Surface interpolation with radial basis functions for medical imaging

Jonathan C. Carr, W. Richard Fright, *Member, IEEE*, Richard K. Beatson

Abstract—

Radial basis functions are presented as a practical solution to the problem of interpolating incomplete surfaces derived from three-dimensional (3-D) medical graphics. The specific application considered is the design of cranial implants for the repair of defects, usually holes, in the skull.

Radial basis functions impose few restrictions on the geometry of the interpolation centers and are suited to problems where the interpolation centers do not form a regular grid. However, their high computational requirements have previously limited their use to problems where the number of interpolation centers is small (< 300). Recently developed fast evaluation techniques have overcome these limitations and made radial basis interpolation a practical approach for larger data sets.

In this paper radial basis functions are fitted to depth-maps of the skull's surface, obtained from X-ray CT data using ray-tracing techniques. They are used to smoothly interpolate the surface of the skull across defect regions. The resulting mathematical description of the skull's surface can be evaluated at any desired resolution to be rendered on a graphics workstation, or to generate instructions for operating a CNC mill.

Keywords— **Surface interpolation, radial basis function approximation, medical graphics, titanium cranioplasty, CT imaging.**

I. Introduction

SURFACE interpolation has a role to play in computer assisted surgical planning through the reconstruction of missing surfaces and the prediction of soft tissue movement. The context for this work is the repair of defects in the skull with cranial implants. In this application medical graphics are used to reveal bone surfaces in X-ray CT data. By fitting a mathematical function which smoothly interpolates across holes in the surface, a hard plastic mold can be milled by a computer numerically controlled (CNC) mill. A titanium prosthesis can then be formed by pressing a flat titanium plate into the mold under high pressure in a hydraulic press. The desired characteristics of the interpolant are therefore that it is a smooth single-valued function, that it can interpolate across irregularly shaped holes and that it tends towards a flat plate far from the interpolation centers.

The paper is organized as follows. Section II backgrounds the cranioplasty application. Section III discusses various approaches to the interpolation problem. Section IV introduces radial basis functions and section V illustrates the application of radial basis functions to the problem of interpolating depth-maps of the skull's surface. Section VI presents examples of fitting surfaces to real data and section VII considers computational requirements. A discussion follows in section VIII.

The notation C^i is used throughout to denote a function which is continuous in its first i derivatives. The notation $\mathcal{O}(f(n))$ indicates that the work or storage (depending on the context) required by a process is bounded by a multiple of $f(n)$. For example, the work in solving an $n \times n$ system of linear equations, with no special structure, is $\mathcal{O}(n^3)$.

Jonathan Carr is with the Department of Electrical and Electronic Engineering at the University of Canterbury, Christchurch, New Zealand. Richard Fright is with the Department of Medical Physics and Bioengineering at Christchurch Hospital, Christchurch, New Zealand. Richard K. Beatson is with the Department of Mathematics and Statistics at the University of Canterbury, Christchurch, New Zealand.

II. Cranioplasty

Cranioplasty is the procedure of repairing defects, usually holes, in the skull with cranial implants. It is difficult to form a cranial prosthesis intra-operatively when a defect is large ($\geq$ 25 cm^2) or occurs in a region where the bones are thin or surround vital tissue, such as the orbit of the eye. In these regions the complex shape of an implant requires that it be formed presurgically.

Prefabrication of a cranial implant requires an accurate model of the defect area to ensure that a good fit is achieved at the time of surgery and the need for alteration is minimized. There is no standard process for the design and manufacture of prefabricated cranial implants, each treatment facility varying in the details of its method, and even then, each case being approached individually. However, most methods involve forming a model to which the implant, or a template for the implant, is manually fitted.

Linney *et al.* [13] have attempted to remove these manual aspects in the case where a defect is unilateral. They reflect the sound side of the anatomy on to the defective side of the body in software and use this as a model for the prosthesis, with the intention of providing symmetry to the shape of the repaired skull. However, this technique is applicable only where a symmetrical undamaged region of the skull exists. Even when this criterion is met, restoring the original skull shape is not always desirable if large cavities result between the brain and the implant. Furthermore, the skull is not always naturally symmetric.

In this paper the problem of designing a cranial implant is viewed as a surface interpolation problem where the surface of the skull is described by a mathematical function. The problem is to fit a function to the bone surface which interpolates across user-defined defect regions. The aim of this work is to create a one-step process for the design of cranial implants by integrating medical graphics and surface interpolation with CNC milling.

