. Comprehensive test results show that our method is more accurate and robust than existing skull reconstruction methods. By incorporating symmetry constraints, it can produce more symmetric and normal results than other methods in reconstructing defective skulls with a large number of defects. It is robust against severe outliers such as radiation artifacts in computed tomography due to dental implants. In addition, test results also show that our method outperforms thin-plate spline for model resampling, which enables the active shape model to yield more accurate reconstruction results. As the reconstruction accuracy of defective parts varies with the use of different reference models, we also study the implication of reference model selection for skull reconstruction.

KEYWORDS
interpolating surface, Laplacian deformation, model resampling, non-rigid registration, skull reconstruction

1 | INTRODUCTION

Skull reconstruction is an important and challenging task in craniofacial surgery planning, forensic investigation and anthropological studies. Existing skull reconstruction methods can be broadly divided into four categories: symmetry-based, statistical, bone repositioning and geometric.

Symmetry-based methods[1-3] rely on the approximate left–right symmetry of human skulls. They regard the reflection of the non-defective parts of a target skull about the lateral symmetry plane as the reconstruction of the defective parts. These methods are not applicable when both sides of a skull are defective.

Statistical methods, particularly active shape models,[4-6] map a target skull to a statistical skull model by computing the model parameters that best fit the non-defective parts of the target, and generate the reconstructed skull from the model parameters. Unlike human face images, it is very difficult to collect a wide variety of 3D models of human skulls ($\gg$ 50) to cover all normal skull variations across age, race and gender. Thus, it is difficult to apply statistical methods to skull reconstruction. Moreover, generating the reconstructed model that fits the overall shape of the target skull is a global optimization process, and thus the reconstructed model may not fit the non-defective parts of the target model closely. Therefore, statistical methods are not very accurate in skull reconstruction.

Bone repositioning methods[7-10] reconstruct a skull by repositioning fractured bone fragments of a defective skull at their correct positions, which is similar to solving a 3D jigsaw puzzle. 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

Wee Kheng Leow, Member, IEEE

surfaces of bone fragments. Moreover, these methods are applicable only to fractured skulls rather than incomplete and deformed ones. They also require every individual fractured bone fragment to be segmented from computed tomography (CT) images. Without an automatic way to segment the bone fragments, it is very tedious and time-consuming to obtain the bone fragments manually.

Geometric methods[11-14] perform non-rigid registration of a single 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 number of corresponding points used in non-rigid registration. Methods that use a small set of manually marked landmarks[11-13] cannot achieve high accuracy. Thus, some methods automatically detect more corresponding points.[14]

Non-rigid registration methods can be grouped into two broad categories based on the goal of registration: approximation and interpolation. Methods that produce **approximating surfaces** such as 3D snake,[15] active balloon,[16] piecewise rigid registration[17] and non-rigid iterative closest point (ICP)[18,19] fit a reference surface to the target by minimizing the 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. On the other hand, methods that produce **interpolating surfaces** such as thin-plate spline (TPS)[20] and Laplacian deformation[21] 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.

Among the interpolating methods, TPS is the most popular for reconstruction of skulls[11-14,22,23] because it can tolerate noise by imposing a surface smoothness constraint through the minimization of surface bending energy. Laplacian deformation preserves local surface curvature and normal, and we are the first to apply it to skull reconstruction.[24,25]

Interpolating methods can produce flipped surfaces when there are conflicts in the hard constraints. Surface flipping refers to the local inversion of surface normals that causes surface self-intersections (Figure 1E,F). They cause severe distortion of surface shape, and are very difficult to remove (Section 3). Note that surface flipping is a direct consequence of surface interpolation with conflicting hard constraints. Imposing a surface smoothness constraint by energy minimization, such as TPS, cannot remove surface flipping (Figure 1F). 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.

This paper proposes a novel method called FAIS (flip-avoiding interpolating surface) that exploits the strength of an interpolating surface while overcoming its shortcoming. It avoids surface flipping by detecting and excluding conflicting hard constraints. Such exclusion is affordable when a very dense set of corresponding points is available (Section 7.2). 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 TPS because Laplacian deformation runs faster with an increasing number of hard constraints, whereas TPS runs slower. Comprehensive test results show that FAIS is more accurate and robust than existing skull reconstruction methods. By incorporating symmetry constraints, it can produce more symmetric and normal results than other methods in reconstructing defective skulls with a large number of defects. FAIS is also robust against severe outliers such as radiation artifacts in CT due to dental implants. In addition, test results also show that FAIS outperforms TPS for model resampling, which enables the active shape model to yield more accurate reconstruction results. As the reconstruction accuracy of defective parts may vary according to different reference models used, we also study the implication of reference model selection for skull reconstruction.

2 | HANDLING SURFACE SELF-INTERSECTION

There are two general approaches to handling surface flipping and self-intersection: (1) detection and resolution, and (2) avoidance. McInemey and Terzopoulos[26] detect self-intersections by examining the deformation result, and resolve self-intersections by rolling back the mesh model to the state before deformation and imposing repulsive forces to keep the potentially intersecting surfaces apart. Lachaud and Montanvert[27] impose proximity conditions between mesh vertices and detects violations of proximity conditions, whereas others[28,29] detect self-intersections through collision detection, and resolve self-intersections by remeshing.

The methods of Choi and Lee[30] and Hagenlocker and Fujimura[31] avoid self-intersection by imposing an injectivity (one-to-one) condition on the free-form deformation function. The injectivity condition confines the free-form deformation of the mesh to regions that do not have self-intersection. Khan et al.[32] and Zhuang et al.[33] apply a

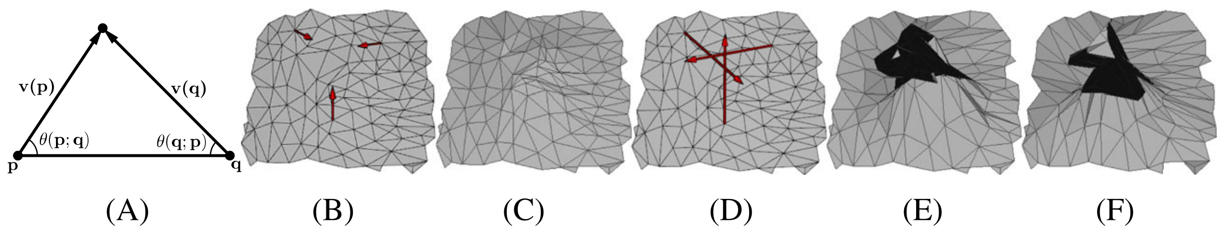

FIGURE 1 Surface flipping. (A), Correspondence vectors **v**(**p**) and **v**(**q**) form a triangle with the line joining **p** and **q** when they meet at the same point. (B), Non-crossing correspondence vectors (arrows) produce no surface flipping (C). D, Crossing correspondence vectors cause Laplacian deformation (E) and also cause TPS to produce flipped and distorted surfaces (F) even when they do not intersect. Black regions are surface patches that have flipped

WILEY- The International Journal of Medical Robotics and Computer Assisted Surgery

diffeomorphic deformation function. A diffeomorphic function and its inverse are both one-to-one and smooth, and the function preserves the topology of the mesh model after deformation, thus avoiding self-intersection. These methods produce approximating surfaces instead of interpolating ones.

Ding et al.[34] devise an ingenious quadrilateral mesh that permits easy detection of possible flippings of mesh edges by arranging the mesh vertices into some forms of total ordering. Edge 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.

Our FAIS is similar in spirit to theirs,[34] except that 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 that of Ding et al.[34]

3 | FLIP-AVOIDING REGISTRATION

3.1 | Overview

FAIS performs non-rigid interpolating registration of a reference model to a defective target model. To achieve the goals discussed in Section 1, FAIS applies the following principles.

1. FAIS uses a small set of landmarks 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 3.2), 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 3.3), 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 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 with correspondence. In particular, Laplacian deformation is used for non-rigid registration.

3.2 | 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 $\mathbf{p}$: $\|\mathbf{p} - \mathbf{p}'\| \leq D_1$, where D_1 is a constant parameter for the search range; and
- $\mathbf{p}'$ and $\mathbf{p}$ have similar surface normals that differ by no more than 10°.

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

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
- $\|\mathbf{p} - \mathbf{p}'\| \leq D_2$, 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 $(\mathbf{p}, \mathbf{p}')$.

3.3 | Flip avoidance

The crossing of correspondence vectors can cause self-intersection or penetration of surfaces, resulting in local inversion of surface normals, which we call surface flipping (Figure 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 model. 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 $\mathbf{q} - \mathbf{p}$ from $\mathbf{p}$ to $\mathbf{q}$ (Figure 1A). 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

$$\|\mathbf{v}(\mathbf{p})\| \cos\theta(\mathbf{p};\mathbf{q}) + \|\mathbf{v}(\mathbf{q})\| \cos\theta(\mathbf{q};\mathbf{p}) = \|\mathbf{p} - \mathbf{q}\|. \tag{1}$$

In general, $\mathbf{p}$ and $\mathbf{q}$ do not meet or intersect at a point in 3D space. Then, the left-hand side of equation 1 is the sum of the projections of $\mathbf{v}(\mathbf{p})$ and $\mathbf{v}(\mathbf{q})$ on the vector $\mathbf{q} - \mathbf{p}$. If $\|\mathbf{p} - \mathbf{q}\|$ is less than the left-hand side of equation 1, $\mathbf{v}(\mathbf{p})$ and $\mathbf{v}(\mathbf{q})$ will cross in 3D space, causing surface flipping (Figure 1D–F). 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 1B, C).

