

# Extrapolation of Missing Craniofacial Skeletal Structure via Statistical Shape Models

Project #1 Seminar Report, EN.600.646 Spring 2014

Robert Grupp  
robert.grupp@gmail.com

April 4, 2014

## 1 Introduction

### 1.1 Reviewed Papers

- S. Benazzi and S. Senck. Comparing 3-Dimensional Virtual Methods for Reconstruction in Craniomaxillofacial Surgery. *Journal of Oral and Maxillofacial Surgery*, 69(4):1184–1194, 2011.
- M. Kazhdan, M. Bolitho, and H. Hoppe. Poisson Surface Reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, 2006.

### 1.2 Motivation

The goal of this project is to design and implement a method for extrapolating missing anatomical craniofacial skeletal structure with the use of a statistical shape model (SSM) of the human cranium. The procedure of interest is craniofacial transplantation, which is the process of transplanting a donor’s craniofacial soft tissue, and possibly bone structure, onto a patient that has been subject to some severe craniofacial deformation. Figure 1 shows preoperative and postoperative views of a transplant recipient. By reliably estimating the “original” skeletal structure of the patient, it is hoped that the surgical planning process may be enhanced, leading to more positive surgical outcomes. The papers by Benazzi, et al. [3] and Kazhdan, et al. [9] propose different methods of surface reconstruction; it is hoped that they may be applied for our purposes in this project.

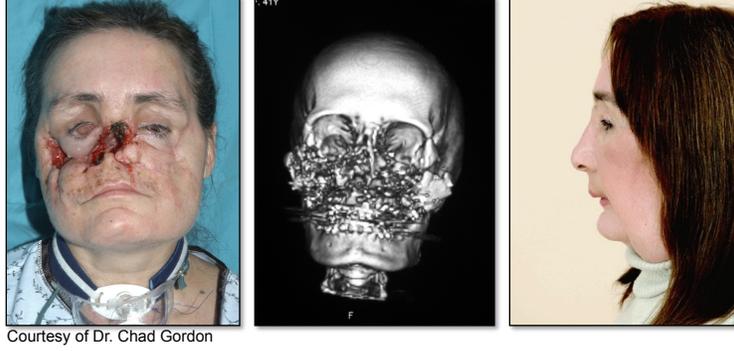


Figure 1: A craniofacial transplant recipient. Left preoperative photograph, middle preoperative CT, right postoperative photograph

### 1.3 Preliminaries

In order to understand precisely how the reviewed papers [3], [9] may be used for our purposes, it is important to first discuss the theoretical concepts for basic shape extrapolation via a SSM. The method described in this section closely follows that described by Chintalapani [6].

Let  $\mathcal{M}_P$  be a triangular mesh representing the skull of a particular patient with some missing, or deformed, region  $\mathcal{R}$ .  $\mathcal{M}_P$  may be represented by the sets of vertices and triangle faces,  $\mathbf{V}_P$  and  $\mathbf{F}_P$ , respectively. Given a SSM of human skull surfaces,  $\mathcal{A}$ , patient-to-atlas registration may be computed to obtain an estimate,  $\mathcal{M}_E$ , of the patient's ideal skull surface (e.g. without a deformity in  $\mathcal{R}$ ). The estimate mesh is also represented by sets of vertices and triangles,  $\mathbf{V}_E$  and  $\mathbf{F}_E$ , respectively. Furthermore,  $\mathcal{M}_E$  may be partitioned into the “unknown” region, represented by  $\mathcal{R}$ , and the remaining “known” region as shown in (1).

$$\mathcal{M}_E = [\mathbf{V}_E, \mathbf{F}_E] = \left[ \begin{pmatrix} \mathbf{V}_E^{\text{known}} \\ \mathbf{V}_E^{\text{unknown}} \end{pmatrix}, \begin{pmatrix} \mathbf{F}_E^{\text{known}} \\ \mathbf{F}_E^{\text{unknown}} \end{pmatrix} \right] \quad (1)$$

Prior to performing the same partition on the original mesh,  $\mathcal{M}_P$ , it is converted to match the mesh topology of  $\mathcal{M}_E$ .  $\mathcal{M}_P$  with updated topology is denoted  $\tilde{\mathcal{M}}_P$ , and is partitioned to match the “unknown” region  $\mathcal{R}$  and remaining “known” region in (2).