Depth-maps of the skull's surface, obtained from X-ray

CT data using ray-tracing techniques, have been used to construct models of cranial defects [23], [7], [13]. A 3-D CT data set in the form of a regular array of voxels is obtained by stacking a set of equally spaced CT slices taken through the head. A set of uniformly spaced parallel rays is cast into the CT data set from a user-determined direction. As a ray propagates through the data set an interpolation strategy is employed to estimate the CT number between voxel centers. A threshold corresponding to the Hounsfield number for bone (typically 500) is chosen as the criterion for terminating a ray. The depth at which each ray terminates is recorded to form a depth-map of the surface of interest.

Accurate determination of depth-map values is important to ensure good fit of the final implant. The Hounsfield threshold must therefore be chosen carefully. Partial voluming artifacts in the CT data mean that ray-tracing may produce errors in the vicinity of thin bones. These difficulties, associated with extracting depth-maps, are discussed elsewhere [8], [19]. In the following sections we address the problem of interpolating depth-maps. Suffice to say that most of our defects occur in regions where the skull is convex and the bone is thick. Applying a global threshold provides sufficient accuracy for our purposes in these cases.

From the resulting depth-map a 3-D model can be milled out of a block of epoxy-resin by a CNC mill. Often only a model of the region immediately surrounding a defect is required and the model can be milled from a single orientation of the workpiece.

A depth-map of a cranial defect can be displayed on a graphics workstation with the aid of shading. A rendered view of the skull surface allows a user to graphically identify a defect region in the skull. Many rendering techniques are found in the medical graphics literature [10], [14], [11], [21], [22]. Surface shading, determined by the gradient of the surface at each ray and the position of hypothesized light sources, can provide a realistic visualization of a 3-D surface on a two-dimensional (2-D) display device. In this work a Phong shading model [21] with a single light source is used to render depth-maps of skull defects. The user identifies a defect region graphically with the aid of a workstation's mouse.

Depth-maps are rendered at a resolution sufficient for the operator to identify defect regions and then subsampled for the purpose of fitting a surface. A depth-map mesh of 2 mm squares provides more than adequate sampling of the surface of the skull for interpolation in most cases. Although the resolution within a CT slice is usually 0.3–0.5 mm, the spacing between slices for cranial examinations is typically 2–3 mm. Consequently, a depth-map which resamples the data at a higher frequency will contain a ripple which originates from the interpolation procedure employed by the ray-tracer. Furthermore, such a depth-map will result in an unnecessarily large number of interpolation nodes, greatly increasing the storage and computational costs of solving the interpolation equations.

Figure 1 illustrates the cranial surface interpolation problem. Figure 1(a) is an example of a ray-traced view of the surface of the skull rendered from a CT data set. Figure 1(b) is a detailed view of the defect in Figure 1(a) where the region identified as being free of defect by the user is highlighted. Figure 1(c) illustrates the lower resolution depth-map corresponding to the detailed view of the defect while Figure 1(d) is the partial depth-map corresponding to the highlighted region in Figure 1(b). This is the depth-map to which an interpolating function is to be fitted. In summary, the original ray-traced depth data lies on a regular rectangular grid. The user then interactively identifies a defect region of irregular shape, and the corresponding data points are discarded. This results in a final depth data set to be interpolated which is of the form of a finite grid, with a hole, or holes.

III. Surface interpolation

The depth-map is regarded as a single-valued function of two variables (2-D), $f : \mathbb{R}^2 \to \mathbb{R}$, where the interpolation nodes at which the value of the function is known do not generally form a complete regular grid. The problem is to approximate f with a smooth function which interpolates the given data and provides at least C^1 continuity in order to avoid creases in the restored surface.

One approach to this problem is to fit a polynomial to the data. However, an invertible system which uniquely defines the interpolant is not guaranteed for all positions of the interpolation nodes. Moreover, such a polynomial interpolant will typically display spurious bumps and wiggles. An alternative is to fit piecewise polynomials. These require that $\mathbb{R}^2$ be divided into suitable patches, typically rectangles and triangles. One employs different polynomials on the different patches, and the pieces are joined in ways that provide continuity of the prescribed low order derivatives. Examples of such methods are tensor product spline interpolants, parametric spline interpolants, C^1 surfaces made up of Clough-Tocher macro elements over triangles, etc. Surfaces of these types are widely used in CAD [5] and in Finite Element methods for solving partial differential equations. However, they are not well suited to scattered data interpolation problems, particularly when there are large data-free regions.