Let D denote the upper bound on the length of the correspondence vectors: $\|\mathbf{v}(\mathbf{p})\| \leq D, \forall \mathbf{p}$. Then, $\mathbf{v}(\mathbf{p})$ and $\mathbf{v}(\mathbf{q})$ will not cross if

$$\cos\theta(\mathbf{p};\mathbf{q}) + \cos\theta(\mathbf{q};\mathbf{p}) < \frac{\|\mathbf{p} - \mathbf{q}\|}{D}. \tag{2}$$

This condition can be simplified as

$$\cos\theta(\mathbf{p};\mathbf{q}) < \frac{\|\mathbf{p} - \mathbf{q}\|}{2D} \quad \text{and} \quad \cos\theta(\mathbf{q};\mathbf{p}) < \frac{\|\mathbf{q} - \mathbf{p}\|}{2D} \tag{3}$$

since condition 3 implies condition 2.

In order that $\mathbf{v}(\mathbf{p})$ does not cross any vector $\mathbf{v}(\mathbf{q})$, condition 3 must be satisfied for all the points $\mathbf{q}$ on the mesh. Since $\cos\theta(\mathbf{p};\mathbf{q}) \leq 1$, condition 3 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, $\|\mathbf{p} - \mathbf{q}\| > 2D$.

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

$$\cos\theta(\mathbf{p};\mathbf{q}) < \frac{\|\mathbf{p}-\mathbf{q}\|}{2D}, \qquad (4)$$
$$\forall \mathbf{q} \in N(\mathbf{p}) = \{\mathbf{q} \,|\, \|\mathbf{p}-\mathbf{q}\| \leq 2D\} \text{ and } (\mathbf{q}, \mathbf{q}') \in C.$$

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

4 | LAPLACIAN DEFORMATION WITH SYMMETRY CONSTRAINTS

Consider the C-shape model shown in Figure 2A. We want to apply Laplacian deformation to deform it such that the bottom landmark is fixed and the left landmark is moved to the right slightly. The result is a deformed shape whose bottom part is fixed but the top part is shifted to the right (Figure 2B).

A similar situation can occur when reconstructing a defective skull with a large number of defective or missing parts (Figure 2C). The reference model shown in Figure 2D happens to be disconnected on the right like the C-shape model, and it is wider than the target. Many reference landmarks on the facial bones have no corresponding target landmarks because of the large number of missing parts in the target. During non-rigid registration, Laplacian deformation reduces the width of the lower jaw of the reference to fit the target's lower jaw according to the landmarks. This process moves the left and right landmarks of the reference inward onto the positions of the corresponding target landmarks. These movements, coupled with the C-shaped reference and lack of correspondence of reference landmarks, cause the craniofacial bones of the reference to shift to the right instead of reducing its width. Consequently, a distorted skull is produced (Figure 2F).

To overcome this problem, we impose a **mid-plane constraint** on Laplacian deformation as follows. Every reference skull model has some landmarks called the **mid-point landmarks** that fall on the mid-line of the skull. These mid-point landmarks form a plane called the mid-plane (Figure 2D). Before deformation, the mid-plane is a vertical, laterally symmetric plane. The mid-plane constraint states that after deformation, the mid-plane should still be a vertical, laterally symmetric plane. With this additional constraint, Laplacian deformation will produce an undistorted result (Figure 2G).

Although the mid-plane stays vertical and laterally symmetric, the model after deformation can still be laterally distorted (Figure 3C). This happens when some reference landmarks on one side of the skull have no corresponding points on the target. To overcome this problem, we introduce the **symmetry constraint**, which constrains every symmetric pair of landmarks to remain symmetric after deformation. With this constraint, Laplacian deformation will produce a symmetric and undistorted result (Figure 3D).

Laplacian deformation[21,35] applies the discrete Laplacian operator $L(\mathbf{p}_i)$ to estimate the surface curvature and normal at vertex i:

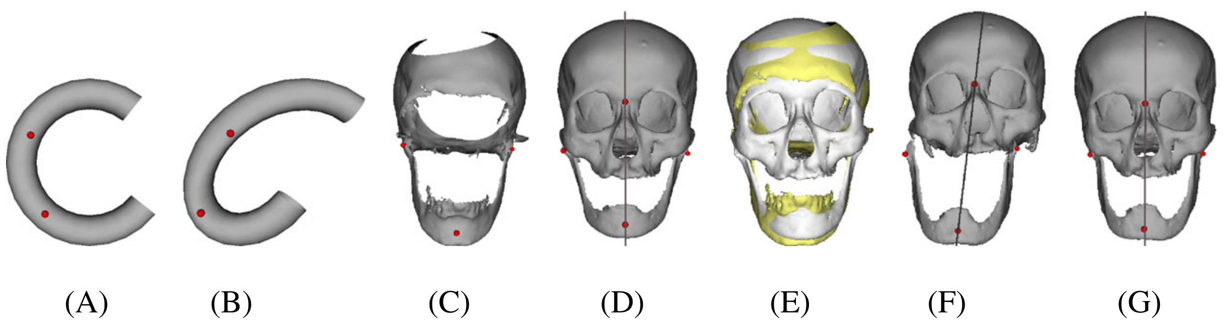

FIGURE 2 Effect of Laplacian deformation. A, C-shape model with 2 landmarks. B, Laplacian deformation of C-shape model with the bottom landmark fixed and the left landmark moved to the right. C, Defective skull model with large missing parts. D, Reference model that is disconnected on the right side like the C-shape model. The straight line denotes the symmetric mid-plane. E, Initial alignment. The reference is white and the target is yellow. F, Registration of reference to target using ordinary Laplacian deformation produces a distorted result. G, Laplacian deformation that preserves the vertical mid-plane produces an undistorted result

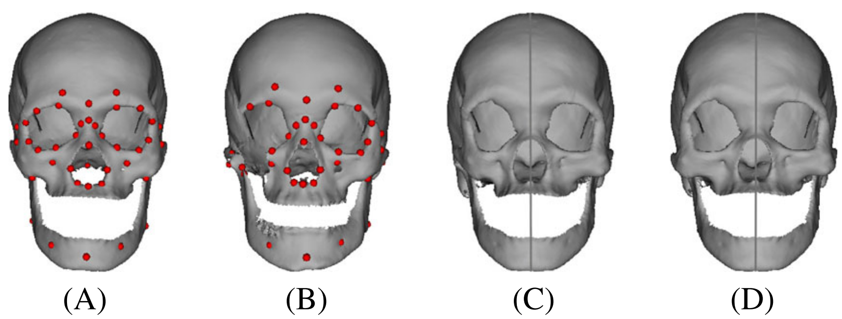

FIGURE 3 Laplacian deformation with the symmetry constraint. A, Reference model. B, Target model with missing facial bone. C, Laplacian deformation with the mid-plane constraint produces a result whose right orbit (in the left side of the image) is distorted, although the mid-plane remains vertical. D, Laplacian deformation with mid-plane and symmetry constraints produces a symmetric and undistorted result

$$L(\mathbf{p}_i) = \sum_{j \in \mathcal{N}_i} w_{ij}(\mathbf{p}_i - \mathbf{p}_j) \tag{5}$$

where $\mathcal{N}_i$ is the set of connected neighbours of vertex i. The weight w_{ij} can be cotangent weight[35] or equal weight $w_{ij} = 1/|\mathcal{N}_i|$. Equal weight leads to simpler optimization equations. With equal weight, the Laplacian operator becomes

$$L(\mathbf{p}_i) = \mathbf{p}_i - \frac{1}{\mathcal{N}_i} \sum_{j \in \mathcal{N}_i} \mathbf{p}_j. \tag{6}$$

Laplacian deformation preserves the model's shape by minimizing the difference of Laplacian operators $L(\mathbf{p}_i^0)$ before deformation and $L(\mathbf{p}_i)$ after it, which is $\|L(\mathbf{p}_i) - L(\mathbf{p}_i^0)\|^2$. These differences for all n mesh vertices are organized into a matrix equation: $\mathbf{Lx} = \mathbf{a}$, where $\mathbf{L}$ is a $3n \times 3n$ matrix that captures the Laplacian constraints, $\mathbf{x}$ is a $3n \times 1$ vector of unknown positions $\mathbf{x}_i$ of mesh vertices, $\mathbf{x} = [\mathbf{x}_1^\top \quad \cdots \quad \mathbf{x}_n^\top]^\top$, and $\mathbf{a}$ is a $3n \times 1$ vector that contains $L(\mathbf{p}_i^0)$ before deformation.

The mid-plane constraint is imposed as follows. Without loss of generality, let the skull model before deformation be oriented such that the mid-plane is located at $x = 0$ and its surface normal is parallel to the x-axis. Moreover, the landmark points coincide with some mesh vertices. Then, the mid-plane constraint requires that the x-coordinates of the mid-point landmarks remain as 0 after deformation, which constrains the mid-plane to remain vertical and laterally symmetric after deformation. For a non-defective target skull, the mid-point landmarks of a reference skull always have corresponding landmarks on the target. On the other hand, for a target skull with a large number of missing or defective facial bones, it is impossible to place mid-point landmarks on the missing or defective parts. In this case, some reference mid-point landmarks will not have corresponding target landmarks. Then, the mid-plane constraint has to be imposed on these reference mid-point landmarks that do not have correspondence. THe mid-plane constraint is organized into a matrix equation: $\mathbf{Mx} = \mathbf{0}$, where $\mathbf{M}$ is a $k \times 3n$ matrix and k is the number of mid-point landmarks without correspondence. The entries in $\mathbf{M}$ that correspond to the x-components of mid-point landmarks without correspondence are set to 1; all other entries are set to 0.