$$\tilde{\mathcal{M}}_P = [\tilde{\mathbf{V}}_P, \tilde{\mathbf{F}}_P] = [\tilde{\mathbf{V}}_P, \mathbf{F}_E] = \left[ \begin{pmatrix} \tilde{\mathbf{V}}_P^{\text{known}} \\ \tilde{\mathbf{V}}_P^{\text{unknown}} \end{pmatrix}, \begin{pmatrix} \mathbf{F}_E^{\text{known}} \\ \mathbf{F}_E^{\text{unknown}} \end{pmatrix} \right] \quad (2)$$

A reconstruction is then possible by replacing the “unknown” vertices of the patient mesh with the “unknown” vertices of the atlas estimate, shown in (3)

$$\tilde{\mathcal{M}}'_P = \left[ \begin{pmatrix} \tilde{\mathbf{V}}_P^{\text{known}} \\ \mathbf{V}_E^{\text{unknown}} \end{pmatrix}, \mathbf{F}_E \right] \quad (3)$$

The major issue with  $\widetilde{\mathcal{M}}'_P$  is the possible presence of a transition that is not smooth, or “jaggy,” [6] as shown in figure 2. In order to perform accurate surgical planning and biomechanical analysis, it is desired to minimize, or remove, these non-smooth transitions. It is hoped that the reviewed papers may be used for this purpose.

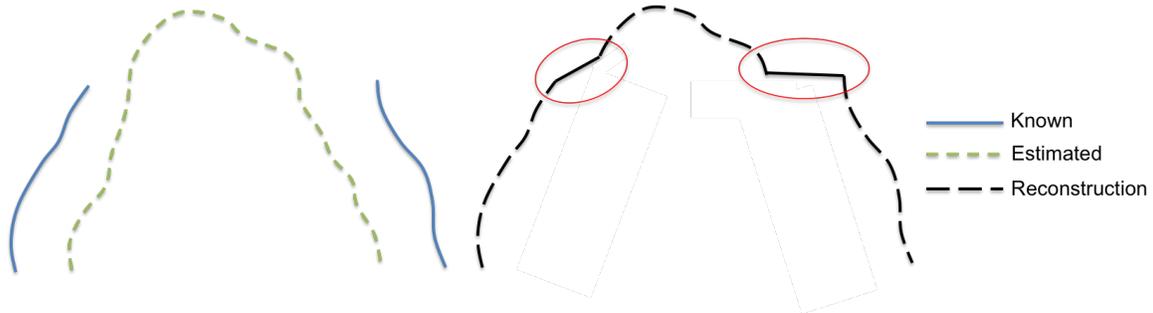


Figure 2: A 2D example of the atlas based extrapolation method. The left figure shows the original patient CT in solid blue and the atlas estimation in dashed green. The right figure depicts the extrapolated surface and the non-smooth transitions highlighted with red ellipses.

## 2 Benazzi, et al.

### 2.1 Introduction

The paper “Comparing 3-Dimensional Virtual Methods for Reconstruction in Craniomaxillofacial Surgery” by Benazzi, et al. proposes several methods for virtually reconstructing the missing skeletal structure after a virtual zygomatic osteotomy. Although the explicit application of this process is not explicitly discussed, similar virtual reconstructions have been shown to be useful for surgical planning and creation of a cutting guide; specifically for the surgical repositioning of the zygomatic bone following some trauma [8]. The authors are well established in the field of forensic anthropology and the reconstruction of human and pre-human skull fossils [10][2][1][4]. This publication appears to be their first foray into the field of surgery. The methods used for landmarks and curves by the authors have been previously used in the medical field [7], therefore it is reasonable to expect positive results.

### 2.2 Materials and Methods

The authors obtained CT scans of 15 modern human skulls of varying sex, age, and geographic origin for this experiment [3]. For each skull a virtual osteotomy of one zygomatic bone is performed, with the intent to reconstruct the missing area. After performing the