The radial basis approximation method offers several advantages over piecewise polynomial interpolants. The geometry of the known points is not restricted to a regular grid and there is no need to define a mesh of patches. Also, the resulting system of linear equations is guaranteed to be invertible under very mild conditions. Finally, polyharmonic radial basis functions have variational characterizations which make them eminently suited to interpolation of scattered data, even with large data-free regions.

IV. Radial basis function approximation

We begin by introducing radial basis function approximation generally before considering the specific case of approximating a function of two variables. Consider $f : \mathbb{R}^d \to \mathbb{R}$ a real valued function of d variables that is to be approximated by $s : \mathbb{R}^d \to \mathbb{R}$, given the values $\{f(\mathbf{x}_i) : i = 1, 2, \ldots, n\}$, where $\{\mathbf{x}_i : i = 1, 2, \ldots, n\}$ is a set

of distinct points in $\mathbb{R}^d$ called the nodes of interpolation.

We will consider approximations of the form

$$s(\mathbf{x}) = p_m(\mathbf{x}) + \sum_{i=1}^{n} \lambda_i \phi(||\mathbf{x}-\mathbf{x}_i||), \qquad \mathbf{x} \in \mathbb{R}^d, \quad \lambda_i \in \mathbb{R}, \tag{1}$$

where p_m is a low degree polynomial, or is not present, $||\cdot||$ denotes the Euclidean norm and ϕ is a fixed function from $\mathbb{R}^+$ to $\mathbb{R}$. Thus the radial basis function s is a linear combination of translates of the single radially symmetric function $\phi(||\cdot||)$, plus a low degree polynomial. We will denote by π_m^d the space of all polynomials of degree at most m in d variables. Then the coefficients, λ_i, of the approximation s are determined by requiring that s satisfy the interpolation conditions

$$s(\mathbf{x}_j) = f(\mathbf{x}_j), \qquad j = 1, 2, \ldots, n, \tag{2}$$

together with the side conditions

$$\sum_{j=1}^{n} \lambda_j q(\mathbf{x}_j) = 0, \qquad \text{for all } q \in \pi_m^d. \tag{3}$$

Some examples of popular choices of ϕ and the name of the corresponding radial function are given below.

$$\left.\begin{array}{ll} \phi(r) = r & \text{(linear)} \\ \phi(r) = r^2 \log r & \text{(thin-plate spline)} \\ \phi(r) = e^{-ar^2} & \text{(Gaussian)} \\ \phi(r) = (r^2 + c^2)^{1/2} & \text{(multiquadratic)} \end{array}\right\}, \qquad r \geq 0, \tag{4}$$

where a and c are positive constants.

Some typical conditions on the nodes under which the interpolation conditions (2) and (3) uniquely specify the radial basis function (1) are given in Table I.

In the case when the polynomial term is absent from the radial basis function (1) there are no corresponding side conditions (3). Also "not coplanar" in this context means that the nodes do not all lie in a single hyperplane, or equivalently that no linear polynomial in d-variables vanishes at all of the nodes. The surveys of Powell [18] and Light [12] are excellent references for these and other properties of radial basis functions.

In this paper we are particularly concerned with 2-D (depth-map) data and will consider linear and thin-plate spline interpolants. Then (2) and (3) imply that the coefficients of the radial basis function and the polynomial $p_1(\mathbf{x})$ can be found by solving the linear system

$$\begin{bmatrix} A & Q \\ Q^T & 0 \end{bmatrix} \begin{bmatrix} \lambda \\ c \end{bmatrix} = \begin{bmatrix} f \\ 0 \end{bmatrix}. \tag{5}$$

where

$$A = (a_{ij}) = (\phi(||\mathbf{x}_i - \mathbf{x}_j||)), \tag{6}$$

$$Q = \begin{bmatrix} 1 & x_1 & y_1 \\ 1 & x_2 & y_2 \\ \vdots & \vdots & \vdots \\ 1 & x_n & y_n \end{bmatrix}, \tag{7}$$

$$\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_n)^T, \tag{8}$$

$$c = (c_0, c_1, c_2)^T, \tag{9}$$

$$p_1(\mathbf{x}) = c_0 + c_1 x + c_2 y \tag{10}$$

and

$$f = (f_1, f_2, \ldots, f_n)^T \tag{11}$$

The thin-plate, or 2-dimensional biharmonic spline we consider models the deflection of an infinite thin plate [6]. While the linear radial basis function will interpolate the data, the thin-plate spline is more attractive since it also provides C^1 continuity and minimizes the energy functional

