

Computer Assisted 3D Craniofacial Reconstruction

by

David William Bullock

B.Sc. (Hon.), Simon Fraser University, 1996

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF


Master of Science

in

THE FACULTY OF GRADUATE STUDIES

(Department of Computer Science)

We accept this thesis as conforming
to the required standard



The University of British Columbia

October 1999

© David William Bullock, 1999

Abstract

This thesis describes a methodology for forensic craniofacial reconstruction, defined as the reproduction a subject's facial appearance from the skull and mandible for the purposes of identification. The focus of this work is on reproducing three dimensional facial shape from the shape of the skull and mandible of a subject. This task can be broken into two parts. First, a technique for determining soft tissue depth at any point on the skull and mandible surface is presented. Second, two complementary methods of producing three dimensional facial shape are presented: a method which generates a surface using the skull and mandible as a skeleton of an implicit surface and a method which produces a hierarchical B-spline surface fit to facial surface locations. Surface material properties and facial features can then be added and modified to these models before being rendered and distributed. Results of these processes are presented.

Contents

Abstract	ii
Contents	iii
List of Tables	v
List of Figures	vi
Acknowledgements	viii
1 Introduction	1
1.1 Previous Work	2
1.2 The Reconstruction Pipeline	4
2 Interpolating Tissue Depths	7
2.1 Formalizing the Reconstruction Process	7
2.2 Interpolating Tissue Depth	8
3 Reconstructing Facial Models with Implicit Surfaces	14
3.1 Defining an Implicit Surface	15
3.2 Polygonizing Isosurface	17

3.2.1	Marching Cubes	19
3.2.2	Bounding Box Optimization	20
3.2.3	Min-Max Optimization	22
3.3	Problems and Solutions	24
3.4	Results	28
4	Reconstructing Facial Models with Hierarchical B-Splines	31
4.1	Overview of Hierarchical B-Spline Surfaces	32
4.2	Fitting Hierarchical B-Spline Surface	35
4.3	Adding Facial Features	38
4.4	Results	39
5	Conclusion	40
5.1	Results	40
5.2	Summary	42
5.3	Future Work	43
	Bibliography	46
	Appendix A Tissue Thickness Measurements	49

List of Tables

A.1	Tissue thickness measurements (in mm), for American Caucasoids	50
A.2	Tissue thickness measurements (in mm), for American Blacks	51
A.3	Tissue thickness measurements (in mm), for Southwestern Indians	52

List of Figures

1.1	Standard skull and mandible landmarks	3
1.2	Digitized skull	5
2.1	Standard dowel network	9
2.2	A cutting plane between two dowels	10
2.3	Barycentric interpolation technique on two dowel areas	12
2.4	Interpolation method applied to skull	13
3.1	Single triangle implicit surface	15
3.2	Multiple triangle implicit surface	16
3.3	Marching Cubes cases	19
3.4	Emitter cache	24
3.5	Original and filled skull for varying t	26
3.6	Polygon reduced implicit surfaces	29
3.7	The final implicit surface	30
4.1	Local refinement	33
4.2	Generic face hierarchical B-spline surface	35
4.3	H-spline surface fit to standard dowel locations	36

4.4	H-spline surface fit to interpolated dowels	39
5.1	Results of the implicit surface and hierarchical B-spline fitting techniques	41
5.2	Implicit surface method results for different weight types	41
5.3	Hierarchical B-spline fitting method results for different weight types	42

Acknowledgements

I would like to thank the following people for their help and support while completing my thesis:

- Jack Snoeyink for accepting me as a student so late in my degree, his help in developing the material, and quick feedback on my drafts;
- David Forsey for all his help and guidance and meeting with me so often after moving to industry;
- Alain Fournier for his feedback on my thesis, and mentorship throughout my time at UBC;
- Bill Gates, Robert Scharein, Lisa Streit, Roger Tam, and Robert Walker for their reviews of my thesis drafts, and thanks to them again plus Paul Fearing and Gene Lee for offers of couches to sleep on; the moral support from Imager lab members has been amazing;
- Donald and Florence Bullock, my parents, not only for their support through this degee, but for instilling a love for learning throughout my whole life; and
- Tracey Chew. At the best of times I could not possibly convey what your love and support has meant to me over the last two years, let alone after being awake for 50 hours. Thank you for all you've given me.

DAVID WILLIAM BULLOCK

The University of British Columbia

October 1999

Chapter 1

Introduction

This thesis describes a methodology for forensic craniofacial reconstruction, the reproduction of human facial appearance from the skull and mandible. We will focus on reproducing 3D facial shape from the 3D shape of the skull.

Our goal is to perform a task that is commonly executed in forensic craniofacial reconstruction: accurately reconstruct the look and shape of a human head from a skull and mandible. We do not wish to replace the forensic artist; the long-term goal of this project is to provide the practical foundation for tools which forensic artists could use to produce facial reconstructions. Section 1.1 will provide background information and previous work on this subject. Section 1.2 will describe the reconstruction process, both in the traditional and computational setting, to provide a context for this work. This will be followed by Chapter 2 which outlines a method for interpolating tissue depth given any location on the skull or mandible. This fundamental concept is used throughout the remainder of the thesis. Chapters 3 and 4 each describe a different method of reconstructing facial shape using implicit surface and hierarchical B-spline techniques, respectively. Final conclusions

and future work are provided in Chapter 5.

1.1 Previous Work

Craniofacial reconstruction has been used for over a century for identification of remains and to reconstruct busts of historical figures, for archaeological and anthropological applications and for identification of human remains in the forensic sciences. It is the latter application that this work focuses on. Forensic craniofacial reconstructions are generated from the remains of a subject in an attempt to acquire a likeness the subject when alive. The reconstruction is then used to attain leads which, in turn, will hopefully result in the positive identification of the subject in conjunction with other techniques such as dental records [NM98]. For the remainder of this thesis craniofacial reconstruction and facial reconstruction are used to refer to “the scientific art of visualizing faces on skulls for the purpose of individual identification” [Geo87].

There are four general techniques used for facial surface reconstruction: (i) two dimensional artistic techniques, (ii) photographic or video super-imposition, (iii) three dimensional facial reconstruction, and (iv) volumetric deformation techniques [NM98]. Techniques (i), (ii), and (iii) rely on associating skull and mandible landmark locations¹ with facial soft tissue locations. These landmark tissue depths are acquired experimentally based on race, age, gender, and body type using a variety of techniques [QCS⁺97, KI86, Geo93]. Refer to Figure 1.1 and table A.1 for the standard landmark locations and tables of tissue depths used throughout this thesis. Technique (iv) uses a variety of volumetric deformation techniques to transform

¹These are the standard landmarks from [Rhi84, KI86]. Landmark definitions and locations differ slightly from study to study.

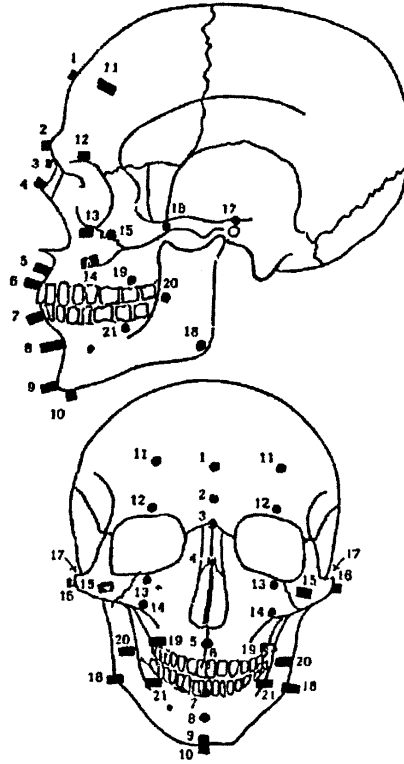


Figure 1.1: Standard skull and mandible landmarks (from [Rhi84]).

an “average” human face to the facial shape of the subject based on the volumetric transform of the “average” skull to the subject’s skull [NM98, QCS⁺97].

Techniques (iii) and (iv) are both intended to produce a fully three dimensional representation of the subject’s pre-death facial soft tissue. Reconstruction methods explored in this thesis will be based on the standardized Krogman method [KI86], which employs technique (iii). For a complete discussion of all these techniques see [TEC⁺97, HRPM93, NM98, KI86, Ste79].

Computer-aided facial reconstruction has been growing in popularity over the last decade in both two dimensional and three dimensional domains. Volumetric deformation methods of technique (iv) are inherently computational. Previous work

on computer-aided variants of technique (iii) can be found in [Arc97, VBL⁺89].

It is hoped that tools used to produce reconstructions will eventually allow forensic artists to produce reconstructions in less time and with greater flexibility in presentation than current methods. As well, we wish to provide an easier means of producing alternative reconstructions when factors such as subject race, weight, and facial features are not known.