virtual osteotomy, the spatial constraints of the missing region are saved for future reference [3]. Landmarks are then identified on this CT as shown in figure 3, with landmarks near, and in, the mid-sagittal region used to form a mid-sagittal plane of best fit [3]. The CT scan of the patient is then mirrored about the estimate of the mid-sagittal plane and the spatial constraints from the osteotomy procedure are used to recover the “donor,” or “implant,” bone in the corresponding mirrored CT [3]. This mirroring process is depicted in figure 3, along with what we have estimated to be a possible mid-sagittal plane, and “donor” region. The “donor” zygomatic bone is then virtually transplanted onto the un-mirrored CT with three approaches: a “copy and paste” insertion, an insertion via rigid registration, and an insertion via iterative TPS warping [3]. The authors state the iterative closest point (ICP) algorithm is used for the rigid registration approach, but do not identify and further details [3]. The first two approaches are relatively straight-forward, requiring no additional landmark identification, whereas the TPS method requires the identification of additional landmarks, semi-landmarks, and curves on the mirrored surface [3]. Iterative TPS implementations from the Amira and Edgewarp 3D software packages are used [3]. These methods allow the adaptive updating (or sliding) of the semi-landmarks in order to obtain homologous correspondences [3]. After reconstruction, the surface deviation of the reconstructed area, compared to the original surface, is computed to characterize results. The surface deviation accounts for the oriented estimate lying above/below the true surface with positive/negative measurements in millimeters. The mean value of these deviations, along with the standard deviation of these deviations, is used to indicate the quality of the results.

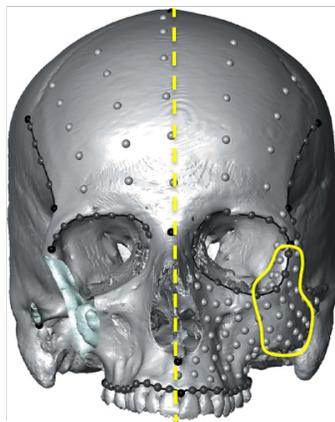


Figure 3: A depiction of landmarks, semi-landmarks, and curves identified on a mirrored skull that will be used to drive the TPS reconstruction. A possible axis of reflection about the mid-sagittal plane is shown as a dotted yellow line, and a possible region to use for reconstruction is overlaid by the yellow curve. (Graphic is from [3])

### 2.3 Results

The authors provide graphic descriptions of the experiments with two of the skulls and a tabular summary for the full set of skulls. Figure 4 shows a subset of the results depicted in [3]. It is clear from these graphics that simple mirroring fails, and the TPS method performs significantly better. It is also obvious from the tabular summary that simple mirroring performance lags behind the rigid and TPS approaches [3]. The authors state that there was no significant difference between the rigid registration results and the TPS registration results, however the graphics included in the article depict much “smoother” edge transitions with the TPS approach; a desired property for our experiment.

The use of *mean* surface deviation is somewhat deceiving. For example, it is possible for half of the estimated surface to lie a constant distance,  $\delta$ , below the surface and for the other half to lie a constant  $\delta$  above the surface, resulting in a zero mean surface deviation. It would have been more appropriate to compute the mean distance of the estimate from the true surface; this would result in a value of  $\delta$ , instead of zero, for the previous example. It would be interesting to also determine if the rigid registration and TPS registration results would remain statistically equivalent if a distance metric was used in place of deviation.

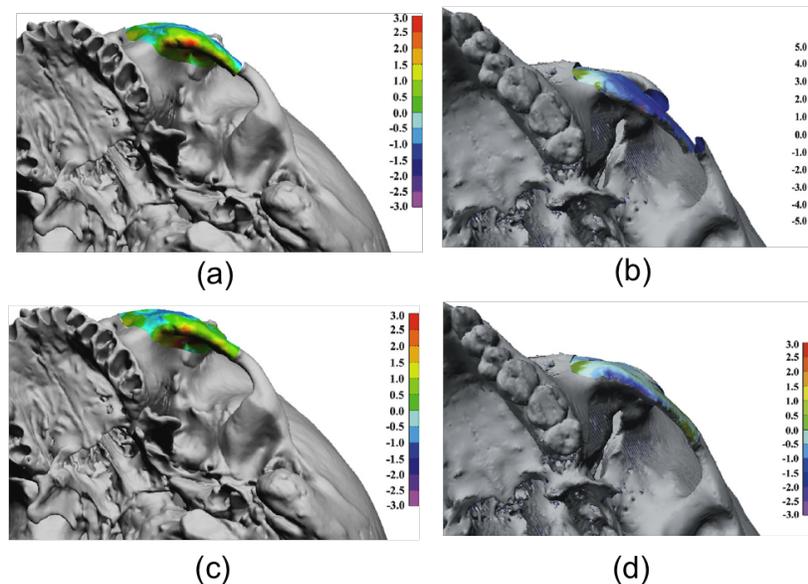


Figure 4: Results of the mirroring and TPS reconstruction methods on two skulls. Surface deviation error is reported in mm according to the color legend. The skull on the left exhibits a larger amount of symmetry than the skull on the right. (a), (b) depict the mirroring reconstruction technique and (c), (d) depict the TPS reconstruction technique. (Graphics are from [3])