$$E(s) = \int_{\mathbb{R}^2} \left(\frac{\partial^2 s}{\partial x^2}\right)^2 + 2\left(\frac{\partial^2 s}{\partial x \partial y}\right)^2 + \left(\frac{\partial^2 s}{\partial y^2}\right)^2 dxdy \tag{12}$$

over all interpolants for which the energy functional (12) is well defined. In this sense the thin-plate spline is the *smoothest* interpolator of f. Higher order polyharmoic splines achieve continuity of higher derivatives. Perrin *et al.* [17], [16] have used the tri-harmonic thin-plate spline ($\phi(\mathbf{x}) = ||\mathbf{x}||^4 \log ||\mathbf{x}||$) in electroencephalography to interpolate with C^2 continuity distributions of potential and current scalp density measured on the surface of the head.

V. Radial basis interpolation of depth-maps

In this section we discuss the application of radial basis functions to the interpolation of cranial depth-maps. The first task is to calculate the interpolant, or rather its coefficients. Our current solution procedure is as follows. The nodes are first scaled uniformly in x and y, and shifted so that the new nodes lie in the unit square. The variational characterization of the thin-plate spline interpolant means that it is preserved under this transformation. Then the interpolation problem corresponding to the transformed data is solved using the double precision diagonal pivoting method for symmetric indefinite systems from Lapack [9]. The implementation invokes iterative refinement to improve the accuracy of an initial solution if warranted by the magnitude of the residual. In our experiments the infinity norm condition number of the system was typically around 6×10^6, the initial relative residual was around 2×10^{-13}, and the initial relative error was around 1.5×10^{-11}, and our aim was to obtain a relative error (against the "true" thin-plate interpolant) of 10^{-4}. Overall, we found that more than adequate accuracy in the fitted surface was readily obtained by solving in double precision, even without iterative refinement.

The natural criterion for assessing a restored cranial surface is how closely it matches the original surface prior to the introduction of the defect. However, the original surface is rarely known and does not exist in the case of a congenital defect. Therefore, we simulated the performance of the method by seeing how well it restored artificial holes in complete skulls. The interpolant fitted to the incomplete depth-map was then compared with the original surface.

Figure 2 illustrates the test data employed in the following interpolation example in which an artificial hole is

introduced into an initially defect free depth-map of a cranial surface. Figure 2(a) is a rendered view of a CT data set which is free from defect. Figure 2(b) is an enlarged view of part of Figure 2(a) with an artificial hole superimposed. Figure 2(c) is the corresponding depth-map. A thin-plate spline radial basis interpolant was fitted to 683 points in Figure 2(c). Figure 3 compares the fitted surface with the original data. Figure 3(a) is the complete depth-map of the original data, while Figure 3(b) is the depth-map of the interpolated surface. Figures 3(c) and 3(d) are the rendered views of the respective depth-maps in 3(a) and 3(b).

In Figure 4(a) the difference between the fitted thin-plate spline surface and the original is displayed. The scales are in millimetres. The maximum error for the thin-plate spline interpolant is 0.60 mm. Figure 4(b) is the corresponding comparison between the original surface and a linear radial basis interpolant. The largest error from the original is 0.56 mm. The differences between the two interpolants and the original surface are dependent upon the surface data chosen in Figure 2. This example is nevertheless typical of many of the smooth, low curvature surfaces which are to be reconstructed. Although the magnitude of the error in both fitted surfaces is similar, the variational characterization of the thin-plate spline, and its C^1 continuity make it the natural choice for cranial implant design.

A. Surface evaluation

Once the linear system is solved, s can be evaluated at any resolution, anywhere on $\mathbb{R}^2$. Figure 5 shows the thin-plate surface from Figure 3 evaluated outside the convex hull of the interpolation nodes. Figure 5(a) is a rendered view taken from above the surface, while Figure 5(b) is the corresponding depth-map. Apparent is the planar component of the solution which is described by the low degree polynomial term in (1).

At Christchurch Hospital, cranial implants are made from flat titanium plate which is pressed into a mold in a hydraulic press. The interpolant determines the shape of the mold. The tendency of the interpolant to a flat plane outside the interpolation nodes is an ideal characteristic in this application. The smooth convergence of the mold surface to a flat plane reduces the likelihood of thinning or tearing of the metal at the periphery as it is drawn in to the mold.