The symmetry constraint is imposed as follows. For every pair (l, r) of landmarks that are symmetric with respect to the mid-plane, which is the y–z plane, their coordinates after deformation should have the relationships: $x_l + x_r = 0$, $y_l - y_r = 0$ and $z_l - z_r = 0$. These relationships can be organized into a matrix equation: $\mathbf{Sx} = \mathbf{0}$, where $\mathbf{S}$ is a $3s \times 3n$ matrix and s is the number of symmetric landmark pairs. The entries in $\mathbf{S}$ that correspond to x_l, x_r, y_l and z_l are set to 1; those that correspond to y_r and z_r are set to -1; all other entries are set to 0.

The Laplacian constraint, the mid-plane constraint and the symmetry constraint are combined together into the following objective function to be minimized:

$$\|\mathbf{Ax} - \mathbf{b}\|^2 = \left\| \begin{bmatrix} \mathbf{L} \\ \mathbf{M} \\ \mathbf{S} \end{bmatrix} \mathbf{x} - \begin{bmatrix} \mathbf{a} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \right\|^2. \tag{7}$$

The positional constraints of the corresponding points between the reference and the target are organized into a matrix equation of the form $\mathbf{Cx} = \mathbf{d}$, where $\mathbf{C}$ indicates the mesh vertices with positional constraints and $\mathbf{d}$ contains the desired vertex positions. Without loss of generality, we can arrange the mesh vertices with positional constraints as vertices 1 to have $m < n$. Then, $\mathbf{C}$ is a $3m \times 3n$ matrix that contains a $3m \times 3m$ identity matrix and a $3m \times 3(n-m)$ zero matrix: $\mathbf{C} = [\mathbf{I}_{3m}\ \mathbf{0}]$. Correspondingly, the top $3m$ elements of $\mathbf{x}$ are the mesh vertices with positional constraints, the bottom $3(n-m)$ elements are those without positional constraints, and $\mathbf{d}$ is a $3m \times 1$ vector of the coordinates of the desired vertex positions. Then, Laplacian deformation with the mid-plane constraint and symmetry constraint solves the following problem:

$$\min_{\mathbf{x}} \|\mathbf{Ax} - \mathbf{b}\|^2 \text{ subject to } \mathbf{Cx} = \mathbf{d}. \tag{8}$$

That is, the Laplacian, mid-plane and symmetry constraints are soft constraints whereas the positional constraints are hard ones.

This Laplacian deformation problem is an *equality-constrained least-squares* problem, which can be solved using QR factorization[35,36] as follows: $\mathbf{C}^\top$ has QR factorization $\mathbf{C}^\top = \mathbf{QR}$, where $\mathbf{Q} = [\mathbf{Q}_1\ \mathbf{Q}_2]$ is orthogonal and $\mathbf{R} = [\mathbf{R}_1^\top \quad \mathbf{0}^\top]^\top$ is upper-triangular. Define vectors $\mathbf{u}$ and $\mathbf{v}$ such that

$$\mathbf{x} = \mathbf{Q} \begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix} = \begin{bmatrix} \mathbf{Q}_1 & \mathbf{Q}_2 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix}. \tag{9}$$

Then, the objective function of (8) becomes

$$\|\mathbf{Ax} - \mathbf{b}\|^2 = \|\mathbf{AQ}_1\mathbf{u} + \mathbf{AQ}_2\mathbf{v} - \mathbf{b}\|^2. \tag{10}$$

Since $\mathbf{C} = [\mathbf{I}_{3m}\mathbf{0}]$, the QR factorization of $\mathbf{Q}^\top$ is

$$\mathbf{C}^\top = \begin{bmatrix} \mathbf{I}_{3m} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{3(n-m)} \end{bmatrix} \begin{bmatrix} \mathbf{I}_{3m} \\ \mathbf{0} \end{bmatrix}. \tag{11}$$

That is,

$$\mathbf{Q}_1 = \begin{bmatrix} \mathbf{I}_{3m} \\ \mathbf{0} \end{bmatrix}, \qquad \mathbf{Q}_2 = \begin{bmatrix} \mathbf{0} \\ \mathbf{I}_{3(n-m)} \end{bmatrix}, \qquad \mathbf{R}_1 = \mathbf{I}_{3m}.$$

With QR factorization of $\mathbf{C}^\top$, the positional constraint equation $\mathbf{Cx} = \mathbf{d}$ becomes

$$\mathbf{Cx} = \mathbf{R}^\top\mathbf{Q}^\top\mathbf{x} = \mathbf{R}^\top \begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix} = \mathbf{R}_1^\top\mathbf{u} = \mathbf{d}. \tag{12}$$

The right-hand side of equation 12 yields $\mathbf{I}_{3m}\mathbf{u} = \mathbf{u} = \mathbf{d}$. So,

$$\mathbf{Q}_1\mathbf{u} = \begin{bmatrix} \mathbf{I}_{3m} \\ \mathbf{0} \end{bmatrix} \mathbf{u} = \begin{bmatrix} \mathbf{I}_{3m} \\ \mathbf{0} \end{bmatrix} \mathbf{d} = \begin{bmatrix} \mathbf{d} \\ \mathbf{0} \end{bmatrix}, \tag{13}$$

$$\mathbf{Q}_2\mathbf{v} = \begin{bmatrix} \mathbf{0} \\ \mathbf{I}_{3(n-m)} \end{bmatrix} \mathbf{v} = \begin{bmatrix} \mathbf{0} \\ \mathbf{v} \end{bmatrix}. \tag{14}$$

Organize the matrix $\mathbf{A}$ as $[\mathbf{A}_1\mathbf{A}_2]$. Then,

$$\mathbf{AQ}_1\mathbf{u} = \mathbf{A}_1\mathbf{d}, \qquad \mathbf{AQ}_2\mathbf{v} = \mathbf{A}_2\mathbf{v}. \tag{15}$$

Then, the objective function (10) becomes

$$\|\mathbf{A}_2\mathbf{v} - (\mathbf{b} - \mathbf{A}_1\mathbf{d})\|^2. \tag{16}$$

Minimization of objective function (16) with linear least squares yields

$$\mathbf{v} = (\mathbf{A}_2^\top\mathbf{A}_2)^{-1}\mathbf{A}_2^\top(\mathbf{b} - \mathbf{A}_1\mathbf{d}). \tag{17}$$

Then, the positions of the mesh vertices after deformation can be computed as

$$\mathbf{x} = \mathbf{Q}_1\mathbf{u} + \mathbf{Q}_2\mathbf{v} = \begin{bmatrix} \mathbf{d} \\ \mathbf{v} \end{bmatrix}. \tag{18}$$

As $\mathbf{d}$ contains known desired positions, Laplacian deformation only needs to solve for $\mathbf{v}$, the coordinates of the mesh vertices without positional constraints after deformation. Therefore, Laplacian deformation runs faster with an increasing number of positional constraints.

5 | FLIP-AVOIDING INTERPOLATING SURFACE

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. FAIS is summarized in Algorithm 1.

Algorithm 1: Flip-avoiding interpolating surface

Input: Reference F, target T, manually marked positional constraints of landmarks C^*.

1 Rigidly register F to T using reference landmarks with correspondence C^*.
2 Orientate F and T so that the mid-plane of F is located at $x = 0$ and its normal is parallel to the x-axis.
3 Non-rigidly register F to T with positional constraints C^* using LD, then set R as registered F.
4 **for** *k from 1 to K* **do**
5   Find correspondence C from R to T using first correspondence search method.
6   Choose a sparse subset C^+ from $C^* \cup C$.
7   Non-rigidly register R to T with constraints C^+ using LD in incremental steps.
8 **end**
9 Find correspondence C from R to T using second correspondence search method.
10 Remove crossings in $C^* \cup C$ giving C^+.
11 Non-rigidly register R to T with constraints C^+ using LD.

Output: Reconstructed model R.

Step 1 rigidly registers the reference F to target T with known correspondence C^* to normalize their global size, location and orientation. Step 2 orientates R and T for computing mid-plane and symmetry constraints.

Step 3 non-rigidly registers reference F to target T using Laplacian deformation (LD), 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 4 to 8 perform K iterations of non-rigid registration in small steps. First, step 5 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 3.2). Step 6 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 $D \leq D_1$. 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 $\|\mathbf{p} - \mathbf{q}\| \leq 2D$, 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 7 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 towards $\mathbf{p}'$ incrementally, allowing FAIS to recover from possible wrong correspondence in subsequent iterations.

Step 9 finds correspondence C from R to T using the second correspondence search method (Section 3.2). Step 10 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 6 (Section 7.2). Finally, step 11 performs the final registration of R to T with C^+ as the positional constraints.

FAIS differs from non-rigid ICP,[18,19] although they have similar iterative structure. Non-rigid ICP performs locally affine registration of the approximating surface, which has no surface flipping problem. On the other hand, FAIS performs non-rigid registration of the interpolating surface and needs to avoid surface flipping.

6 | CLUSTERING OF REFERENCE MODELS

To reconstruct a target accurately, the shape of the reference model used should be as close to the target as possible so that the reconstruction of defective parts matches its expected normal shape. To achieve this goal, one can exhaustively try every reference candidate for the same target and pick the reference model with the smallest reconstruction error. However, this exhaustive way of choosing a reference model is time-consuming. Alternatively, it is possible to group the reference models into a small number of clusters, and use the cluster prototypes as the reference candidates. This will greatly speed up the process of reference selection because the number of cluster prototypes is much smaller than the number of reference models in total.

The most important components of a clustering algorithm are (1) the distance measure between two items, and (2) the cluster prototype, which is typically the average of the items in a cluster. In our application, the skull models are mesh models constructed from patients' CT images. The most natural distance $d(S;R)$ between two mesh models S and R is the error of registering R to S, which can be measured as the average distance from the points on the registered R to the surface of S. Since S and R have different mesh connectivities in general, the difference $d(S;R)$ is positive but non-symmetric.