1.2 The Reconstruction Pipeline

This work focuses on the reproduction of facial shape. The reconstruction process begins with a subject's skull and mandible from which we wish to ascertain the subject's pre-death appearance. Forensic scientists use clues from the scene, skull shape, and local demographic information to deduce the subject's gender, race, and approximate weight and age.

We will first describe the standardized Krogman method used to produce a facial reconstruction on a physical skull and mandible followed by a computational process to accomplish the same task. In the Krogman method, dowels are cut and placed on the standard skull and mandible landmark positions. The length of each dowel is determined by charts of soft tissue thickness indexed by the subject's gender, race, and weight. Next, strips of sculpting material such as clay are placed between the dowels to interpolate the soft tissue thickness between adjacent landmarks. Strips are further placed in a systematic fashion between landmarks or strips until there are only small gaps in the surface. These gaps are then filled in and the whole surface is smoothed to produce the final facial surface. Additional facial features such as hair are added to complete the model.

To produce a computational reconstruction, the skull and mandible are digi-



Figure 1.2: Digitized skull.

tized by laser scanning, computed tomology (CT) scans, magnetic resonance imaging (MRI) scans, or other 3D digitization techniques. This produces a full three dimensional polygonal model of the skull and mandible which is used as the basis of the reconstruction. Figure 1.2 shows the skull used for reconstruction throughout this thesis. This is the skull of an average weight male caucasoid.

The next step is to produce a representation of the facial soft tissue. First, virtual dowels are placed on the polygonal skull model at standard landmarks. Again, the length of each dowel is dictated by experimentally determined charts of soft tissue depths for the standard landmarks. Then, a pre-death facial model is produced in the form of a polygonal model using the processes such as those discussed in Chapters 2, 3, and 4. The interpolation between dowel locations will be

more consistent and less prone to human error with this computational model.

After the facial shape is obtained, the facial model can be exported to a commercial 3D animation package where material properties can be applied to the skin. Hair and ears can be added to produce a more realistic approximation of the subjects appearance. Animations of the reconstruction could then be used and distributed for the purpose of identification. As well, there will be a practical means of distributing models in a 3D form in the future as conventions such as VRML mature and become more widely available.

Alternatively, static images of the reconstructed face can be produced. These images could then be coloured and enhanced by a forensic artist using professional digital paint programs.

Chapter 2

Interpolating Tissue Depths

This chapter addresses the problem of determining tissue depth in regions between dowels. The first section formalizes terms needed in the reconstruction process; the next section describes how to calculate tissue depth of an arbitrary point.

2.1 Formalizing the Reconstruction Process

We will now provide a formalization of the techniques used for craniofacial reconstruction as described in Section 1.2. Three dimensional craniofacial reconstruction starts with the placement of dowels at standard facial landmark locations (Figure 1.1) of set length (Tables A.1, A.2, and A.3). The computational reconstruction system simulates the process of physically placing dowels on the skull and mandible by allowing the user to attach virtual dowels to the skull model. By default, dowel direction is normal to the surface it is placed on but can be adjusted.

Definition 2.1 *A dowel d_i is defined as a tuple $(\vec{p}_i, \vec{n}_i, \ell_i)$ where \vec{p}_i is the base, \vec{n}_i is the direction in the form of a unit vector, and ℓ_i is the length of d_i .*

The endpoint of a dowel can be calculated as $\vec{p}_i + \vec{n}_i \ell_i$.

In the standardized Krogman method, the forensic artist fills in the facial skin once these dowels have been placed. The forensic artist runs modelling material such as clay from dowel to dowel in a standard network and then fills in the areas of this network to produce a smooth surface. These areas are filled in using standard rules for interpolation and the forensic artist's knowledge of facial form. The system of strips is used to reduce the subjectivity of the interpolation. Since tissue depths can be attained computationally, and therefore more objectively, a global interpolation method based on this network of dowels is provided. This network is formalized as follows:

Definition 2.2 *A **dowel area** is an area of influence defined for three dowels. Tissue depths in the interior of a given **dowel area** are an interpolation of the three dowel tissue depths.*

Definition 2.3 *A **dowel network** is a collection of dowel areas incident on a set of dowels.*

The dowel network is standardized and shown in Figure 2.1.

2.2 Interpolating Tissue Depth

Given a point \vec{p} on the skull model, the tissue depth ℓ associated with \vec{p} must be determined. The first step is to determine which dowel area \vec{p} belongs to. That dowel area is then used to interpolate tissue depth at \vec{p} . The soft tissue surface can be thought of as sitting at $\vec{p} + \ell \vec{n}$, where \vec{n} is the unit vector normal at \vec{p} . The normal \vec{n} can be calculated either by using the skull model surface normal at \vec{p} or by interpolating the dowel directions as well.




Figure 2.1: Standard dowel network (from [Rhi84], modified).

We will start with a few definitions before defining which dowel area a point belongs to.

Definition 2.4 *A cutting plane is defined as a plane that divides adjacent dowel areas (i.e., dowel areas that share a common edge). If d_i and d_j are the dowels shared between two adjacent dowel areas, the **cutting plane** between these dowel areas is defined to be the plane that is parallel to $n_i + n_j$ and contains points p_i and p_j (p_i and n_i are defined on d_i as described in Section 2.1).*

Figure 2.2 provides a visual example of a cutting plane between two dowels.

The distance of p to a dowel area a is defined as the *shortest distance* between




Figure 2.2: A cutting plane between two dowels.

p and the plane prescribed by the base points of the three dowels for a . We say p is inside a cutting plane of a if p lies in the same half-space of the cutting plane as the points defining a .

We say p belongs to a dowel area a when the following conditions are met:

1. p lies inside each of the three cutting places of a ; and
2. the distance of p to a is less than the distance of p to any other dowel area that meets condition 1.

Once we determine that p belongs to some dowel area a , we calculate the tissue depth as an interpolation of the three dowels of a , d_1 , d_2 , and d_3 . The interpolation is conducted on the plane prescribed by p_1 , p_2 , and p_3 (the bases of d_1 , d_2 , and d_3 , respectively); we will call this the interpolation plane. The projection of point p is calculated first. Then, the tissue depth of p is calculated using barycentric interpolation in the triangle $\Delta(p_1, p_2, p_3)$ as described below.

Dowel area a contains three cutting planes. Cutting plane $c_{1,2}$ is the cutting plane defined on dowels d_1 and d_2 ; cutting planes $c_{2,3}$ and $c_{3,1}$ are similarly defined.

We can formulate the cutting planes using the standard plane equations as follows:

$$A_{1,2}x + B_{1,2}y + C_{1,2}z + D_{1,2} = 0$$

$$A_{2,3}x + B_{2,3}y + C_{2,3}z + D_{2,3} = 0$$

$$A_{3,1}x + B_{3,1}y + C_{3,1}z + D_{3,1} = 0$$

or in matrix form:

$$M \begin{bmatrix} x \\ y \\ z \end{bmatrix} + \begin{bmatrix} D_{1,2} \\ D_{2,3} \\ D_{3,1} \end{bmatrix} = 0$$

$$\text{where } M = \begin{bmatrix} A_{1,2} & B_{1,2} & C_{1,2} \\ A_{2,3} & B_{2,3} & C_{2,3} \\ A_{3,1} & B_{3,1} & C_{3,1} \end{bmatrix}$$

If $\det|M| = 0$ then all three dowel normals are parallel¹. In this case, \vec{p}' , the projection of \vec{p} on the interpolation plane, will be the intersection of the line $\overline{\vec{p}, \vec{p} + \vec{n}_1}$ and the interpolation triangle.

Otherwise, the cutting planes will intersect at a point \vec{p}_c since $\det|M| \neq 0$.

$$\vec{p}_c = M^{-1} \begin{bmatrix} -D_{1,2} \\ -D_{2,3} \\ -D_{3,1} \end{bmatrix}$$

¹Note that if all dowel normals are \vec{n} , then the three plane normals are $(2\vec{n}) \times (\vec{p}_2 - \vec{p}_1)$, $(2\vec{n}) \times (\vec{p}_3 - \vec{p}_2)$, and $(2\vec{n}) \times (\vec{p}_1 - \vec{p}_3) \Rightarrow$ the plane normals are coplanar $\Rightarrow \det|M| = 0$.




Figure 2.3: Barycentric interpolation technique on two dowel areas. Interpolated tissue depths of the four dowels (represented as narrow cones) are added to the location of the points on the plane to produce a new surface. The regions of influence of the two dowel areas is evident.

This \vec{p}_c will be used as a center of projection. The intersection of the line $\overline{\vec{p}, \vec{p}_c}$ and $\Delta(p_1, p_2, p_3)$ will be the interpolation point \vec{p}' .