It is desirable to orient the mold surface so that the planar component is flat with respect to the plane of the press. The normal of the planar component of the fitted surface (Figure 6) is derived directly from the coefficients c_1 and c_2 of the polynomial term $p_1(\mathbf{x})$,

$$\mathbf{n} = (-c_1, -c_2, 1). \tag{13}$$

The rotation angles about the x and y axes, denoted by θ and ψ respectively, required to re-orient the interpolated depth-map so that the normal aligns with the z axis (viewing axis) are given by

$$\theta = \tan^{-1}\left(\frac{-c_2}{\sqrt{1+c_1^2}}\right) \qquad \text{and} \qquad \psi = \tan^{-1} c_1 \tag{14}$$

In practice the plane of the fitted surface is usually close enough to the (x-y) plane of the press to be used without risk of tearing the titanium plate or press membrane. If it is not, then ray-tracing is repeated with an improved viewpoint determined by the orientation angles θ and ψ. A new surface is fitted to the depth-map corresponding to the improved viewpoint.

VI. Examples

Two examples are presented to illustrate the application of radial basis interpolation to real CT data of skull defects requiring repair.

The first example involves the repair of a very large (150 cm^2) hole in a basically convex region of the skull. Figure 7(a) is a rendered view of the CT data set which consists of 38 slices at 3 mm spacing with a pixel size of 0.47 mm. The rendered threshold corresponds to a Hounsfield number of 500. Figure 7(b) is a detailed view of the defect with the support region identified by the user highlighted. Figure 7(c) is a mesh-plot of the support shown in Figure 7(b). The mesh resolution is approximately 1 mm. This is the data to which a surface is fitted. Figure 8(a) is a rendered view from directly above the thin-plate spline surface fitted to the data in Figure 7(c). Figure 8(b) is a mesh plot of the fitted surface viewed from the same direction as Figure 7(c).

The mold produced from the interpolated surface data is shown in Figure 9(a) next to the finished titanium plate. CNC instructions were generated from the surface data to machine the mold from a very hard epoxy-resin. In Figure 9(b) the finished plate is shown beside a model of the defect area which has been machined from hard plastic using depth-map data shown in Figure 7.

The second example involves the repair of a hole close to the orbital margin and other regions of high curvature such as the zygomatic arch. Consequently, a smooth concave mold could not be produced from a mirror-image depth-map of the unaffected side. Figure 10(a) is a rendered view of the CT data set which consists of 47 slices at 3 mm spacing with a pixel size of 0.89 mm. The rendered threshold corresponds to a Hounsfield number of 500. Figure 10(b) is a detailed view of the defect with the support region identified by the user highlighted. Figure 10(c) is a mesh-plot of the support shown in Figure 10(b). The mesh resolution is approximately 1 mm. Figure 11(a) is a rendered view from directly above the thin-plate spline surface fitted to the data in Figure 10(c). Figure 11(b) is a mesh plot of the fitted surface viewed from the same direction as Figure 10(c). The surface continues smoothly outside the support of the defect region, which is essential for pressing titanium into the mold without tearing.

VII. Computational considerations

Radial basis functions have previously been proposed as interpolants for 3-D data where the interpolation centers do not lie on a regular grid, by Nielson [15]. However, Nielson notes that conventional methods limit the application of radial basis functions to data sets of moderate size (n=300 to 500). Cutting *et al.* [4] have used the thin plate spline radial basis function to warp a wire mesh template of a skull on to cranial surface data extracted from CT and MRI data. This is done by manually identifying 44 homologous landmarks in the two data sets. Each landmark in the template dataset forms an interpolation node and has a warping vector associated with it which maps it on to its pair in the skull dataset. Radial basis functions have also been used in electroencephalography to interpolate potentials measured on the surface of the head to produce topographical maps [17], [16], [20]. However, the number of electrodes (interpolation nodes) used in this application is again small, typically fewer than 40.