## 2.4 Application to CIS Project

It is plausible to use the techniques proposed by Benazzi, et al. in order to seamlessly incorporate the extrapolated “unknown” facial region, estimated via patient-to-atlas registration, into a full reproduction of the patient’s pre-trauma cranium. The mirrored zygomatic bone used in the aforementioned experiments would be replaced by the atlas estimate of the “unknown” facial region with some additional overlap of the “known” region. Landmarks, semi-landmarks, and curves may then be identified on regions of overlap and the TPS reconstruction method may be applied. This would require a non-trivial amount of work to complete. Landmarks, etc. would need to be identified manually and *accurately*, and a TPS registration implementation utilizing “sliding” landmarks would need to be found elsewhere or implemented. Edgewarp 3D is a potential implementation of this TPS registration method, and is freely available, however the PowerPC OSX binaries will not run on a recent Apple operating system, therefore a suitable Linux distribution must be found to run the available Linux binaries [5].

## 3 Kazhdan, et al.

### 3.1 Introduction

The paper “Poisson Surface Reconstruction” by Kazhdan, et al. introduces a new method to reconstruct a 3D surface from a collection of oriented point samples [9]. The general goal of the algorithm is to estimate an indicator function,  $\chi_M$ , which is 1 valued for spatial locations lying in the interior of the surface,  $M$ , and 0 valued for locations lying external to the surface. Once  $\chi_M$  has been estimated the surface boundary,  $\partial M$ , is implicitly defined. The key enabler in this algorithm is (4) which expresses the Laplacian of the indicator function as equal to the divergence of the vector field described by the oriented points [9]. The details of the derivation of (4) and its use will be summarized in the next section.

$$\Delta \tilde{\chi} = \nabla \cdot \vec{V} \quad (4)$$

The authors claim that the Poisson approach has the advantages of a global solution considering “all data at once,” but not susceptible to local approximations or heuristics. Additionally, the method is claimed to be robust in the presence of noise, due to the use of a “hierarchy of locally supported functions,” compared to the global support provided by a radial basis function (RBF) approach [9]. This hierarchy is provided by an octree of linearized Gaussians and will be further discussed in the following section.

As a general comment, it would be nice to see this method published as a journal paper, where more explanation may be given to the mathematical derivations of the algorithm. This would most likely reduce the time required to fully understand the approach and its mathematical subtleties.

### 3.2 Materials and Methods

The reconstruction method is derived through a series of mathematical formalisms that shall be summarized here. (5) describes the initial approximation of the gradient of the indicator function,  $\vec{V}$ . The authors show that the gradient of the smoothed true indicator function,  $\chi_M$ , is equivalent to the integration of the smoothed surface normals [9]. When discretizing the problem, the equality becomes an approximation as indicated in (5). The vector field is sampled at a greater number of points than in the actual point cloud. For each sampled point,  $q$ , the vector field value is computed as the weighted sum over all normal vectors in the original point cloud.  $|\mathcal{P}_s|$  represents the size of the surface partition containing point  $s$ ,  $\tilde{F}_{s,p}(q)$  represents a spatial weighting, or smoothing, function centered at the original point  $s$  relative to the current sample point  $q$ , and  $s.\vec{N}$  represents the normal vector at the original point  $s$ . It is important to note that this is a theoretical result, and the actual computation will be shown later in (10).

$$\nabla(\chi_M * \tilde{F})(q) \approx \sum_{s \in S} |\mathcal{P}_s| \tilde{F}_{s,p}(q) s.\vec{N} \equiv \vec{V}(q) \quad (5)$$

Now, the estimate of the gradient of the indicator function is equal to the computed vector field, as shown in (6). The authors remark that  $\vec{V}$  is not usually integrable, therefore the divergence operator ( $\nabla \cdot$ , note the dot used to distinguish from the gradient) is applied to (6) to obtain a Poisson equation representation shown in (7) [9].

$$\nabla \tilde{\chi} = \vec{V} \quad (6)$$

$$\Delta \tilde{\chi} = \nabla \cdot \nabla \tilde{\chi} = \nabla \cdot \vec{V} \quad (7)$$