The difference in mesh connectivity makes it impossible to compute the average of multiple meshes by averaging the positions of their mesh vertices. To compute the average in this way, it is necessary to first resample the mesh models so that they have the same mesh connectivity. This resampling process is not only tedious but also time-consuming. Therefore, a different definition of cluster prototype is required.

We define the cluster prototype P_j of cluster C_j as the model in C_j that has the shortest average distance to all other models in C_j:

$$P_j = \arg\min_{R \in C_j} \frac{1}{|C_j|} \sum_{S \in C_j} d(S;R). \tag{19}$$

This definition is a form of *generalized mean*. In fact, the vector mean in vector space satisfies this definition of generalized mean, with the squared error of the vector as the distance measure. This fact can be proved as follows.

Proof. The average squared error of $\mathbf{v}_i$, $i = 1, \ldots, n$, with respect to a vector $\mathbf{u}$ is given by

$$\frac{1}{n}\sum_{i=1}^{n}\|\mathbf{v}_i - \mathbf{u}\|^2. \tag{20}$$

Expanding equation (20) gives

$$\frac{1}{n}\sum_{i=1}^{n}(\mathbf{v}_i - \mathbf{u})^\top(\mathbf{v}_i - \mathbf{u}) = \frac{1}{n}\sum_{i=1}^{n}(\mathbf{v}_i^\top\mathbf{v}_i - 2\mathbf{v}_i^\top\mathbf{u} + \mathbf{u}^\top\mathbf{u}). \tag{21}$$

To obtain the minimal $\mathbf{u}$, differentiate equation (21) with respect to $\mathbf{u}$ and equate to 0, which yields

$$\frac{1}{n}\sum_{i=1}^{n}\mathbf{u} = \frac{1}{n}\sum_{i=1}^{n}\mathbf{v}_i. \tag{22}$$

Therefore, the left-hand side of equation (22), which is $\mathbf{u}$, is equal to the mean of $\mathbf{v}_i$. □

With the distance measure and generalized mean given above, the generalized k-means clustering algorithm for skull models can be summarized as in Algorithm 2. After clustering, given a target model T, the cluster prototype P_j that is nearest to T, measured in terms of distance $d(T; P_j)$, is the selected reference. We use $k = 4$ in the experiment.

Algorithm 2: Generalized k-means clustering

Input: A set U of non-defective skull models.
1 Compute distances $d(S; R)$ between all pairs of S and R in U.
2 Randomly pick k models in U as cluster prototypes P_j.
3 **repeat**
4 Assign each model S in U to the nearest cluster C_j with the smallest distance $d(S, P_j)$.
5 Update cluster prototype P_j of each cluster C_j according to equation 19.
6 **until** *convergence*;
Output: Cluster prototypes $P_j, j = 1, \ldots, k$.

7 | EXPERIMENTS

7.1 | Data preparation and PC configuration

Sixty-two 3D mesh models of non-defective skulls were constructed from patients' CT images. Two of the non-defective skulls, one with teeth and one without, were used as the reference models (Figure 4A,B). Ten of the non-defective skulls were each used to synthesize 5 types of defective skulls with different sizes and locations of fractures (Figure 4C–G), giving a total of 50 defective testing skulls. These synthetic skulls were generated by displacing bone fragments in a way similar to real fractures. More specifically, types S1–S3 of the synthetic skulls have increasing sizes of fractures; S1 and S5, as well as S2 and S4, have roughly the same fracture size but different locations. The 10 non-defective skulls served as the ground truth. The other 50 non-defective skulls were used to build an active shape model for performance comparison. In addition, 6 skull models of trauma patients with real fractures (Figure 8A) and 3 defective skull models with complications (Figure 9A) were used for the reconstruction test. To evaluate the robustness of FAIS, 5 non-defective skulls and 2 defective skulls with dental artifacts were used for testing. Each of the above skull models had up to 56 landmarks manually placed on them. Landmarks on defective parts were omitted because defective parts provide wrong information that may mislead the reconstruction algorithms.

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

7.2 | Dense correspondence

In this experiment, FAIS was tested in turn with Laplacian deformation and TPS as the non-rigid registration method (at steps 3, 7 and 11) on a normal testing skull. Figure 5A shows that FAIS with Laplacian deformation finds more corresponding points than FAIS with TPS. This is because the former is more accurate than the latter (Section 7.3). During the iterative stage from step 4 to step 8, up to 90% of the corresponding points are rejected by the simple no-crossing condition. At step 10, 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[11-14] 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 5B shows that Laplacian deformation runs faster and TPS runs slower with increasing number of hard constraints. This is because Laplacian deformation's running time is dependent on the number of mesh vertices without positional constraints (Section 4), whereas TPS's is dependent on the number of mesh vertices with positional constraints. 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.

FAIS with Laplacian deformation takes about 3 minutes to run, with Laplacian deformation requiring two-thirds of the running time and correspondence search and flip avoidance taking one-third. Therefore, most of the computation time is spent on performing Laplacian deformation. On the other hand, TPS deformation alone takes a total running time of more than 8 minutes for the first 20 iterations. In other words, TPS runs more than 4 times slower than Laplacian deformation. Moreover, a test was performed with another linear algebra library called Armadillo. The test result shows that Armadillo runs 10 times slower than Intel MKL.

7.3 | Reconstruction of synthetic fractured skulls

This experiment compares the reconstruction accuracy of FAIS with other skull reconstruction algorithms on synthetic fractured skulls given the same reference model. Fifty synthetic fractured skulls were used, whose fractured parts were marked manually. The normal skulls

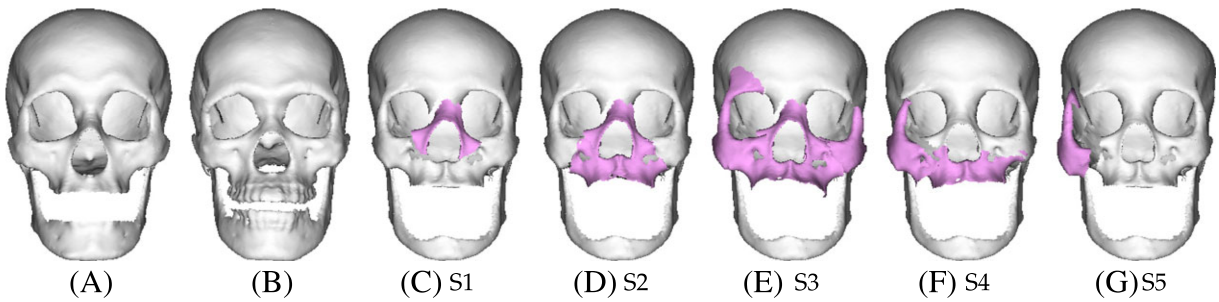

FIGURE 4 Skull models. A, Reference model without teeth. B, Reference with teeth. C–G, Synthetic skull models of types S1–S5 with fractures of different sizes and locations generated from normal skull models. Defective parts are marked in pink

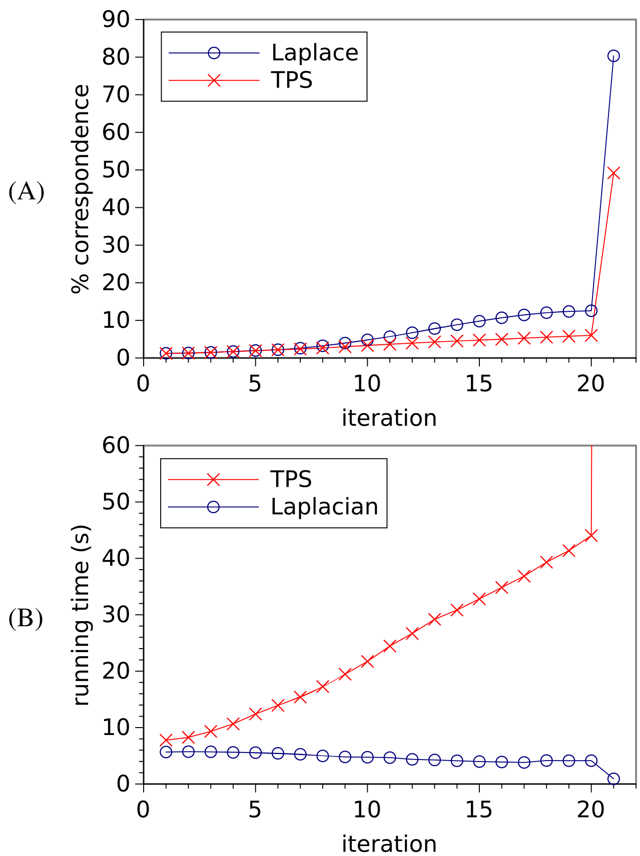

FIGURE 5 Correspondence search. A, Amount of correspondence at various iterations. B, Running time is dependent on the amount of correspondence

used to generate the synthetic skulls served as the ground truths. Each testing skull was reconstructed by the following algorithms:

- FAIS: The proposed method with symmetry constraints.
- FAIS-0: The proposed method without symmetry constraints.[24] This is achieved using Laplacian deformation without mid-plane and symmetry constrains.
- LD: Laplacian deformation with mid-plane and symmetry constraints.
- TPS-1: Same as steps 1–3 of FAIS except TPS is used as the non-rigid registration algorithm; similar to Deng et al.,[11] Lapeer and Prager[12] and Rosas and Bastir.[13]
- TPS-2: Same as steps 1–8 of FAIS with $K = 1$ except TPS is used as the non-rigid registration algorithm; similar to Zhang et al.[14]
- ASM-F: Active shape model using FAIS-0 for resampling.
- ASM-T: Active shape model using TPS-2 for resampling; as in Zhang et al.[6]

To construct the active shape model (ASM), 50 normal skull models were resampled so that they had the same number of vertices and the same mesh connectivity. Two resampling methods were used to test their effects on ASM's reconstruction accuracy, namely FAIS-0 and TPS-2. For ASM-F and ASM-T, resampling was achieved by registering a normal reference to each normal skull using FAIS-0 and TPS-2, respectively. Since the models used to construct ASM are all non-defective, FAIS-0 gives the same results as FAIS. The reconstruction results of these algorithms were recorded. Reconstruction errors were measured between the reconstructed models and their ground truths. Reconstruction errors of the defective and non-defective parts of testing skulls were measured separately.