Generating a radial basis interpolant involves solving a linear system like (5) for $(\lambda|c)^T$. Since the ϕ's we use are typically non-zero, and growing away from zero, the matrices of these systems display none of the special structures such as sparsity, bandedness, positive definiteness commonly exploited to enable fast solution of linear systems. Thus solution of these systems by direct, or even simple iterative methods, requires $\mathcal{O}(n^3)$ operations and $\mathcal{O}(n^2)$ storage. Taking advantage of the symmetry of the system, will halve the usual operations count for Gaussian elimination but even symmetric solvers will still be impractical when n is large. These requirements used to be a considerable obstacle to the use of radial basis functions when n is large. Fortunately there are already known methods which allow the solution of the thin-plate spline interpolation equations with $\mathcal{O}(n)$ storage and in $\mathcal{O}(n^2 \log n)$ operations (see [3] and the references there). Furthermore, improved methods under development by M.J.D. Powell and others promise solution in $\mathcal{O}(n \log n)$ operations.

The current calculations however involve moderately sized data sets and conventional techniques are satisfactory for solving the corresponding linear systems. The 697 node interpolation problem underlying Figure 8 required approximately 40.5 seconds of CPU time to solve on a SUN Sparc 2, while the 578 node interpolation problem underlying Figure 11 required approximately 23.5 seconds of CPU time to solve on the same machine.

Once a radial basis function has been fitted to the data we face the problem of evaluating it. Due to the behavior of ϕ, the work required for direct evaluation of s at a single extra point, x, is proportional to the number of nodes n. Thus, the amount of computation required for direct evaluation of s on a fine mesh becomes significant, even for moderate node counts n. However, recently developed fast evaluation methods (see Beatson and Newsam [1] and Beatson and Light [2]) exploit hierarchical data structures and series expansions, to reduce the incremental cost of one evaluation of a thin-plate spline to within precision ϵ, to $\mathcal{O}(1 + |\log \epsilon|)$ operations. Of course the setup time before evaluation does depend on the node count, nodal positions, the values of the spline coefficients, and the desired accuracy.

Spline of	205×205 evaluation grid		512×512 evaluation grid	
	direct	fast	direct	fast
Figure 8	176s	4.34s	1103s	9.53s
Figure 11	146s	4.35s	910s	9.33s

TABLE II

Comparison between fast and direct evaluation times of thin-plate spline interpolants.

Table II shows times in seconds for evaluation tasks associated with the figures of this paper. The improvement to be had by using a fast evaluation code is dramatic. All computations were carried out in double precision on a SUN Sparc 2 machine, and the fast evaluator was required to evaluate the fitted surface with an infinity norm relative error, ϵ, of less than 10^{-4}.

VIII. Discussion

This paper has demonstrated the application of radial basis functions to cranial implant design. Radial basis function approximation is suited to this type of problem due to the few constraints placed upon the geometry of the nodes of interpolation. Few alternative techniques are available when the data do not lie on a complete regular grid. The thin-plate spline is a particularly appropriate choice since it fits a smooth C^1 surface to the data and is the smoothest C^1 interpolant in the sense that it minimizes the energy functional (12). The property that the interpolant tends to a flat plate far from the interpolation centers is ideally suited to the manufacture of titanium prostheses in a hydraulic press.

Accurate depth-map values are required because the fitted surface passes through all the specified depths. Although ray-tracing can be in error, for example, in the vicinity of thin bone, a strength of our approach is that the user can arbitrarily exclude points from the depth-map in regions where they do not have confidence in the ray-traced values.

Some complex repairs can not be described by a single-valued function and consequently can not be manufactured by pressing titanium plate in a mold. In these cases we resort to traditional hand forming techniques and sometimes use multiple-part prostheses, where the separate parts are formed and butt welded together [7]. Most of our defects, however, occur over convex regions of the skull.

A development of our current software will provide the surgeon with an interactive facility for sculpting the surface via adjustment of the depth at user specified "control points". For example, restoring the natural shape of the skull across a large hole may lead to a large cavity forming between the brain and the prosthesis. The user may want to indent an initial depth-map in the defect region to

avoid such cavities. Similarly, the surface data around the periphery of the hole may not be sufficient to restore the appropriate shape to the skull. Ultimately, the 'correct' surface is that which is intended by the surgeon. The thin plate spline interpolant provides an appropriate surface in most cases, but sometimes may require adjustment.

Prerequisites for the new interactive software are an efficient means for updating an initial matrix factorization as interpolation nodes are added and deleted, and also a fast means of recalculating the resulting surface. Fortunately, algorithms for both tasks are available.