Once we have \vec{p}' , we will calculate tissue depth ℓ_i using barycentric interpolation.

$$\ell_i = u\ell_1 + v\ell_2 + w\ell_3$$

where

$$\begin{aligned} u &= \frac{Area(\vec{p}, \vec{p}_2, \vec{p}_3)}{Area(\vec{p}_1, \vec{p}_2, \vec{p}_3)} \\ v &= \frac{Area(\vec{p}_1, \vec{p}, \vec{p}_3)}{Area(\vec{p}_1, \vec{p}_2, \vec{p}_3)} \\ w &= \frac{Area(\vec{p}_1, \vec{p}_2, \vec{p})}{Area(\vec{p}_1, \vec{p}_2, \vec{p}_3)} \end{aligned}$$

One can see how this interpolation works on a plane in Figure 2.3.




Figure 2.4: Interpolation method applied to skull. Shown is (a) the original skull, (b) the skull with the tissue depths added to each point of the skull, and (c) the surface produced by extracting those faces from the skull that were effected.

Finally, we can apply this interpolation to the full skull model. Given a point \vec{p}_i with normal direction \vec{n}_i and interpolated tissue depth ℓ_i , we can move this point by a delta $\ell_i \vec{n}_i$. Figure 2.4 (b) shows us a skull model with these interpolated tissue depths added as a delta to each point. We compute a new point $p_i^{\vec{n}ew} = \vec{p}_i + \ell_i \vec{n}_i$ for each p_i . Separating this delta mesh from the original skull model clearly displays how these deltas affect the mesh; see 2.4 (c).

Now that we can interpolate tissue depth at any location of the skull, we must provide a smooth and accurate surface which interpolates these soft tissue locations. This is demonstrated in Chapters 3 and 4.

Chapter 3

Reconstructing Facial Models with Implicit Surfaces

In this chapter we describe how to reconstruct the facial model with the use of implicit surfaces. We are attempting to produce a facial model that sits a given distance (i.e., tissue depth) from the skull model. The facial model must not contain the noise and sharp corners of the polygonal skull model. Emission-based implicit surface modelling techniques seem well suited to solving this problem; a set of skeleton emitters generally produce a smooth surface, even for jagged skeletons.

In Section 3.1 we show how to define an implicit surface for the purposes of facial reconstruction. In Section 3.2 we discuss a standard method used to polygonize the implicit surface and extensions to it to speed the polygonization process. Next, Section 3.3 describes problems with the implicit surface technique and solutions to these problems. Finally, Section 3.4 shows results using this technique.




Figure 3.1: A triangle and the implicit surface produced from it.

3.1 Defining an Implicit Surface

First we define an implicit surface. Given a function $f(\vec{p})$ defined on 3-dimensional Euclidean space, an implicit surface is the isosurface associated with a particular real value of f . If we want the isosurface at value v , we are interested in the surface produced by the union of the points where $f(\vec{p}) = v$.

The skull model is used as the implicit surface skeleton. Each polygon i will emit a tissue depth ℓ_i calculated as the average tissue depth of the polygon points. Tissue depth of the polygon points is calculated as described in Chapter 2.

The field of an emitter¹ is defined for polygon i as

$$e_i(\vec{p}) = \begin{cases} 0 & \text{if } \vec{p} \text{ is behind polygon } i \\ (\frac{\ell_i}{d})^k & \text{otherwise} \end{cases} \quad (3.1)$$

where d is the shortest distance between \vec{p} and polygon i , and k is a constant.

The isosurface of this single emitter is calculated at $e_i(p) = 1$ with $e_i(p) > 1$ inside and $e_i(p) < 1$ outside the object. The interior of the object produced by the

¹The terms emitter and skeleton are used interchangeably in implicit surface literature. Though the term skeleton is more commonly used, emitter produces less confusion given the context of this thesis.

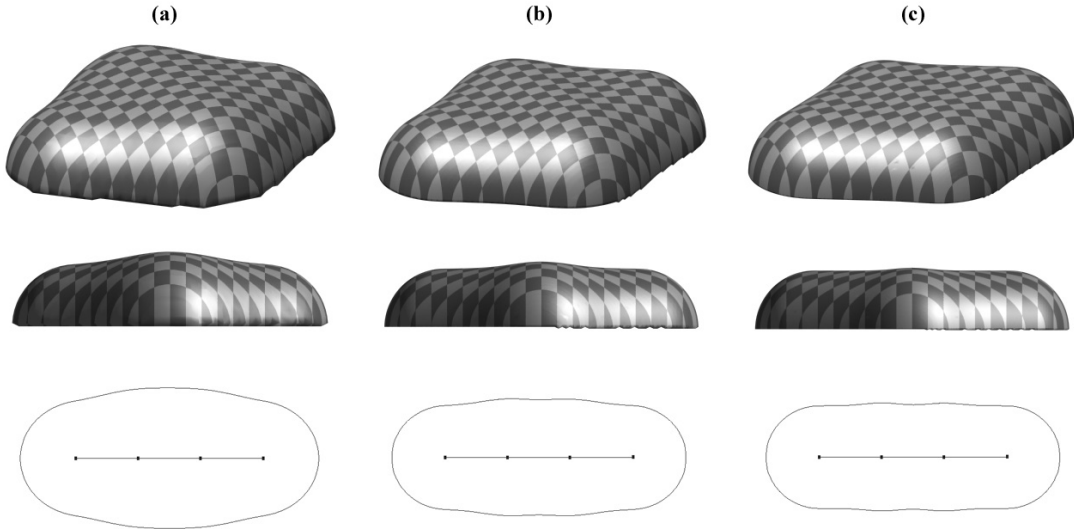


Figure 3.2: Multiple triangle implicit surface. The top two rows depict the implicit surface produced by two adjacent triangles. The bottom row depicts the implicit surface cross section of three adjacent triangles. For column (a) $k = 3.0$, (b) $k = 6.0$, and (c) $k = 15.0$.

polyhedron is the volume union of the polyhedron extruded by ℓ_i in its facing direction, cylinders along the edges, and balls on the points of radius ℓ_i as seen in Figure 3.1. The volume is halved along the plane of the emitter polyhedron.

The sum of the emitters at a point \vec{p} produces the total field function e as

$$e(\vec{p}) = \sum_i (e_i(\vec{p})) - 1$$

with the isosurface taken at $e_i(\vec{p}) = 0$; a threshold value of 0 is used to follow convention. This summing of emitter values produces a blending between the individual isosurfaces. The constant k encapsulates the idea of locality of emitter influence. The higher the value of k , the faster $e_i(\vec{p})$ decreases outside distance ℓ_i , and therefore the less $e_i(\vec{p})$ can contribute to the field past a distance ℓ_i . As k increases, the resulting isosurface comes closer to the union of the isosurfaces of the individual

emitters. Figure 3.2 shows an example of two emitters with different values of k . Note that as k is increases from 3.0 in 3.2(a) to 15.0 in 3.2(c) that the center bulge where the two emitters meet is reduced.

3.2 Polygonizing Isosurface

For the purpose of forensic reconstruction we need to convert this implicit function to an editable and renderable form. Given an implicit function, a polygonal surface can be produced at the desired isosurface level; the subsystem that accomplishes this task will be referred to as the polygonizer. There are a number of techniques that can be used to accomplish this including Marching Cubes (MC) [LC87] (and variants), and Shrinkwrap [vOW93].

While these techniques work well there are some problems for this particular application. The Shrinkwrap technique works on only one surface at a time. Although we only wish to produce one full facial model, incomplete and segmented models are sometimes produced during the reconstruction process due to insufficient sample rates, small tissue depths, and interior disjoint elements that the skull and mandible models can contain. As well, due to the nature of the implicit surfaces, an inside and outside model is produced; one of the models is the desired skin model while the other is produced by the back-faces of the emitters. Changes in topology also add complexity to the Shrinkwrap algorithm.

There are problems with using traditional MC as well. The ratio of tissue depth to maximal skull extent can reach levels of two-hundred to one. A sampling rate of at least this ratio is required in each extent of the marching volume in order to ensure no surface is skipped over. Thus, MC would have to check n^3 or 8 million cubes during the march. Even worse, all emitters need to be evaluated for each

corner of the cubes. This would be a prohibitive amount of work.

To reduce the amount of work required to produce the polygonized implicit surface, two extensions to the MC algorithm will be presented. Traditionally, the isosurface generator can only obtain information from the field emitters in the form of a function $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ that evaluates the implicit function given a point. This is a nice separation in terms of design; it provides a simple interface between the implicit surface evaluator and the polygonizer. In addition, isosurface polygonization algorithms that use such an interface also work on volumetric data. However, the work required for a march can be reduced by allowing the isosurface generator to query the field individual emitters.

An additional interface is provided between the implicit function evaluator and the polygonizer. The polygonizer evaluates each of the individual emitters and is responsible for summing them up. As well, the polygonizer can ask for additional information from each emitter as described in Sections 3.2.2 and 3.2.3.

An octree data structure [Sam89, pp. 1–9] is used to keep track of the spatial influence of each emitter and to query the value of the implicit function.