The authors use an octree, composed of a “multi-resolution” of basis smoothing functions, in order to efficiently approximate the sampled vector field,  $\vec{V}$  [9]. For each node,  $o$ , in the octree,  $\mathcal{O}$ , the node function is defined as in (8).  $o.c$  represents the center location of that node’s area and  $o.w$  represents its width. The function,  $F$ , is a linearized approximation of a smoothing function defined in (9); it is the successive convolution of 3D box function with itself to form a “3D pyramid.” The octree is defined with a maximum depth,  $D$ , and nodes are created so that every sample point will fall within a leaf node’s represented volume [9].

$$F_o(q) \equiv F\left(\frac{q - o.c}{o.w}\right) \frac{1}{o.w^3} \quad (8)$$

$$F(x, y, z) \equiv (B(x)B(y)B(z))^{*n} \quad \text{with} \quad B(t) = \begin{cases} 1 & |t| < 0.5 \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

Once the octree is created, the vector field,  $\vec{V}$ , is computed via (10). This computes  $\vec{V}(p)$  as the weighted sum of each of the point cloud normals, taking into account the eight

neighboring 3D octree nodes about each point cloud point.  $\alpha_{o,s}$  represents the trilinear interpolation weights computed for the eight 3D neighbors [9].

$$\vec{V}(q) \equiv \sum_{s \in \mathcal{S}} \sum_{o \in \text{Ngb}_{\mathcal{D}}(s)} \alpha_{o,s} F_o(q) s \cdot \vec{N} \quad (10)$$

Once  $\vec{V}$  is computed, then the Poisson problem may be formulated by applying the divergence operator. A vector,  $v \in \mathbb{R}^{|\mathcal{O}|}$  (dimensionality equal to the number of octree nodes), is then computed so that each component is defined as shown in (11). The authors are not clear on the exact definition of the inner product used here, however it is likely defined as the standard “dot” product where  $F_o$  is evaluated at each of the spatial points represented by  $\nabla \cdot \vec{V}$ .

$$v_o = \left\langle \nabla \cdot \vec{V}, F_o \right\rangle \quad (11)$$

The matrix,  $L \in \mathbb{R}^{|\mathcal{O}| \times |\mathcal{O}|}$ , is computed so that each entry is defined in (12). Once again, the authors are not clear about the exact inner product definition. Due to the finite support of  $F$ , this matrix is sparse, allowing for efficient computations using it [9].

$$L_{o,o'} \equiv \left\langle \frac{\partial^2 F_o}{\partial x^2}, F_{o'} \right\rangle + \left\langle \frac{\partial^2 F_o}{\partial y^2}, F_{o'} \right\rangle + \left\langle \frac{\partial^2 F_o}{\partial z^2}, F_{o'} \right\rangle \quad (12)$$

The linear least-squares optimization problem, shown in (13), is solved and finally used to compute the estimate of the surface indicator function, shown in (14). The isosurface is then extracted by assigning boundary points when  $\tilde{\chi}$  is equal to the average value of  $\tilde{\chi}$  for each original point cloud sample [9].

$$x^* = \underset{x \in \mathbb{R}^{|\mathcal{O}|}}{\text{argmin}} \|Lx - v\|^2 \quad (13)$$

$$\tilde{\chi} = \sum_{o \in \mathcal{O}} x_o^* F_o \quad (14)$$

The authors adjust the theory described by the previous formulations to account for non-uniformly sampled data [9], however it is beyond the scope of this write-up.

### 3.3 Results

The results and output products of the Poisson approach are mostly favorable. In visual comparisons, the Poisson reconstructions achieve a sufficiently smooth surface, while also providing sufficient edge detail and producing minimal artifacts due to noise compared to other methods. Figure 5 highlights these results with the reconstruction of the “Stanford Bunny.” Figure 6 shows an advantage of the Poisson approach and also one of the drawbacks. The zoomed-in display of figure 6 shows that the Poisson reconstruction reproduces



Figure 5: Results of surface reconstruction for the “Stanford Bunny” dataset. The left shows reconstruction using Robust Cocone method, the center shows reconstruction from the Fast RBF method, and the right shows reconstruction using the Poisson method. (Graphics are from [9])

a smoother surface, with less “jaggy” artifacts than another approach. This makes the Poisson approach more desirable for our purposes, as we wish to remove non-smooth surface transitions. The drawback of the Poisson approach is that it does not behave appropriately when there are no samples (no information) in a region, as shown by the connected surface between the Buddha’s legs in figure 6.