Table 1 summarizes the reconstruction errors of synthetic fractured skulls. FAIS has the smallest reconstruction errors E_N on non-defective parts and E_D on defective parts among all the skull reconstruction algorithms. Its E_N is slightly larger than zero because its flip-avoiding method removes some conflicting positional constraints. Thus, only those mesh vertices with positional constraints are registered exactly with zero error.

FAIS's reconstruction accuracy on non-defective parts is not affected by the size and location of defective parts. This is indicated by the roughly identical E_N across all different types of synthetic skulls. FAIS's reconstruction accuracy on defective parts is also not affected by the location of defective parts. This can be seen from the roughly identical E_D between S1 and S5 or between S2 and S4, which have similar size of defective parts. On the other hand, FAIS's reconstruction accuracy on defective parts is affected by the size of the defective parts, as indicated by the increasing E_D from S1 to S3, and from S5 to S3.

7.3.1 | FAIS vs. FAIS-0

Both FAIS and FAIS-0 have the smallest reconstruction errors on non-defective parts (Table 1). Their E_N are comparable but FAIS is more robust than FAIS-0, which is indicated by its smaller standard deviation. Moreover, FAIS has smaller reconstruction errors E_D on the defective parts than does FAIS-0, because it yields more symmetric and normal

TABLE 1 Mean reconstruction errors and standard deviations (in mm) of synthetic fractured skulls

(a) Errors E_N of non-defective parts	S1	S2	S3	S4	S5	All
FAIS	0.13 ± 0.04	0.12 ± 0.03	0.13 ± 0.05	0.12 ± 0.03	0.13 ± 0.03	0.13 ± 0.04
FAIS-0	0.12 ± 0.03	0.11 ± 0.03	0.19 ± 0.14	0.12 ± 0.03	0.12 ± 0.03	0.13 ± 0.07
LD	2.31 ± 0.41	2.31 ± 0.43	2.70 ± 0.57	2.37 ± 0.42	2.37 ± 0.39	2.41 ± 0.46
TPS-1	3.41 ± 1.00	3.43 ± 1.03	2.83 ± 0.69	3.17 ± 0.86	3.12 ± 0.89	3.19 ± 0.89
TPS-2	1.50 ± 0.86	1.40 ± 0.79	0.79 ± 0.32	1.14 ± 0.63	1.12 ± 0.54	1.19 ± 0.68
ASM-F	1.30 ± 0.18	1.30 ± 0.16	1.41 ± 0.26	1.33 ± 0.19	1.29 ± 0.17	1.32 ± 0.19
ASM-T	2.01 ± 0.41	2.00 ± 0.40	1.98 ± 0.32	1.89 ± 0.36	1.87 ± 0.32	1.95 ± 0.35
(b) Errors E_D of defective parts						
FAIS	0.20 ± 0.05	0.52 ± 0.19	1.42 ± 0.37	0.49 ± 0.10	0.18 ± 0.04	0.56 ± 0.49
FAIS-0	0.19 ± 0.03	0.55 ± 0.08	1.85 ± 0.65	0.59 ± 0.11	0.23 ± 0.05	0.68 ± 0.68
LD	2.14 ± 0.52	2.18 ± 0.47	3.96 ± 0.96	2.15 ± 0.48	2.10 ± 0.51	2.51 ± 0.94
TPS-1	3.26 ± 0.90	3.19 ± 0.79	2.76 ± 0.57	2.89 ± 0.64	2.84 ± 0.78	2.99 ± 0.74
TPS-2	1.54 ± 0.65	1.57 ± 0.48	1.51 ± 0.25	1.37 ± 0.30	1.18 ± 0.45	1.43 ± 0.46
ASM-F	1.27 ± 0.18	1.44 ± 0.22	1.61 ± 0.23	1.44 ± 0.18	1.25 ± 0.14	1.40 ± 0.23
ASM-T	1.96 ± 0.30	2.09 ± 0.27	2.10 ± 0.20	1.96 ± 0.25	1.81 ± 0.25	1.98 ± 0.27

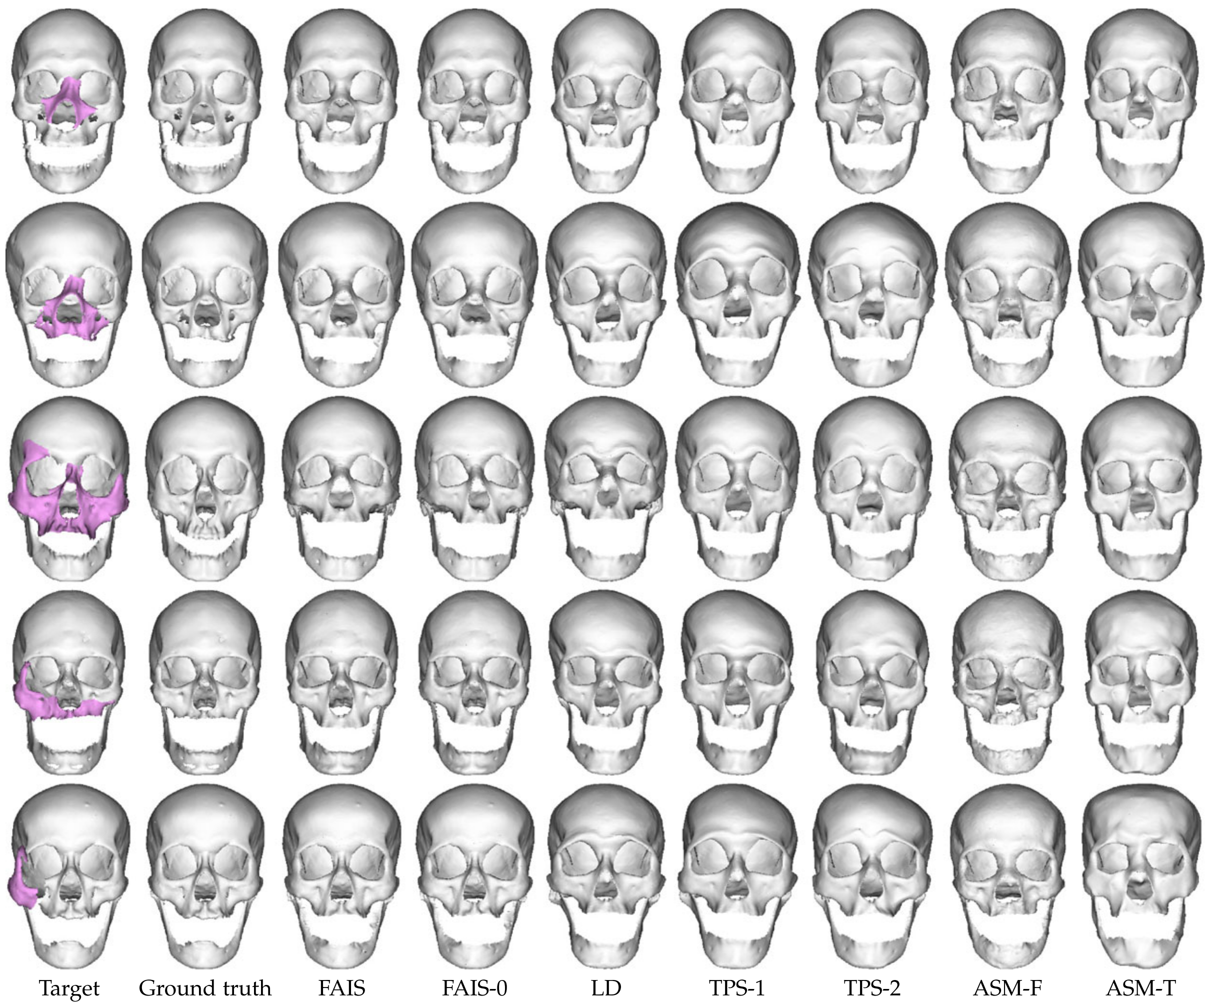

FIGURE 6 Sample reconstruction results of synthetic fractured skulls. Fractured parts of synthetic target skulls are shown in pink. Column 2 illustrates ground truths and columns 3–9 show reconstruction results of various algorithms

reconstruction than FAIS-0. Thus FAIS is more accurate and robust than FAIS-0 for skull reconstruction.

Figure 6 shows sample reconstruction results. Except for the third model, which has a very large number of fractures, FAIS's reconstructed models look identical to the ground truths. The surfaces of reconstructed models are smooth. FAIS's reconstructed models are all symmetric and normal. In comparison, FAIS-0's reconstructed models of the targets in rows 2, 3 and 4 are laterally distorted. Nevertheless, the reconstructed models are visually similar to the non-defective parts of the targets. Therefore, FAIS-0 can be used for resampling of the non-defective parts although it is not as accurate as FAIS for reconstruction of defective parts.

7.3.2 | FAIS vs. other non-rigid registration algorithms

Compared to the other non-rigid registration algorithms such as LD, TPS-1 and TPS-2, FAIS has significantly smaller reconstruction errors E_N on the non-defective parts (Table 1) because it applies a very dense set of correspondence and iterative incremental registration. On the other hand, LD, TPS-1 and TPS-2 use many fewer corresponding points. Moreover, they have larger reconstruction errors E_D on the defective parts than FAIS.