One suitable method for the first task is the matrix updating method which is amongst the things discussed in Beatson and Powell [3]. They first change the interpolation problem into an equivalent problem, with symmetric positive definite matrix, and then consider means for updating the Cholesky factorization of this matrix as the nodes of interpolation change. In the current context this would enable solving the system with a few extra nodes of interpolation at the cost of $\mathcal{O}(n^2)$ extra operations, rather than the $\mathcal{O}(n^3)$ operations required if one started from scratch. A suitable method for the second task is the fast evaluation code described in a section VII.

Acknowledgments

The authors wish to thank Michael Parker who made the titanium plates, neurosurgeons Dr's McFarlane and Bonkowski, and the reviewers for their helpful comments. The authors also wish to acknowledge the late John Hinton who pioneered titanium cranioplasty at Christchurch Hospital and provided the motivation for this work. This work was supported by Telecom New Zealand in the form of a post-graduate scholarship to Jonathan Carr and by Lottery Health Research for the purchase of hardware.

References

[1] R. Beatson and G. Newsam. Fast evaluation of radial basis functions: I. *Computers and Mathematics with Applications*, 24(12):7–19, 1992.

[2] R.K. Beatson and W.A. Light. Fast evaluation of radial basis functions: Methods for 2-dimensional polyharmonic splines. Technical Report 119, Mathematics department, University of Canterbury, Christchurch, New Zealand, December 1994.

[3] R.K. Beatson and M.J.D. Powell. An iterative method for thin-plate spline interpolation that employs approximations to the Lagrange functions. In D.F. Griffiths and G.A. Watson, editors, *Numerical Analysis 1993*, pages 17–39. Longman Scientific, 1994.

[4] C. B. Cutting, F. L. Bookstein, B. Haddad, D. Dean, and D. Kim. A spline-based approach for averaging three-dimensional curves and surfaces. In J. Wilson and D. Wilson, editors, *Mathematical Methods in Medical Imaging II, SPIE Proceedings*, volume 2035, pages 29–42, 1993.

[5] J. Foley and A. Van Dam. *Fundamentals of Interactive Computer Graphics*. Addison-Wesley Publishing Co., 1982.

[6] R. L. Harder. Interpolation using surface splines. *Journal of Aircraft*, 9(2):189–191, February 1972.

[7] J. W. Hinton, M. R. MacFarlane, and G. B. Blake. Repairing people - titanium in reconstructive surgery. *Shadows: New Zealand Journal of Medical Radiation Technology*, 35(1):27–32, March 1992.

[8] Arie Kaufman. *Volume Visualisation*. IEEE Computer Society Press, Los Altimos, CA, 1991.

[9] LAPACK. *LAPACK users guide*. Society for Industrial and Applied Mathematics, Philadelphia, 1992.

[10] M. Levoy. Display of surfaces from volume data. *IEEE Computer Graphics and Applications*, pages 29–37, May 1988.

[11] M. Levoy. A hybrid ray tracer for rendering polygon and volume data. *IEEE Computer Graphics and Applications*, pages 33–40, March 1990.

[12] W.A. Light. Some aspects of radial basis function approximation. In S.P. Singh, editor, *Approximation Theory, Spline Functions and Applications*, pages 163–190. Kluwer Academic Publishers (Dortrecht), 1992.

[13] A. D. Linney, A. C. Tan, J. Richards, S. Gardener, S. Grindrod, and J. P. Moss. Three-dimensional visualisation of data on human anatomy: Diagnosis and surgical planning. *Journal of Audiovisual Media in Medicine*, 16:4–10, 1993.

[14] D. Ney, E. Fishman, D. Magid, and R. Drebin. Volumetric rendering of computed tomography data: Principles and techniques. *IEEE Computer Graphics and Applications*, pages 24–31, March 1990.

[15] G. M. Nielson. Scattered data modelling. *IEEE Computer Graphics and Applications*, pages 60–70, January 1993.

[16] F. Perrin, O. Bertrand, and J. Pernier. Scalp current density mapping: value and estimation from potential data. *IEEE Transactions on Biomedical Engineering*, BME-34(4):283–288, 1987b.

[17] F. Perrin, J. Pernier, O. Bertrand, M. H. Giard, and J. F. Echallier. Mapping of scalp potentials by surface spline interpolation. *Electroencephalography and Clinical Neurophysiology*, 66:75–81, 1987a.

[18] M.J.D. Powell. The theory of radial basis function approximation in 1990. In W.A. Light, editor, *Advances in Numerical Analysis II: Wavelets, subdivision algorithms and radial functions*, pages 105–210. Oxford University Press, Oxford, UK,, 1992.