The performance of the Poisson approach is discussed in terms of maximum memory usage and the number of triangles represented; lower memory usage and a larger number of triangles would indicate good performance relative to another approach. Only the Volumetric Range Image Processing (VRIP) and Multi-level Partition Unity (MPU) reconstruction algorithms produce more triangles in less memory than the Poisson approach, however the MPU was shown (graphically) to have a noisier reconstruction [9]. The authors also show that the algorithm is tailorable based on the desired output quality and processing memory limits, by showing the results of reconstructions with differing levels of octree depths. As expected the reconstructions with smaller octree depths produce noticeably less detailed reconstructions via fewer triangles, but are also computed faster and use less memory [9]. For each change in octree depth, the performance statistics are increased/decreased roughly by a factor of 5 [9].

The authors had previously developed a reconstruction algorithm that was based on the Fast Fourier Transform (FFT), and show that it is theoretically equivalent to the Poisson approach when the vector field  $\vec{V}$  is uniformly sampled [9]. However, when  $\vec{V}$  is sampled in a non-uniform fashion, the Poisson approach becomes more resilient to noise than the FFT method [9]. Additionally, the FFT approach uses about five times more memory than the Poisson approach [9].

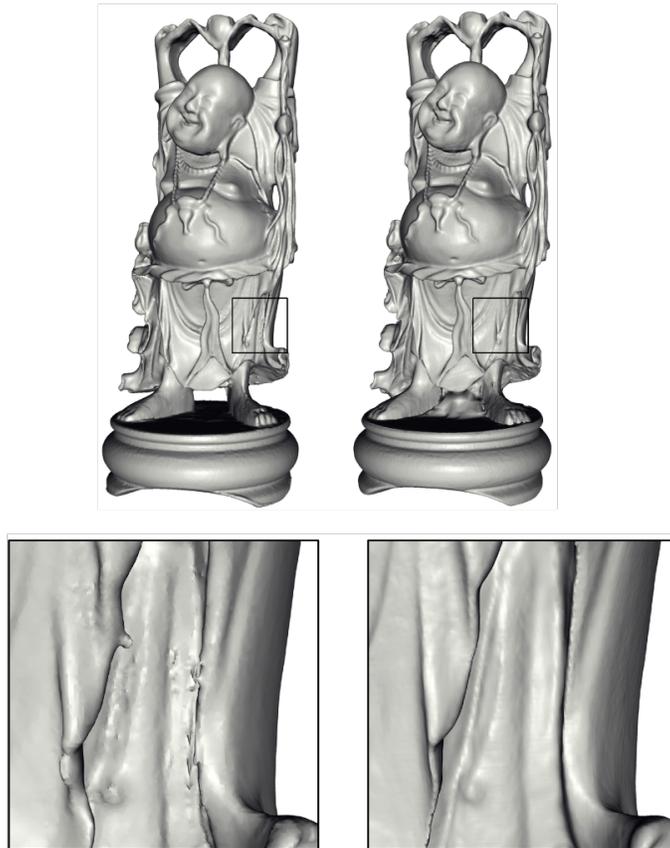


Figure 6: Results of surface reconstruction for the “Happy Buddha” dataset. The left column shows reconstruction using a competing method (VRIP) and the right column shows the Poisson method. The bottom row shows a zoomed-in display of the two methods; the Poisson method shows significantly less non-smooth artifacts. (Graphics are from [9])

### 3.4 Application to CIS Project

Using the basic extrapolated surface,  $\widetilde{\mathcal{M}}'_P$ , from (3), a point cloud may be generated so that the Poisson method may be used to derive a surface. For each vertex in  $\widetilde{\mathcal{M}}'_P$ , a surface normal may be calculated such that it is oriented towards the interior of the surface (either  $\mathcal{M}_P$  or  $\mathcal{M}_E$ ). This yields a cloud of oriented points,  $\mathcal{P}$ , defined in (15), where  $\mathbf{n}_{\mathbf{v},S}$  denotes the normal vector originating at vertex  $\mathbf{v}$  and oriented towards the interior of surface  $S$ .

$$\mathcal{P} = \{(\mathbf{v}, \mathbf{n}_{\mathbf{v},P}) \mid \mathbf{v} \in \mathbf{V}_P^{\text{known}}\} \cup \{(\mathbf{v}, \mathbf{n}_{\mathbf{v},E}) \mid \mathbf{v} \in \mathbf{V}_E^{\text{unknown}}\} \quad (15)$$

Figure 7 shows a hypothetical result that the Poisson method may output given a hypothetical  $\mathcal{P}$ . Since the use of the Poisson algorithm only requires that normal vectors be computed for each vertex of the basic extrapolated surface, it should be a straightforward method to apply to our data. The C++ source code for the Poisson method described in [9] is provided on the author’s website, making this a relatively low risk option to explore.