The reconstructed models of FAIS look visually close to the ground truths. In comparison, those of LD, TPS-1 and TPS-2 do not look like their corresponding ground truths. Therefore, FAIS is more accurate than the other non-rigid registration algorithms for skull reconstruction.

7.3.3 | FAIS vs. active shape model

ASM-F has smaller reconstruction errors E_N and E_D than ASM-T. This is expected because ASM-F uses FAIS-0 for mesh resampling, and FAIS-0 has significantly smaller reconstruction errors E_N than TPS-2, which is used by ASM-T for mesh resampling. ASM-F's smaller reconstruction errors are also attributed to its smaller resampling errors. The resampling error was measured as the mean distance between the resampled skull and its corresponding normal target skull. The resampling errors of ASM-F and ASM-T are 0.26±0.23 mm and 2.11±1.10 mm, respectively. ASM-F's resampling is more accurate and stable than those of ASM-T. This means that FAIS-0 is a better choice than TPS-2 for model resampling. For the reconstruction of defective parts, ASM-F has smaller E_D than LD, TPS-1, TPS-2 and ASM-T. But its E_N and E_D are much higher than those of FAIS.

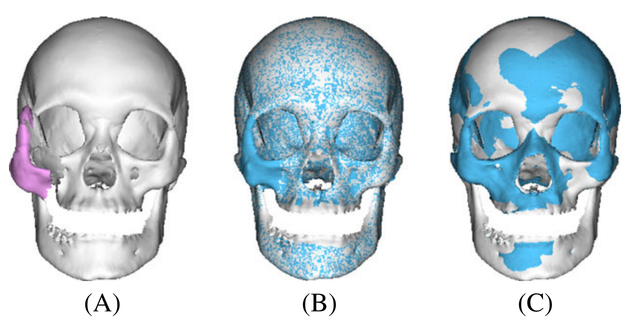

FIGURE 7 FAIS vs. ASM-F. A, Target. B, Superimposition of the non-defective parts of the target (white) and the reconstruction of FAIS (blue). C, Superimposition of the non-defective parts of the target (white) and the reconstruction of ASM-F (blue)

FAIS's reconstruction errors E_N are about 10 times smaller than those of ASM-F and ASM-T (Table 1). Its reconstruction errors E_D are about 3 times smaller than those of ASM-F and ASM-T. Figure 6 shows that FAIS's reconstruction is close to the ground truths, whereas those of ASM-F and ASM-T are not. Figure 7 shows that the reconstruction of FAIS fits the non-defective parts of the target very tightly, which is indicated by the alternating colour pattern on the superimposition. On the other hand, the reconstruction of ASM-F does not fit the target well. This is because FAIS registers an interpolating surface to the target whereas ASM generates an approximating surface. This suggests that ASM-F can be used to estimate the overall normal shape of the target but cannot be used for accurate reconstruction. In comparison, FAIS is more accurate for skull reconstruction.

7.4 | Reconstruction of real defective skulls

This experiment compares the reconstruction accuracy of FAIS, FAIS-0 and ASM-F, the 3 most accurate algorithms in the previous tests. These 3 algorithms were tested on 6 real defective skulls T1–T6 with post-op models and 3 defective cases T7–T9 with complications (Figure 9) but without post-op models. The test procedure was the same as that in Section 7.3. Discrepancy between the reconstructed parts generated by the algorithms and the post-op models was measured as the average distance between their surfaces.

Figure 8 shows the reconstruction results in order of increasing discrepancy for FAIS. T2 and T4 have minor fractures of the cheek bones. For these cases, the reconstruction results of FAIS and FAIS-0 are visually similar. For T1, T3 and T5, which are severely fractured, the reconstruction results of FAIS-0 are laterally distorted because it does not incorporate symmetry constraints. In comparison, FAIS's reconstruction results of these models look normal and symmetric. T6 has a defective left frontal bone due to a tumour. Again, FAIS-0's reconstruction is not symmetric whereas FAIS's reconstruction is more symmetric. In all cases, ASM-F's reconstruction results look normal and symmetric but they do not match their corresponding targets well. Instead, all its reconstruction results look more like the reference model (Figure 4A).

Table 2 shows that the discrepancy between FAIS and post-op models is small for T1 to T5. Its discrepancy for T6 is large because the surgically operated left frontal-sinus bone looks flat and the top of the left eye socket is lower than the right side. The post-op model shows that it is very difficult for the surgical operation on this skull to achieve good results without a good reference model. Our reconstructed model could have improved the surgical outcome if it was available during the surgical operation. The discrepancies of FAIS-0 are larger than those of FAIS. In particular, its discrepancy for T1 is very large because its reconstruction of T1 is highly distorted (Figure 8). ASM-F's discrepancies of T3–T6 are slightly smaller than those of FAIS. But as discussed above, ASM-F's reconstruction does not match the target shape well.

Figure 9 illustrates the reconstruction of FAIS, FAIS-0 and ASM-F on real defective skulls with complications. Target T7 is the defective skull

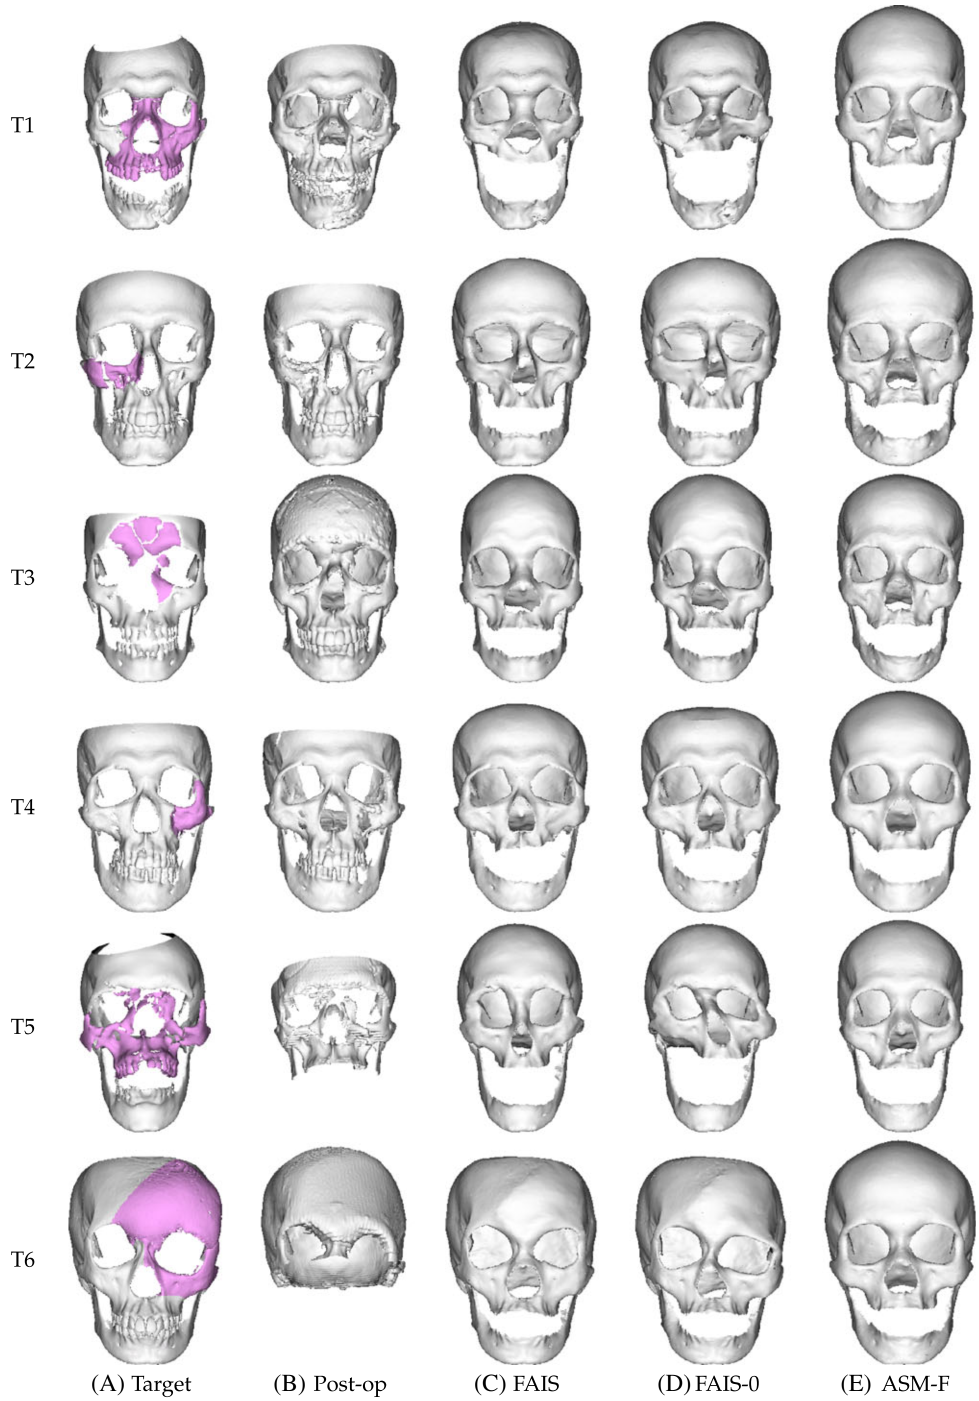

FIGURE 8 Reconstruction results of real defective skulls with post-op models. Defective parts of targets are shown in pink. B, Post-ops. C–E, Reconstruction results of various algorithms

TABLE 2 Discrepancy (in mm) between reconstructed parts generated by the algorithms and the post-op models

	T1	T2	T3	T4	T5	T6
FAIS	1.88	1.93	2.13	2.54	2.85	5.84
FAIS-0	8.44	2.74	2.42	2.94	4.70	5.31
ASM-F	3.22	4.43	1.71	2.39	2.48	5.15

of a 3-year-old child with a fractured orbital bone. Even though its shape differs greatly from that of the reference model, which is an adult skull, the reconstructed model of FAIS is still symmetric and normal. On the other hand, the reconstructed model of ASM-F looks different from the target. Target T8 is the defective skull of a patient who had a very severe orbital tumour. FAIS is able to reconstruct a normal and symmetric skull that fits the non-defective parts well. In comparison, FAIS-0's reconstruction is laterally distorted, and ASM-F's reconstruction looks reasonable for this case. Target T9 is a deformed skull with sunken facial bone. The shape of its cranium is also deformed. With almost no information to utilize, FAIS can still reconstruct a skull with normal facial bone. In comparison, FAIS-0's reconstruction is laterally distorted, and ASM-F's reconstruction is highly distorted because the abnormal shape of the target is not captured by the statistical model.