Definition 3.1 *An **octree** is a recursive spatial tree structure defined on a rectilinear bounding volume v . Each node represents a volume; the root of the tree represents all of v . Each internal node has eight children which contain each of the eight equal size octants of the parent node's volume.*

The MC algorithm and optimizations to MC are presented in the following sections.




Figure 3.3: Marching Cubes cases. Shown are the 15 cases for cubes in the Marching Cubes algorithm; all possible combinations of inside/outside corners are rotation and inverse equivalent to these cases. Black corners represent points inside the object, white corners outside the object.

3.2.1 Marching Cubes

The standard MC algorithm will be presented in this section. Before the “march”² is conducted, a rectilinear volume is calculated based on the maximal extents of the surface we wish to produce. This volume is partitioned into n rectilinear boxes in each dimension to produce a three dimensional matrix of n^3 boxes.

To polygonize the surface, each box is sequentially evaluated as follows. Each of the corners of the box is evaluated using the emission function to determine whether the corner is inside or outside the object. There exist $2^8 = 256$ possible combinations of inside/outside corners on the 8 corners a box. These 256 variations

²The term “march” will be used to refer to the process of polygonizing the surface using the MC algorithm or one of its variants.

are rotation and inverse equivalent to the 15 cases shown in Figure 3.3. A box contributes polygons to the isosurface mesh based on which case the box is equivalent to; no polygons are added if all points are inside or outside the object (i.e., case 0). The actual location of the point on an edge of the box can be calculated by linear interpolation of the corner values or by binary search of the cross-over point; the latter technique is used for this application.

In order to avoid recalculating the values at the box corners and edge point locations of adjacent boxes, an n^2 level of boxes is kept in memory during the march. The march is conducted one level at a time. While marching a level, point values and edge points can be stored and reused. When the level is completed, the corner values and edge points adjacent to the next n^2 level of boxes are retained.

3.2.2 Bounding Box Optimization

In this optimization the polygonizer is allowed to query each emitter for a bounding box of influence. This bounding box will be considered to be the maximal volume of influence of the emitter. Emitters are assumed to provide no field strength outside this volume.

While emitter field strength never reaches zero, we can choose an arbitrarily small maximal contribution an emitter can make by increasing the size of the bounding box surrounding it. A distance d is chosen such that it would take m emitters to produce a surface at distance d . By the definition of emitter strength, $(\frac{\ell}{d})^k m = 1 \Rightarrow d = \ell m^{\frac{1}{k}}$, where ℓ and k are defined for the emitter; we set $m = 500$ for the purposes of this application. This d is then used to calculate the emitter bounding box which in turn is used in the octree data structure.

Definition 3.2 *An MCOctree is an octree whose nodes contain a list of emitters*

incident on the node volume (i.e., whose bounding boxes intersect the node volume).

An MCOctree can be constructed from a given list of emitters and a bounding volume. The initial bounding volume is computed from the extents of the emitter bounding boxes. A node in the octree is only expanded during octree construction if the node's list of emitters is greater than 1 and the depth of the node is less than $\lceil \log_2 n \rceil$, where n is the sampling rate of the march.

This rule for expansion works for depths of up to eight; tree memory usage beyond this point is prohibitive, even if intersecting emitter lists are only kept at the leaves and bounding volumes implicitly stated. Since MC only requires one n^2 level of cubes be in memory storage at any time during the march, only the nodes that intersect the current level of the march will be expanded in order to preserve memory. As we move from one n^2 level to the next, the octree is updated with nodes outside the level deleted and nodes inside expanded.

Once the octree is constructed, a point can be evaluated by traversing the tree until the leaf that contains the point is found. The emitters for the leaf are evaluated and summed on the point and used as the field strength. If there are no emitters in the leaf, then 0 is returned.

If we could attain bounding boxes of the interior of the emitters the MC-Octree could be constructed so that any node volume that is entirely covered by the interior of emitters would not be expanded. No surface would be produced in this volume since the emitter value would always be larger than the threshold inside the inner emitter bounding box. This further improvement to the algorithm would produce no gain for this application as the emitter size is small relative to the smallest node volume.

3.2.3 Min-Max Optimization

The bounding boxes of emitters are quick and easy to compute; however, this representation is overly conservative. This will be illustrated in an example. Suppose we have two emitters whose bounding boxes intersect a volume but whose maximal value in the volume is $\frac{1}{3}$. Even though a surface cannot be produced within this volume, the volume would not be pruned from the MCOctree.

Instead the maximal and minimal values that an emitter can contribute for a given volume are evaluated. Take emitter i with emission function e_i . Given a rectangular bounding volume v , these minimum and maximum values can be calculated with the functions $\max_i(v)$ and $\min_i(v)$ respectively as follows:

$$\begin{aligned} \max_i(v) &= \begin{cases} \epsilon_{max} & \text{if } v \text{ intersects emitter polygon} \\ e_i(\vec{p}_{closest}) & \text{otherwise} \end{cases} \\ \min_i(v) &= \min\{e_i(\vec{p}) : \vec{p} \text{ is corner of bounding box } v\} \end{aligned}$$

where $\vec{p}_{closest}$ is the closest point in v to emitter polygon for emitter i .

The maximal value will be capped at some $\epsilon_{max} \gg 1$. Any value smaller than some ϵ_{min} will be considered to be equivalent to zero and therefore not to affect the volume. We can now define a new octree with a new expansion rule.

Definition 3.3 *A MinMax MCOctree is an octree whose nodes contain a list of emitters whose maximal contribution is at least ϵ_{min} for the node.*

During tree construction, a node with volume v and emitter list L is only expanded if $\sum_{i \in L} \max_i(v) > 1$ and $\sum_{i \in L} \min_i(v) < 1$. No surface could possibly be produced if this is not the case as the crossover value does not exist in v . Just

as with the MCOctree, the MinMax MCOctree need only be expanded for a single marching level at a time in order to conserve memory.

An additional optimization is required in order to achieve a speedup with the min-max method. Since the emitter's minimal and maximal values must be computed at the corners of the octree bounding boxes and since these corners can be incident on up to eight nodes of the octree for up to $\lceil \log_2 n \rceil$ levels of the octree, we would like to cache these emitter values. We will align the octree structure so that its lowest level bounding volumes coincide with the marching cubes themselves. The top/front/left box of the lowest level of the octree will be aligned with the $(0,0,0)$ box of the marching matrix. There will be $\lceil \log_2 n \rceil$ levels to the march with the extra $2^{\lceil \log_2 n \rceil} - n$ octree boxes in each extent extruding past the end of the marching volume. The emitter value would then only need be computed once for any location. Keep in mind that we would still have to compute the emitter values along the edges of the marching cubes, but this is only done once for each edge during the march.

Emitter values are cached at the corners of the octree. An index (i, j, k) into the $(n + 1)^3$ list of corners of the octree and marching matrix is converted into an index (i, j, d) into the cache. Let (i, j) be the index into the current n^2 level of boxes during the march. The value of d is set to be the minimum depth of any of the corners of the k^{th} level $(n + 1)^2$ list of corners. This can be calculated directly from k such that $d = \lceil \log_2 n \rceil - i$ where i is the index of the lowest order bit set in k . Because of the order in which the cubes are marched, k will either be at the top or bottom of our current n^2 level of boxes. We need to maintain the emitter values for the $\lceil \log_2 n \rceil$ depths of the octree that surround the current n^2 level of marching boxes. These octree depths are reused when the octree is re-evaluated as we move




Figure 3.4: Emitter cache. On the left we see a sample of a marching volume for $n = 8$. On the right is the octree surrounding the third level of the march. (i, j, k) is the index into the array of corners and d is the minimum depth of a point of a level of corners. Note that levels for the octree are completely filled in assuming emitters are active on all boxes of the current level.

from one marching level to the next. These $\lceil \log_2 n \rceil$ octree depths will include the corners of the current marching level.

3.3 Problems and Solutions

This section discusses some of the problems associated with the implicit surface reconstruction method and solutions to these problems.

The implicit surface method cannot be used on the skull and mandible directly. The indents around the skull temples and the gap between the mandible and skull (Figure 3.5(a)) show up in the final implicit surface (Figure 3.5(c) and (e)). These gaps in the surface can be manually filled in (Figure 3.5(b)) to produce a more realistic resulting surface (Figure 3.5(d) and (f)).

Another problem associated with this form of reconstruction is that final facial shape is not independent of the number of emitters. Even if two emitter

skeletons form exactly the same shape, the resulting surface could be quite different. Bulging can occur if the emitter skeleton size is small compared to its field distance as shown in Fig. 3.2. As 3.2 shows, bulging can be reduced by increasing the value k in the emitter formula; however, even high values of k do not remove this bulging as this does not eliminate the source of the problem.