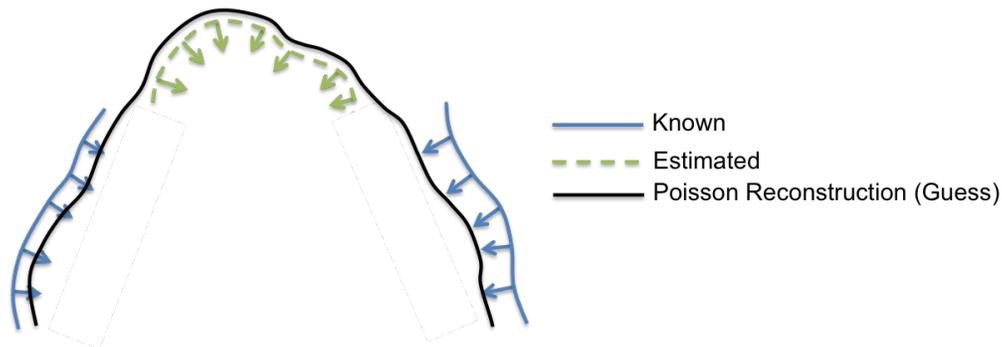


Figure 7: A potential result of using the Poisson method for reconstruction of our 2D “toy” surface originally shown in figure 2. The arrows indicate inward facing normal vectors at each vertex, the black line indicates a surface that the Poisson method could conceivably generate given the vertices and normals from the “known” surface and the “unknown” estimate surface.

## References

- [1] S. Benazzi, F. L. Bookstein, D. S. Strait, and G. W. Weber. A new {OH5} reconstruction with an assessment of its uncertainty. *Journal of Human Evolution*, 61(1):75–88, 2011.
- [2] S. Benazzi, L. Fiorenza, S. Kozakowski, and O. Kullmer. Comparing 3d virtual methods for hemimandibular body reconstruction. *The Anatomical Record: Advances in Integrative Anatomy and Evolutionary Biology*, 294(7):1116–1125, 2011.

- 
- [3] S. Benazzi and S. Senck. Comparing 3-dimensional virtual methods for reconstruction in craniomaxillofacial surgery. *Journal of Oral and Maxillofacial Surgery*, 69(4):1184–1194, 2011.
- [4] S. Benazzi, E. Stansfield, C. Milani, and G. Gruppioni. Geometric morphometric methods for three-dimensional virtual reconstruction of a fragmented cranium: the case of angelo poliziano. *International Journal of Legal Medicine*, 123(4):333–344, 2009.
- [5] F. Bookstein and W. Green. Edgewarp 3d, 2002. Available from: *brainmap.stat.washington.edu/edgewarp/*. Accessed April, 8, 2010.
- [6] G. Chintalapani. *Statistical atlases of bone anatomy and their applications*. Johns Hopkins University, 2010.
- [7] C. Cutting, R. Taylor, F. Bookstein, D. Khorramabadi, B. Haddad, A. Kalvin, H. Kim, and M. Noz. Computer aided planning and execution of craniofacial surgical procedures. In *Engineering in Medicine and Biology Society, 1992 14th Annual International Conference of the IEEE*, volume 3, pages 1069–1070. IEEE, 1992.
- [8] C. Herlin, M. Koppe, J.-L. Béziat, and A. Gleizal. Rapid prototyping in craniofacial surgery: using a positioning guide after zygomatic osteotomy—a case report. *Journal of Cranio-Maxillofacial Surgery*, 39(5):376–379, 2011.
- [9] M. Kazhdan, M. Bolitho, and H. Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, 2006.
- [10] S. Senck, M. Coquerelle, G. W. Weber, and S. Benazzi. Virtual reconstruction of very large skull defects featuring partly and completely missing midsagittal planes. *The Anatomical Record*, 296(5):745–758, 2013.