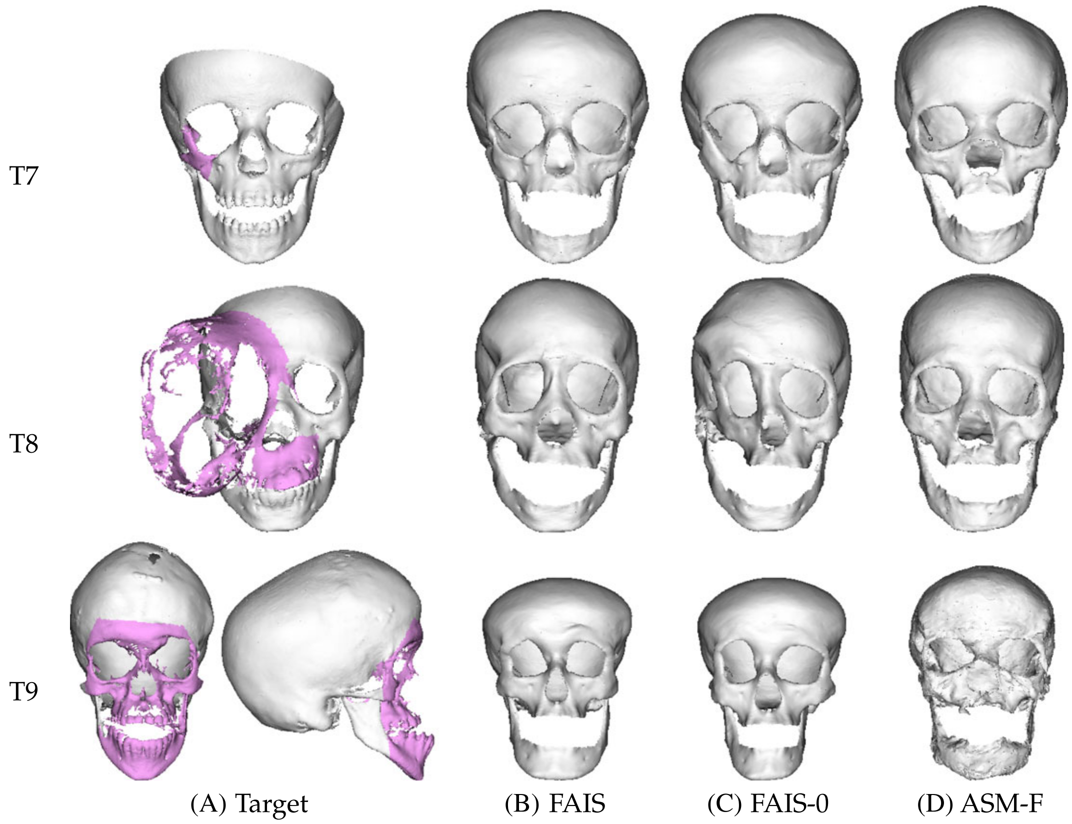

FIGURE 9 Reconstruction results of real defective skulls with complications. Post-op models are not available. Defective parts of targets are shown in pink. B–D, The reconstruction results of various algorithms

7.5 | Robustness against outliers

CT images that are used to construct 3D skull models can contain radiation artifacts caused by metallic dental implants[37] that are very difficult to remove. Thus, 3D skull models segmented and constructed from CT images often contain metal artifacts (Figure 10). This test evaluates the robustness of FAIS and FAIS-0 against outliers such as metal artifacts. The test was performed on 7 actual skulls with metal artifacts. Among the testing skulls, 5 were normal and the other 2 were fractured. A normal skull with teeth (Figure 4B) was used as the reference model. For this test, the algorithms differ slightly such that the first, instead of the second, correspondence search method was applied on the mesh vertices in the teeth region at step 9 of the algorithms.

Test results in Figure 10 show that the first correspondence search is robust enough to exclude metal artifacts as possible corresponding points. Consequently, FAIS and FAIS-0 do not register the reference to the metal artifacts, and the reconstruction results are free of metal artifacts. Some of the teeth reconstructed by FAIS and FAIS-0 are slightly distorted because the skull mesh models have insufficient resolution to model each tooth accurately. For the normal target skulls (Figure 10(M1–M5)), the reconstruction results of FAIS and FAIS-0 look visually similar. For the defective target skulls (Figure 10(M6,M7)), the reconstruction results of FAIS are more symmetric and normal than those of FAIS-0, which is consistent with the test results on the reconstruction of synthetic skulls.

7.6 | Effect of reference model

This experiment investigates how the selection of reference models affects FAIS's reconstructed results. The 5 synthetic skulls of type S3, which are the most difficult cases, were used as target models, denoted as D1–D5. We measured FAIS's reconstruction errors E_D of defective parts given the reference models generated in the following ways.

- Ideal. Among the 50 normal skulls, the model with the smallest E_D is used as the reference model. In practice, this method is not feasible because E_D cannot be measured without ground truths.
- Exhaustive. Among the 50 normal skulls, the model with the smallest E_N is used as the reference model. This method is feasible in practice.
- Best prototype. Among the 4 cluster prototypes, the one with the smallest E_N is used as the reference model.
- ASM-F's reconstruction. The reconstructed model of ASM-F is used as the reference model.

Table 3 shows that, on average, the exhaustive method has a reconstruction error close to that of the ideal method. The method that uses the best prototype has a reconstruction error comparable to but larger than that of the exhaustive method. In comparison, using ASM's reconstructed model as the reference for FAIS results in the largest reconstruction error on average. Compared to the test results in Table 1(b), we can conclude that the reference model used for the tests that produce Table 1(b), which is chosen randomly, is better than the best prototype but slightly worse than the one selected by the exhaustive method.

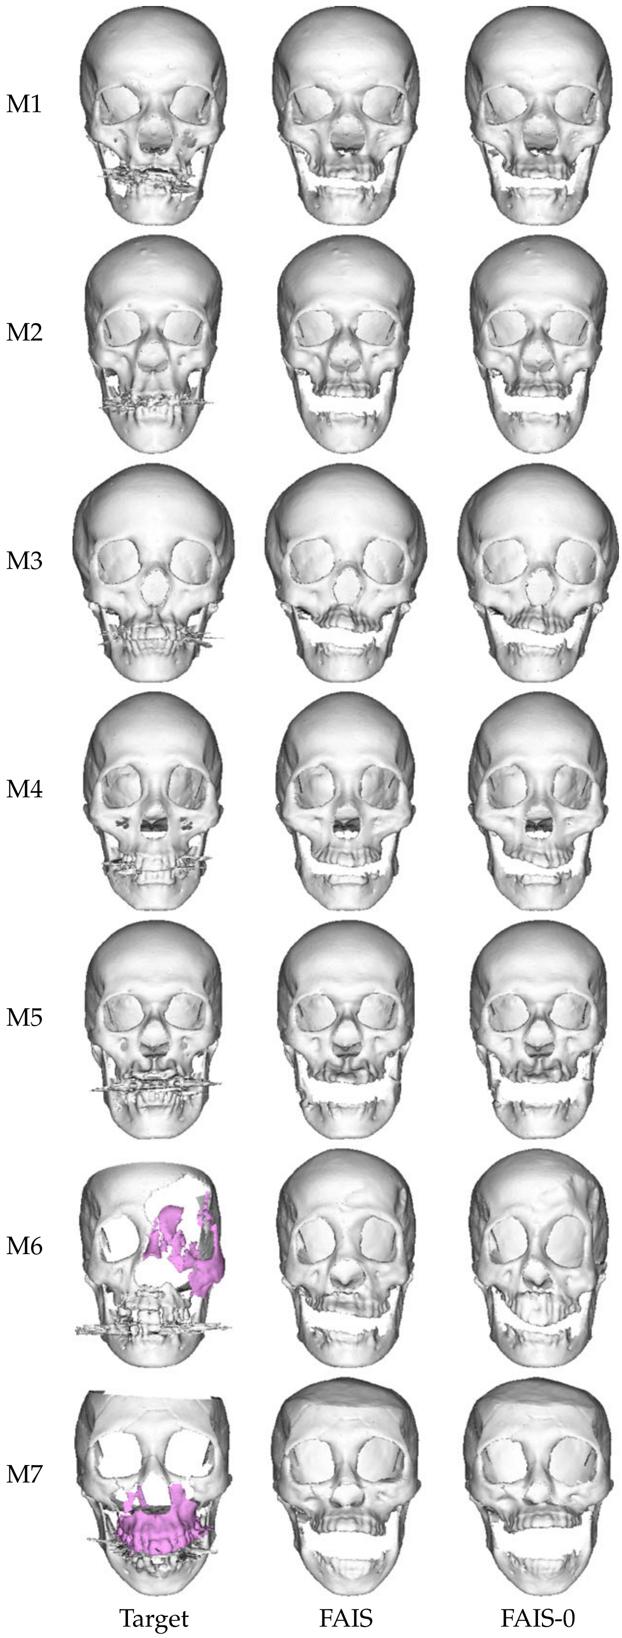

FIGURE 10 Robustness against metal artifacts. M1–M5, Non-defective skulls with metal artifacts. M6–M7, Defective skulls with metal artifacts. Defective parts are shown in pink

Therefore, in practice, it is advisable for the surgeon to visually choose a reference that is as similar to the target as possible in order to achieve accurate reconstruction.