A new emitter is defined for polygon i based on a *split distance* t , $0 \leq t \leq 1$, as follows:

$$e_i(\vec{p}) = \begin{cases} \epsilon_{\max} & \text{if } 0 \leq d_i < t\ell_i \text{ and } \vec{p} \text{ projects on polygon } i \\ \left(\frac{(1-t)\ell_i}{d}\right)^k & \text{if } d_i \geq t\ell_i \\ 0 & \text{otherwise} \end{cases} \quad (3.2)$$

where d_i is the signed distance of \vec{p} from polygon i 's plane with positive d_i indicating distances in front of the plane and negative d_i giving distances behind the plane and d is the distance of \vec{p} to a polygon constructed by moving the points of the emitter polygon by $t\ell_i$ in the direction of the emitter polygon normal.

If we set $t = 0$, the emitter is exactly as defined in equation 3.1. However, as we increase the value of t , we limit the scope of the emitter. This new emitter can be visualized as follows. Take the volume of polygon i extruded along its normal a distance $t\ell_i$. If \vec{p} is in this volume, we return ϵ_{\max} ; if \vec{p} is past distance $t\ell_i$ in the facing polygon direction, we evaluate the emitter exactly as if it were an emitter of length $(1-t)\ell_i$ for a polygon skeleton placed at the end of the volume extrusion. If \vec{p} is anywhere else in the space, we say the point is outside the surface by assigning the value 0.

Figure 3.5 shows the resulting implicit surface for different values of t .

Another problem encountered when using emitters is bulging around edges

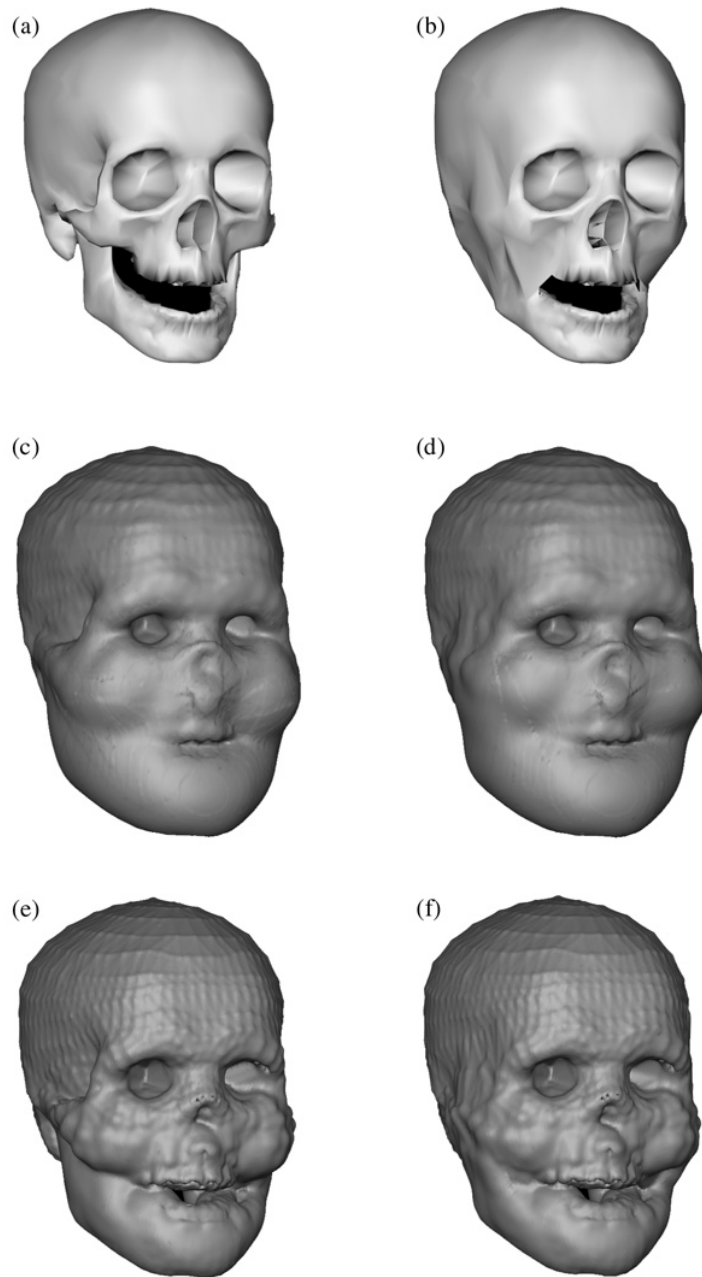


Figure 3.5: Original and filled skull for varying t . (a) Original skull and mandible; (b) skull with jaw and temples filled in; (c) and (d) $n = 150$, $k = 6.0$, $t = 0.0$; (e) and (f) $n = 150$, $k = 6.0$, $t = 0.5$. (c) and (e) use original skull; (d) and (f) use the filled skull.

where emitters meet. This occurs because emitters are in proximity to each other along these edges (and in fact are usually coincident) causing larger field strength sums than for points that project onto the middle of emitter polygons. This can be seen in both 2D and 3D in Figure 3.2. Although the new emitter described by equation 3.2 reduces this bulging, this is an inherent shortcoming of skeletal emitters. The value of k can be increased to reduce this bulging, but too high a value will reduce the surface blending effect of the implicit surface for emitters in close proximity to each other. It is anticipated that a new emitter definition can be constructed to remove this restriction. At the very least it should be possible to define a polygonal emitter so that the resulting surface is flat when adjacent polygon emitters are co-planar.

Polygon reduction is also used on the resulting mesh. Since we must conduct the march at a highest resolution, the resulting mesh has an extremely high polygon count of tiny polygons. The method described by Garland and Heckbert [GH97] is used to reduce the resulting mesh which also provides the added bonus of smoothing out the aforementioned ridges produced by adjacent polygons. Since these ridges are small features compared to the overall facial shape they disappear as a result of the mesh reduction process. Figure 3.6 shows examples of meshes before and after these mesh decimations.

As the implicit surface sampling rate gets arbitrarily high, the act of visiting all n^3 cubes will begin to affect the speed of the implicit surface polygonizer. This is true even if most cubes can be trivially ignored and don't require emitter evaluation on their corners. To avoid this, we can traverse the octree to visit only those boxes which have emitters that contribute to the implicit function on the current n^2 level of boxes. A list of these visited boxes can be maintained so that only these boxes

are initialized between marching levels.

3.4 Results

The final produced implicit surface can be seen in Figure 3.7. This is marched at a resolution of 150 marching cubes in each extent and produced with the parameters $k = 6.0$ and $t = 0.5$ to produce a model of 350,679 polygons. This is reduced using a iterative contraction of vertex pairs [GH97] down to 14000 polygons.

Facial features can be added to the implicit model by fusing polygonal facial features to the polygonized implicit surface. These facial features can be constructed in standard 3D modelling packages or by acquiring 3D scans of template subjects.




Figure 3.6: Polygon reduced implicit surfaces. The above shows the fully polygonized implicit surface model each with greater than 340,000 polygons (left) and the corresponding polygon reduced model to 14,000 polygons (right). $t = 0.0$ for (a) and (b), $t = 0.5$ for (b) and (c); $n = 150$ and $k = 6.0$ for all models.




Figure 3.7: The final implicit surface. $n = 150$, $k = 6.0$, $t = 0.5$, reduced to 14,000 polygons.

Chapter 4

Reconstructing Facial Models with Hierarchical B-Splines

Performing a craniofacial reconstruction with a hierarchical B-spline (H-spline) surface is the focus of this chapter. In [Arc97], Archer proposed a first formalization of computational 3D craniofacial reconstruction and an algorithm to fit hierarchical B-splines to standard dowel locations. She found the number of dowels in the standard dowel set to be insufficient to produce a satisfactory fit. This chapter is a continuation of that work.

Section 4.1 provides background information on hierarchical B-splines followed by Section 4.2 which describes Archer's fitting algorithm and the fitting process. Section 4.3 discusses issues related to facial features. The chapter concludes with results in Section 4.4.

4.1 Overview of Hierarchical B-Spline Surfaces

Hierarchical B-splines are an extension to B-splines.¹ Refer to [FvDFH90] for a good overview of B-splines and [FB88] for a complete description of hierarchical B-splines. A brief overview of hierarchical B-splines is provided for the convenience of the reader.

Refinement of a B-spline surface will be described first. If we start with an $m \times n$ array of control vertices which defines a B-spline surface, we can redefine the surface by increasing the size of the array in either the m or n extent. This produces a surface of $m + 1 \times n$ or $m \times n + 1$ control vertices that represent the same surface. This refinement process is discussed in [WW92]. This is referred to as **refinement** of the surface. Hierarchical B-splines provide a hierarchical means of refining and editing a surface locally.

Definition 4.1 *A local refinement about a B-spline control vertex \vec{x} is a 7×7 array of B-spline control vertices centered at \vec{x} which redefines exactly the four B-spline patches adjacent to (u_0, v_0) , the parametric location of \vec{x} .*

Definition 4.1 implies that the u and v spacing of this new local surface is $\frac{1}{2}$ that of the spacing for the original surface. See Figure 4.1 for a pictorial example of a local refinement.

Definition 4.2 *An overlay is a local refinement of a surface such that allows the central control point, but no other, to be edited.*

The essential idea behind H-splines is that of the overlay. This concept will be illustrated by an example. Take a control point p on a surface. If we wished to

¹Throughout this work, the term “B-splines” is used to refer to bicubic B-Splines.




Figure 4.1: Local refinement. The original control vertices and patches are shown as circles and solid lines and the locally refined control vertices and patches are shown as crosses and lines (both solid and dashed), respectively.

edit p on a higher level of detail, we would produce a new B-spline surface around p that exactly matches the patches adjacent to p as discussed above. Note that as long as p is the only point edited at this higher resolution, only the region of influence of p (i.e, the four adjacent lower level patches) can be modified. Evaluating this surface using only the most refined patch surrounding a given parametric (u, v) location results in a C^2 continuous surface. By only moving the central control, only the surfaces created by the local refinement in Definition 4.2 will be modified.

Overlays can be extended into composite overlays.

Definition 4.3 *A composite overlay is created when two composite overlays of a given level overlap (i.e., share control vertices). The two composite overlays are merged into a new composite overlay with coincident control vertices treated as one. A control vertex may only be edited in a composite overlay if it is the center control*

vertex of a 7×7 array of control vertices. An overlay is the simplest form of composite overlay.

These composite overlays can in turn be treated as B-spline surfaces and be refined further using overlays. A surface can then be specified with an arbitrary amount of detail by progressively refining the surface.

There remains an important aspect of H-splines to be covered that pertains to editing: **offset referencing**. If we kept the H-spline representation as a series of composite overlays in absolute coordinates it would be difficult to edit a control vertex which has been refined (i.e., which has the same (u, v) value as a higher level control vertex). If this point is moved we would like all lower level overlays to move with the lower detail surface in an intuitive manner.

To provide an easy means of editing, vertices at higher levels of detail are stored by means of offset referencing. Instead of storing detail vertex locations in world coordinates, a vertex \vec{x} is stored as an offset \vec{o} . Let (u_0, v_0) be the parametric location in which \vec{x} has the greatest influence. The offset is calculated from \vec{r} , the next-lower surface evaluated at (u_0, v_0) , in a frame of reference defined by the u direction and v direction tangent vectors at (u_0, v_0) . Let F be the transform from this frame to world coordinates. The actual location of \vec{x} can then be calculated as:

$$\vec{x} = \vec{r} + F\vec{o}$$

When control vertices at lower levels of detail influence location (u_0, v_0) on the surface, frame F is implicitly modified in a way that ensures the influence of \vec{x} follows the surface.




Figure 4.2: Generic face hierarchical B-spline surface. The original face model is shown (a) with its parametric associations to the standard dowel and (b) rendered separately.

4.2 Fitting Hierarchical B-Spline Surface

The algorithm which fits the hierarchical B-spline surface to dowel endpoints is treated as a black box for the purposes of the Craniofacial Reconstruction program. In this work, the algorithm from [Arc97] is used to fit the hierarchical B-spline to the dowel locations. A quick overview of this algorithm is provided.

The fit starts with a generic facial model as shown in Figure 4.2. Each of the dowels is treated as a data point to be interpolated by the hierarchical B-spline. A parametric (u, v) location is associated with each dowel manually; this parametric association on the standard dowel set is used for all fits on all skull models. The head is translated and scaled around the skull so that the hierarchical B-spline nose is approximately positioned in its final position and the standard dowels are relatively close to their associated (u, v) locations on the hierarchical B-spline.

Fitting is done as an iterative process. A pair of parameters are passed to each iteration: L , the level of the H-spline to fit, and f , a fit fraction, $0 \leq f \leq 1$.

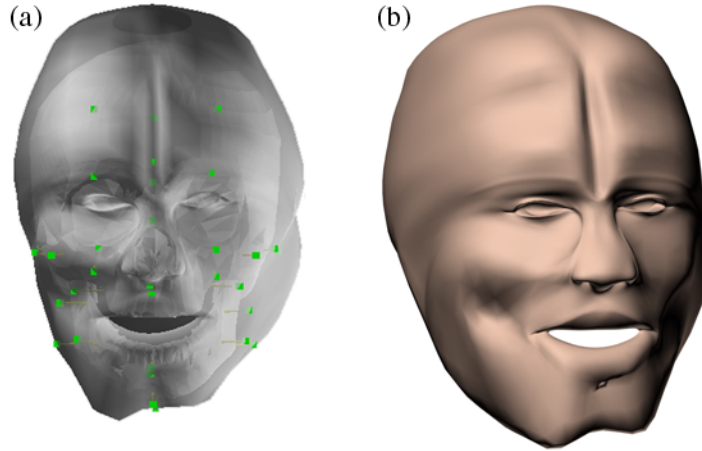


Figure 4.3: H-spline surface fit to standard dowel locations. The fit facial model is shown (a) with its parametric associations to the standard dowel and (b) rendered separately.

The L^{th} level of the H-spline is treated as a simple B-spline and fit by least-squares. The fit fraction f is used to limit the effect of the iteration on the overall fitting process. The position of control vertex i after an iteration of the fit is $\vec{c}_i + f(\vec{c}'_i - \vec{c}_i)$, where \vec{c}_i and \vec{c}'_i are the positions of control vertex i before and after the least squares fit.

The standard dowels are not sufficient to produce a satisfactory fit as demonstrated in Figure 4.3. While the H-spline surface does take shape about the standard dowel location from iterations of the fitting algorithm on lower detail levels of the H-spline surface, there are many higher detail control vertices that have no parametric associations over their area of influence. Such control vertices will not be moved during the fitting process. Since only the high detail control vertices near standard dowel locations can be moved, ridges and dents occur in the H-spline surface around these locations.

To ensure a reasonable fit, dowels must be inserted both outside and inside

the dowel areas' region of influence. Dowels outside this area of influence are given a default small tissue depth; for the purposes of this application this tissue depth was the same as that for the supraglabella dowel on the mid-upper forehead. Since the hierarchical B-spline surface is approximately placed in the correct location and scaled around the skull, we can calculate a (u, v) location to associate with the dowel by intersecting a ray cast by the dowel with the H-spline surface. This association can be modified if necessary to reduce distortion of the fit on the hierarchical B-spline surface.

When a dowel is inserted inside a dowel area's region of influence its tissue depth is interpolated as described in Section 2.2. The (u, v) coordinate of the dowel is interpolated as well; however this will not always produce a correct location. Since the parametric surface is not spaced evenly, the distance along the parametric surface will not directly correspond to the interpolated u and v . The dowel's corresponding parametric location can be edited manually if necessary.

The quality of the fit is highly sensitive to the quality of the dowel placement. Too few data points can result in both surface buckling and poor overall surface shape as seen in Figure 4.3. Too many data points will overload the fitting algorithm and again produce a poor interpolation. It appears the best results are attained when a data point exists for each patch at the highest level of fitting. This is level two of the hierarchical B-spline.²

Another problem encountered when fitting the generic H-spline face is distortion of facial features. Since we wish to preserve the structure of the face in order to easily edit facial features, it is important to keep the facial features mostly symmetrical and unchanged. Offsets at levels of the H-spline higher than two are

²Hierarchical B-spline levels are by convention counted from 0, the lowest level of detail.

not modified during the fitting process.

Even if one takes these pitfalls into account, it is inevitable that the model will require touch-ups. This task can be completed with relative ease as hierarchical B-splines allow for easy editing of the model.

4.3 Adding Facial Features

Facial features such as the nose, eyes, and mouth are an implicit part of the initial facial model. The only issues to address are the addition of different facial features and the effects of fitting.

Facial features can be modified and distorted during the fitting process. It is therefore desirable to provide a means of limiting the effect to such regions. At first, it would seem like a good idea to ensure that the fitting algorithm does not move such vertices; however, this often produces poor results as a greater weight is put on other control vertices near facial feature offsets. Instead, fitting is conducted as if facial feature offsets can be moved; the original offsets are restored after fitting is complete. This produces the most visually satisfying results.

Changing facial features is handled in a similar manner. A library of different facial features could be maintained as a database of control vertex offsets. Though not implemented for this work, a tool for blending between facial feature offsets would provide added flexibility to the artist. As well, the hierarchical B-spline surface can be touched up after facial features have been added.




Figure 4.4: H-spline surface fit to interpolated dowels. Presented are (a) a semi-transparent fit surface with interpolation points, (b) a rendering of the h-spline surface fit to the dowel locations, and (c) the model in (b) edited to reduce facial distortions.

4.4 Results

A final fit model is presented in Figure 4.4(b); the same model, modified manually to remove visual distortions, is shown in Figure 4.4(c). Given as tuples in the form (fit level, fit fraction), the iterations of the fit were $(0, 0.4)$, $(1, 0.35)$, $(2, 0.35)$, $(0, 0.35)$, $(1, 0.3)$, and $(2, 0.3)$.