TABLE 3 Effect of reference selection on FAIS's reconstruction errors E_D (in mm) of defective parts

	D1	D2	D3	D4	D5	Mean
Ideal	1.28	1.35	1.14	1.11	1.01	1.18
Exhaustive	1.50	1.53	1.14	1.14	1.35	1.33
Best prototype	1.62	1.53	1.73	1.14	1.35	1.47
ASM-F's reconstruction	1.84	1.20	1.34	1.80	1.43	1.52

8 | CONCLUSION

This paper has presented a novel method called FAIS for skull reconstruction. It exploits the strength of a non-rigid registration algorithm with interpolating surface while overcoming its weakness using a flip-avoiding technique. With an interpolating surface that regards positional constraints as hard constraints, FAIS can register a reference model to the non-defective parts of a target model exactly. It uses an iterative incremental registration method to obtain a very dense set of corresponding points for non-rigid registration. The flip-avoidance technique removes only about 20% of the mesh vertices for dense correspondence. FAIS's correspondence set is 2 orders of magnitude larger than those of existing methods such as TPS and ASM. As a result, FAIS can achieve practically zero error for the reconstruction of non-defective parts and smaller errors for the reconstruction of defective parts than existing methods. By incorporating symmetry constraints, it can produce more symmetric and normal results than other methods in reconstructing defective skulls with large numbers of defects. It is robust against severe outliers such as radiation artifacts in CT due to dental implants. As for existing methods, FAIS's reconstruction accuracy depends on how similar the reference model is to the target model. Proper reference selection can improve reconstruction accuracy in practice.

ACKNOWLEDGEMENTS

This research is supported by MOE grant MOE2014-T2-1-062

ORCID

Shudong Xie http://orcid.org/0000-0002-4764-7667

REFERENCES

1. Cevidanes LHC, Tucker S, Styner M, et al. 3D surgical simulation. *Am J Orthod Dentofacial Orthop*. 2010;138(3):361-371.
2. De Momi E, Chapuis J, Pappas I, et al. Automatic extraction of the mid-facial plane for cranio-maxillofacial surgery planning. *Int J Oral Maxillofac Surg*. 2006;35(7):636-642.
3. Wei L, Yu W, Li M, Li X. Skull assembly and completion using template-based surface matching. In: International Conference on 3D Imaging, Modeling, Processing, Visualization and Transmission; 2011; Hangzhou, China. 413-420.
4. Lüthi M, Albrecht T, Vetter T. Building shape models from lousy data. *Med Image Comput Comput Assist Interv*. 2009;12(Pt 2):1-8.
5. Zachow S, Lamecker H, Elsholtz B, Stiller M. Reconstruction of mandibular dysplasia using a statistical 3D shape model. In: Proc. Computer Assisted Radiology and Surgery (CARS); 2005; Berlin, Germany. 1238-1243.

6. Zhang K, Leow WK, Cheng Y. Performance analysis of active shape reconstruction of fractured, incomplete skulls. In: Proc. Int. Conf. on Computer Analysis of Images and Patterns. Malta: Springer; 2015:312-324.
7. Wei L, Yu W, Li M, Li X. Skull assembly and completion using template-based surface matching. In: 2011 International Conference on 3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT) IEEE; 2011; Hangzhou, China. 413-420.
8. Yin Z, Wei L, Li X, Manhein M. An automatic assembly and completion framework for fragmented skulls. In: 2011 IEEE International Conference on Computer Vision (ICCV). IEEE: Barcelona, Spain; 2011:2532-2539.
9. Yu W, Li M, Li X. Fragmented skull modeling using heat kernels. *Graph Models*. 2012;74(4):140-151.
10. Cheng Y. Computer-aided craniomaxillofacial surgery planning for fractured skulls. *Ph.D. Thesis*: National University of Singapore; 2013.
11. Deng Q, Zhou M, Shui W, Wu Z, Ji Y, Bai R. A novel skull registration based on global and local deformations for craniofacial reconstruction. *Forensic Sci Int*. 2011;208:95-102.
12. Lapeer RJA, Prager RW. 3D shape recovery of a newborn skull using thin-plate splines. *Comput Med Imaging Graph*. 2000;24(3):193-204.
13. Rosas A, Bastir M. Thin-plate spline analysis of allometry and sexual dimorphism in the human craniofacial complex. *Am J Phys Anthropol*. 2002;117:236-245.
14. Zhang K, Cheng Y, Leow WK. Dense correspondence of skull models by automatic detection of anatomical landmarks. In: *CAIP 2013, Part I. LNCS, Vol. 8047*, Wilson R, Hancock E, Bors A, Smith W, eds. Heidelberg: Springer; 2013:229-236.
15. Terzopoulos D, Witkin A, Kass M. Constraints on deformable models: recovering 3D shape and nonrigid motion. *Artif Intell*. 1988;36:91-123.
16. Cohen LD, Cohen I. Finite-element methods for active contour models and balloons for 2-D and 3-D images. *IEEE Trans PAMI*. 1993;15(11):1131-1147.
17. Bonarrigo F, Signoroni A, Botsch M. Deformable registration using patch-wise shape matching. *Graph Models*. 2014;76(5):554-565.
18. Amberg B, Romdhani S, Vetter T. Optimal step nonrigid ICP algorithms for surface registration. In: IEEE Conference on Computer Vision and Pattern Recognitio, 2007. Minneapolis, Minnesota, USA: CVPR'07. 2007:1-8.
19. Hontani H, Matsuno T, Sawada Y. Robust nonrigid ICP using outlier-sparsity regularization. In: Proceedings CVPR; 2012; Providence, Rhode Island. 174-181.
20. Bookstein FL. Principal warps: thin-plate splines and the decomposition of deformations. *IEEE Trans. PAMI*. 1989;11:567-585.
21. Sorkine O, Cohen-Or D, Lipman Y, Alexa M, Rössl C, Seidel H-P. Laplacian surface editing. In: Proc. Eurographics/ACM Siggraph Symp. Geometry Processing. 2004; Nice, France. 175-184.
22. Turner WD, Brown REB, Kelliher TP, Tu PH, Taister MA, Miller KWP. A novel method of automated skull registration for forensic facial approximation. *Forensic Sci Int*. 2005;154:149-158.
23. Marreiros FMM, Heuzé Y, Verius M, Unterhofer C, Freysinger W, Recheis W. Custom implant design for large cranial defects. *Int J Comput Assist Radiol Surg*. 2016;11(12):2217-2230.
24. Xie S, Leow WK. Flip-avoiding interpolating surface registration for skull reconstruction. In: Proc. icpr; 2016.
25. Xie S, Leow WK, Lim TC. Laplacian deformation with symmetry constraints for reconstruction of defective skulls. In: Proc. CAIP; 2017; Ystad, Sweden. 24-35.
26. McInemey T, Terzopoulos D. Topology adaptive deformable surfaces for medical image volume segmentation. *IEEE Trans Med Imaging*. 1999;18(10):840-850.
27. Lachaud J-O, Montanvert A. Deformable meshes with automated topology changes for coarse-to-fine three-dimensional surface extraction. *Med Image Anal*. 1998;3(2):187-207.
28. Delingette H, Montagnat J. Shape and topology constraints on parametric active contours. *Comput Vis Image Underst*. 2001;83(2): 140-171.
29. Jung W, Shin H, Choi BK. Self-intersection removal in triangular mesh offsetting. *Comput-Aided Des Applic*. 2004;1:477-484.
30. Choi Y, Lee S. Injectivity conditions of 2D and 3D uniform cubic B-spline functions. *Graph Models*. 2000;62(6):411-427.
31. Hagenlocker M, Fujimura K. CFFD: a tool for designing flexible shapes. *Vis Comput*. 1998;14:271-287.
32. Khan A, Aylward E, Barta P, Miller M, Beg MF. Semi-automated basal ganglia segmentation using large deformation diffeomorphic metric mapping. In: Proc.MICCAI; 2005; Berlin, Heidelberg. 238-245.
33. Zhuang X, Rhode K, Arridge S, et al. An atlas-based segmentation propagation framework using locally affine registration—application to automatic whole heart segmentation. In: Proc. MICCAI; 2008; New York, NY, USA. 425-433.
34. Ding F, Yang W, Leow WK, Venkatesh SK. 3D segmentation of soft organs by flipping-free mesh deformation. In: Proc. WACV; 2009; Snowbird, Utah, USA. 1-7
35. Masuda H, Yoshioka Y, Furukawa Y. Interactive mesh deformation using equality-constrained least squares. *Comput Graph*. 2006;30(6):936-946.
36. Golub GH, Van Loan CF. *Matrix computations*, 3rd ed. Johns Hopkins; 1996.
37. Mouton A, Megherbi N, Van Slambrouck K, Nuyts J, Breckon TP. An experimental survey of metal artefact reduction in computed tomography. *J X Ray Sci Technol*. 2013;21(2):193-226.

How to cite this article: Xie S, Leow WK, Lee H, Lim TC. Flip-avoiding interpolating surface registration for skull reconstruction. *Int J Med Robotics Comput Assist Surg*. 2018;14:e1906. https://doi.org/10.1002/rcs.1906