As mentioned previously, the fitting process is highly sensitive to the parametric locations to which dowels are associated. Currently, the parametric locations of the dowels are manually modified in order to produce an undistorted fit. This process can be improved in the future. One possible approach would be an iterative relaxation-based approach; parametric locations can be moved along the surface based on a penalty formula. The fitting process would be conducted at each iteration, and (u, v) locations moved in areas that cause surface creases, bunching, and distortions. Such a technique would reduce the subjectivity of parametric locations.

Chapter 5

Conclusion

This chapter presents the final results of the computational facial reconstruction techniques in Section 5.1, an summary of the thesis in Section 5.2, and finishes with a discussion of future directions for the research in Section 5.3.

5.1 Results

For comparative purposes, the final facial model for both the implicit surface and hierarchical B-spline fitting techniques is provided in Figures 5.1(a) and (b), respectively. These figures show the raw outputs of each technique without modifications to the facial features or model touch-ups. Note that the hierarchical B-spline includes facial features in the generic surface and these features are carried on to the final result. As well, Figures 5.2 and 5.3 show the implicit surface and hierarchical B-spline fitting techniques for different subject weights.

Each of the two techniques presented has its strengths and weaknesses. The implicit surface technique provides an almost automated means of producing facial shape. Once the standard dowels are placed, the algorithm generates the the three




Figure 5.1: Results of (a) the implicit surface and (b) hierarchical B-spline fitting techniques for an average weight male caucasoid.




Figure 5.2: Implicit surface method results for different weight types.

dimensional model. In contrast, the H-spline fitting technique requires manual placement of additional dowels, modification of dowel parametric placement to remove facial distortions, and experimentation with fitting parameters. These additional tasks add a further element of subjectivity to the process.

However, to its merit, the H-spline fitting technique is better suited to the addition of facial features and editing of the output model. Unlike the implicit surface reconstruction method for which facial features must be added as a separate

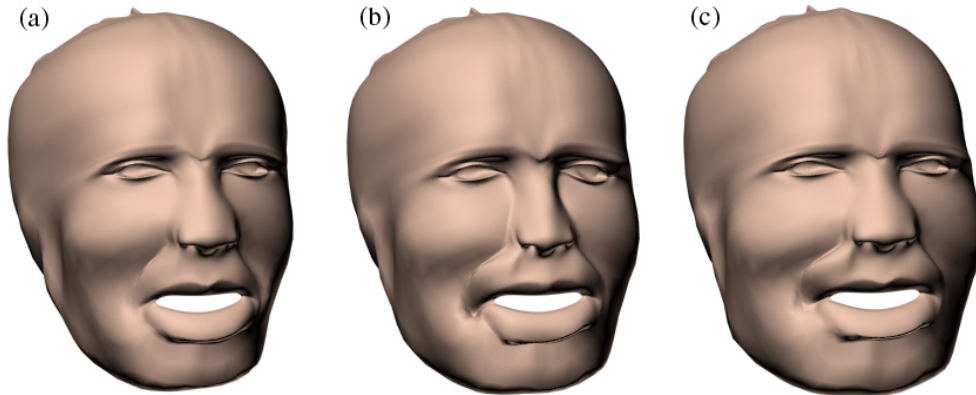


Figure 5.3: Hierarchical B-spline fitting method results for different weight types.

process, facial features are integral elements of both the original general and output H-spline surface. As detailed in Section 4.3, a library of facial features can be maintained and blended between; each facial detail is stored as H-spline offsets for specific control vertices. As well, the resulting H-spline surface is much easier to modify than the raw mesh produced from the implicit surface technique.

5.2 Summary

This work has detailed the problem of computer assisted 3D craniofacial reconstruction in a series of phases. Chapter 1 describes the context of the problem; the long-term goal of this work is to provide the practical foundations of tools to produce a subject's pre-death appearance. This thesis concentrates on the reproduction of three dimensional facial shape from the skull and mandible.

The reproduction of facial shape is accomplished in two steps. The first is to provide a means of determining the facial tissue depth at any point on the skull; this is based on the lengths of the dowels at the standard facial landmarks and is describes in Chapter 2. The second step is to generate the facial model. Two complementary

methods are detailed. Chapter 3 describes a technique to reconstruct the facial shape using the polygons of the skull and mandible as emitters of an implicit surface. Chapter 4 describes the task as the problem of hierarchical B-spline surface fitting. Each of the standard dowels is given a corresponding parametric location on the H-spline surface; this parametric location indicates which location on the H-spline surface is to be fit to the end point of the dowel. Additional dowels are manually added with the length and parametric location determined computationally. The parametric location will sometimes need to be adjusted manually.

5.3 Future Work

Future work could proceed in two directions: validation of the reconstruction process and improvement of reconstruction techniques. Each will be discussed.

Currently, the tissue depth of a given point on the skull is determined by barycentric interpolation of the point projected onto a triangle prescribed by three dowels. The results of this technique, and shown in Figure 2.4, were acceptable according to RCMP forensic artists. While promising, a less subjective analysis of the results is desirable. It could be that expert knowledge of the forensic artist would have to be incorporated into any automatic system for estimating tissue depths. An analysis of interpolated tissue depths should be conducted on cadavers or living subjects. The tissue depth of a subject can be measured at a large number of facial locations, including the standard dowel set. These measurements could be acquired using traditional invasive techniques and CT scans on cadavers or safe tissue depth methods such as MRI scans on living subjects. These measurements can be compared to the interpolated tissue depths; the lengths of the standard dowels would be set to the measured standard landmark tissue depths for the subject. Once there is

a means of validating the interpolation method, other interpolation techniques such as natural neighbour interpolation [WP87, Wat92] can be investigated.

As well, there is currently no validation of the resulting facial model to the original subject. Both a quantitative and qualitative analysis of facial reconstruction results is needed. As already mentioned, an MRI scan can be taken of a living subject. The facial model of the subject, as extracted from the scan, can be compared with the reconstructed facial model. A metric, such as volume between the scanned and reconstructed facial models, can be computed to measure reconstruction error.

In addition, a qualitative study is required to determine whether the reconstructed facial models look similar to the subject according to human perception. This is as much a test of the validity of forensic craniofacial reconstruction as a test of an individual reconstruction technique. Case studies and reviews of physical reconstructions have been conducted (e.g., [HRPM93, TEC⁺97]); however, since reconstructions can be produced computationally, can be done on living subjects, and can be produced with relative ease, a proper statistical experiment for the purposes of identification can be conducted on a scale previously impractical.

Once a means of validating a reconstruction is established, extension and new techniques for reconstruction can be explored. For example, the manual insertion of dowels on the skull in the H-spline fitting method is a time consuming process. Instead of inserting fitting dowels on the skull, points on the output implicit surface model could be automatically added. Regions such as the nose, mouth, and orbitals would be masked so that in these regions would not be added. A better means of associating these points with H-spline parametric locations, such as that discussed in Section 4.4, should be explored.

Other reconstruction methods which should be explored include a relaxation

based spring model approach and volumetric deformation techniques. The relaxation based spring model would represent the facial surface as a network of springs. Vertices of the network would be kept a distance from the skull by either keeping them in place with springs of resting distance equal to the tissue depth or by not allowing the vertices to cross approximate meshes (such as those in Figures 2.4(b) and 3.5) of the facial surface. Particular vertices in the spring network would be locked to specific facial locations such as the standard landmark positions and the edges of facial features such as the eyes, nose, and mouth. Models of facial features could be chosen from a database and automatically added to the resulting surface since facial feature locations are set in the network. Volumetric deformation techniques, as discussed in Section 1.1, can be extended. Different deformation techniques, for both the three dimensional skull/facial models and for two dimensional cross-sections of the volume such as those produced by MRI scans, could be developed and compared.

It is evident the problem of forensic craniofacial reconstruction is an active area of research that requires much future work.

Bibliography

- [Arc97] Katrina Archer. Craniofacial reconstruction using hierarchical B-spline interpolation. Master's thesis, University of British Columbia, 1997.
- [FB88] David R. Forsey and Richard H. Bartels. Hierarchical B-spline refinement. In *Proceedings of SIGGRAPH*, pages 205–212, 1988.
- [FvDFH90] James D. Foley, Andries van Dam, Steven K. Feiner, and John F. Hughes. *Computer Graphics: Principles and Practice*, chapter 11. Addison-Wesley, 1990.
- [Geo87] Robert M. George. The lateral craniographic method of facial reconstruction. *Journal of Forensic Science*, 32:1305–1330, 1987.
- [Geo93] Robert M. George. *Forensic Analysis of the Skull*, chapter 16: Anatomical and Artistic Guidelines for Forensic Facial Reconstruction, pages 215–227. Wiley-Liss, 1993.
- [GH97] Michael Garland and Paul S. Heckbert. Surface simplification using quadric error metrics. In *Proceedings of SIGGRAPH*, pages 209–216, 1997.

- [HRPM93] Richard P. Helmer, S. Rohricht, D. Petersen, and F. Mohr. *Forensic Analysis of the Skull*, chapter 17: Assessment of the Reliability of Facial Reconstruction, pages 229–246. Wiley-Liss, 1993.
- [KI86] W.M. Krogman and M.Y. Iscan. *The Human Skeleton in Forensic Medicine*. Charles C. Thomas, Springfield, IL, 2nd edition, 1986.
- [LC87] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Proceedings of SIGGRAPH*, pages 163–169, 1987.
- [NM98] L.A. Nelson and S.D. Michael. The application of volume deformation to three-dimensional facial reconstruction: A comparison with previous techniques. *Forensic Science International*, 94(3):167–181, June 1998.
- [QCS⁺97] G. Quatrehomme, S. Cotin, G. Subsol, H. Delingette, Y. Garidel, G. Grevin, M. Fidrich, P. Bailet, and A. Ollier. A fully three dimensional method for facial reconstruction based on deformable objects. *Journal of Forensic Science*, 42(4):649–652, 1997.
- [Rhi84] Stanley Rhine. Tissue thickness measures: American caucasoids, american blacks, southwestern indians. Physical Anthropology Laboratories, Maxwell Museum of Anthropology, University of New Mexico, 1982, 1983, 1984.
- [Sam89] Hanan Samet. *Applications of Spatial Data Structures*. Addison-Wesley Publishing Company, 1989.
- [Ste79] T.D. Stewart. *Essentials of Forensic Anthropology*. Charles C. Thomas, springfield, il edition, 1979.

- [TEC⁺97] Andrew J. Tyrell, Martin P. Evison, Andrew T. Chamberlain, Michael A. Green, and F.R.C. Path. Forensic three-dimensional facial reconstruction: Historical review and contemporary developments. *Journal of Forensic Science*, 42(4):653–661, 1997.
- [VBL⁺89] P. Vanezis, R.W. Blowes, A.D. Linney, A.C. Tan, R. Richards, and R. Neave. Application of 3d computer graphics for facial reconstruction and comparison with sculpting techniques. *Forensic Science International*, 42:69–84, 1989.
- [vOW93] Kees van Overveld and Brian Wyvill. Potentials, Polygons and Penguins. An efficient adaptive algorithm for triangulating an equipotential surface . In *Proc. 5th Annual Western Computer Graphics Symposium (SKIGRAPH 93)*, pages 31–62, 1993.
- [Wat92] David F. Watson. *Contouring: A Guide to the Analysis and Display of Spatial Data*. Pergamon, 1992.
- [WP87] D.F. Watson and G.M. Philip. Neighborhood-based interpolation. *Geobyte*, 2(2):12–16, 1987.
- [WW92] Alan H. Watt and Mark Watt. *Advanced Animation and Rendering Techniques: Theory and Practice*. Addison-Wesley Pub Co, 1992.

Appendix A

Tissue Thickness Measurements

This appendix reprints tissue depths from studies conducted by Rhine and Stanley [Rhi84].

No.	Measurement Name	Emaciated		Normal		Obese	
		Male	Female	Male	Female	Male	Female
1	Supraglabella	2.25	2.50	4.25	3.50	5.50	4.25
2	Glabella	2.50	4.00	5.25	4.75	7.50	7.50
3	Nasion	4.25	5.25	6.50	5.50	7.50	7.00
4	End of Nasals	2.50	2.25	3.00	2.75	3.50	4.25
5	Mid Philtrum	6.25	5.00	10.00	8.50	11.00	9.00
6	Upper Lip Margin	9.75	6.25	9.75	9.00	11.00	11.00
7	Lower Lip Margin	9.50	8.50	11.00	10.00	12.75	12.25
8	Chin-Lip Fold	8.75	9.25	10.75	9.50	12.25	13.75
9	Mental Eminence	7.00	8.50	11.25	10.00	14.00	14.25
10	Beneath Chin	4.50	3.75	7.25	5.75	10.75	9.00
11	Frontal Eminence	3.00	2.75	4.25	3.50	5.50	5.00
12	Supraorbital	6.25	5.25	8.25	7.00	10.25	10.00
13	Suborbital	2.75	4.00	5.75	6.00	8.25	8.50
14	Inferior Malar	8.50	7.00	13.25	12.75	15.25	14.00
15	Lateral Orbit	5.00	6.00	10.00	10.75	13.75	14.75
16	Zygomatic Arch, midway	3.00	3.50	7.25	7.50	11.75	13.00
17	Supraglenoid	4.25	4.25	8.50	8.00	11.25	10.50
18	Gonion	4.50	5.00	11.50	12.00	17.50	17.50
19	Supra M ²	12.00	12.00	19.50	19.25	25.00	23.75
20	Occlusal Line	12.00	11.00	18.25	17.00	23.50	20.25
21	Sub M ²	10.00	9.50	16.00	15.50	19.75	18.75

Table A.1: Tissue thickness measurements (in mm), for American Caucasoids.

Measurement		Emaciated		Normal		Obese	
No.	Name	Male	Female	Male	Female	Male	Female
1	Supraglabella	4.00	5.00	5.00	4.50	5.00	3.50
2	Glabella	5.25	6.00	6.25	6.00	7.50	6.00
3	Nasion	5.25	5.25	6.00	5.25	5.25	4.75
4	End of Nasals	3.00	3.25	3.75	3.75	3.25	3.00
5	Mid Philtrum	11.75	10.00	12.25	11.25	11.75	12.00
6	Upper Lip Margin	12.50	12.00	14.25	12.50	12.50	15.25
7	Lower Lip Margin	13.75	12.25	15.50	15.00	15.50	12.00
8	Chin-Lip Fold	11.75	9.50	11.75	12.25	13.00	12.25
9	Mental Eminence	11.25	11.00	11.50	12.50	15.25	13.00
10	Beneath Chin	8.00	6.50	8.25	8.00	9.50	8.50
11	Frontal Eminence	3.75	3.25	5.00	4.00	5.50	5.00
12	Supraorbital	7.75	7.25	8.50	8.00	11.75	8.50
13	Suborbital	5.75	6.50	7.75	8.25	9.25	9.00
14	Inferior Malar	14.00	14.50	16.50	16.75	17.50	18.75
15	Lateral Orbit	10.50	12.00	13.25	13.00	20.00	12.75
16	Zygomatic Arch, midway	6.75	8.00	8.25	9.50	13.75	9.25
17	Supraglenoid	9.50	9.75	11.00	11.50	17.50	17.25
18	Gonion	11.50	11.00	13.00	13.50	24.00	17.50
19	Supra M ²	19.00	20.50	22.00	20.25	24.00	23.50
20	Occlusal Line	16.75	17.75	19.00	19.25	30.00	20.00
21	Sub M ²	13.50	14.25	16.50	17.00	23.50	20.00

Table A.2: Tissue thickness measurements (in mm), for American Blacks.

No.	Measurement Name	Emaciated		Normal		Obese	
		Male	Female	Male	Female	Male	Female
1	Supraglabella	5.75	4.00	5.00	4.50	4.50	4.25
2	Glabella	5.75	4.75	5.75	4.50	6.00	4.50
3	Nasion	5.75	6.50	6.86 ¹	7.00	6.50	5.00
4	End of Nasals	2.75	2.50	3.50	2.50	3.25	3.25
5	Mid Philtrum	7.50	10.00	9.75	10.00	9.25	8.51 ¹
6	Upper Lip Margin	8.25	9.50	9.75	11.00	9.25	10.00
7	Lower Lip Margin	9.25	12.00	11.00	12.25	8.75	11.25
8	Chin-Lip Fold	8.50	9.00	11.50	10.00	9.75	11.00
9	Mental Eminence	8.00	11.00	12.00	13.00	12.50	13.25
10	Beneath Chin	5.25	8.00	8.00	8.00	8.00	7.75
11	Frontal Eminence	4.75	4.75	4.25	4.00	4.50	4.20 ¹
12	Supraorbital	6.75	5.00	9.00	8.50	8.50	8.25
13	Suborbital	3.75	3.25	7.50	6.25	7.75	6.75
14	Inferior Malar	10.00	9.00	14.00	12.00	15.75	15.00
15	Lateral Orbit	8.00	8.25	12.50	11.50	11.75	13.75
16	Zygomatic Arch, midway	6.00	5.75	7.50	7.00	8.75	9.00
17	Supraglenoid	5.75	4.50	8.50	6.25	9.75	7.75
18	Gonion	7.75	6.25	13.25	10.50	15.40 ¹	12.75
19	Supra M ²	14.25	11.75	21.50	18.00	23.50	19.00
20	Occlusal Line	15.50	12.25	20.75	17.50	22.75	19.25
21	Sub M ²	12.50	10.50	19.25	17.00	18.50	15.75

Table A.3: Tissue thickness measurements (in mm), for Southwestern Indians.

¹This may be a typographical error within the original data.