[19] S. P. Raya and J. K. Udupa. Shape-based interpolation of multidimensional objects. *IEEE Transactions on Medical Imaging*, 9(1):32–42, March 1990.

[20] B. L. Satherley. *Zero-based ensemble deconvolution and EEG spectral topography*. PhD thesis, Department of Electrical and Electronic Engineering, University of Canterbury, Christchurch, New Zealand., 1994.

[21] U. Tiede, K. Heinz, M. Bomans, A. Pommert, M. Riemer, and G. Wiebecke. Investigation of medical 3D-rendering algorithms. *IEEE Computer Graphics and Applications*, pages 41–53, March 1990.

[22] J. K. Udupa and D. Odhner. Shell rendering. *IEEE Computer Graphics and Applications*, pages 58–67, November 1993.

[23] M. C.. Van Putten and S. Yamada. Alloplastic cranial implants made from computed tomographic scan-generated casts. *Journal of Prosthetic Dentistry*, 68(1):103–108, July 1992.

ϕ	spatial dimension d	polynomial degree m	Restriction on nodes
linear	any	1	nodes not coplanar
thin-plate	2	1	nodes not colinear
Gaussian	any	absent	none
multiquadric	any	absent	none

TABLE I

CONDITIONS IMPOSED ON NODES FOR VARIOUS RADIAL BASIS INTERPOLANTS.

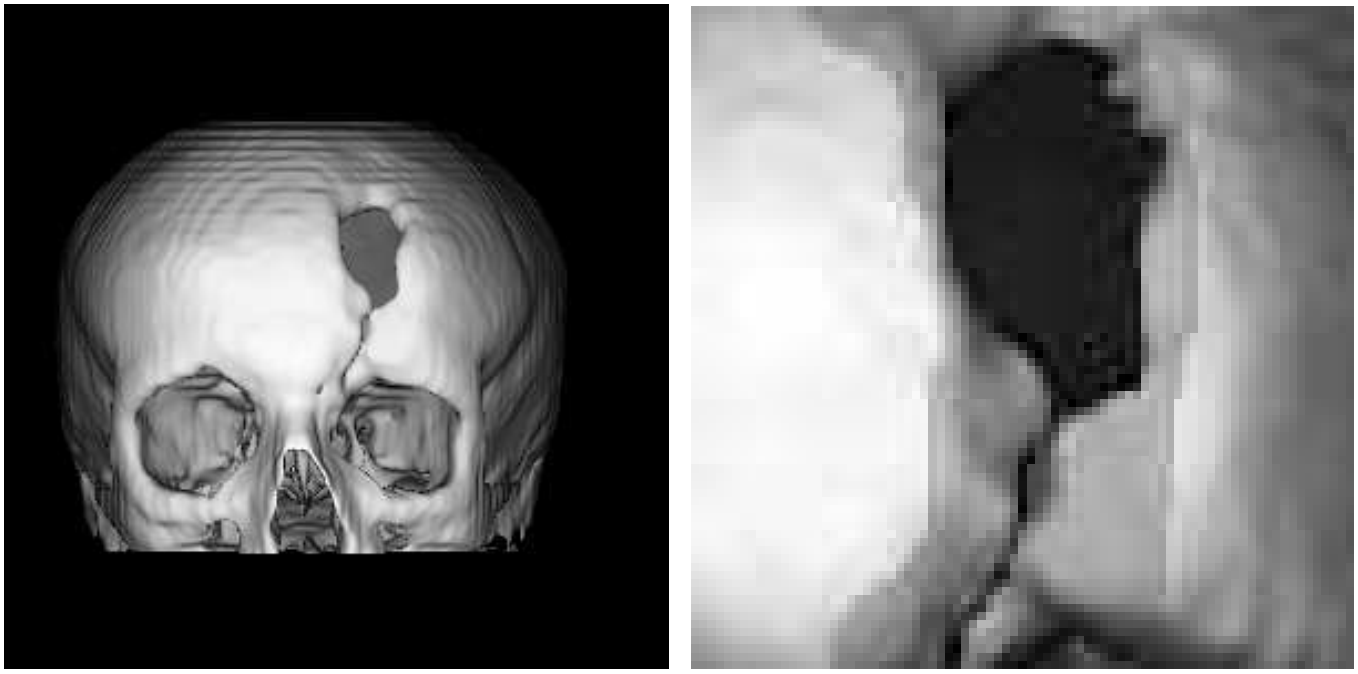

(a) (b)

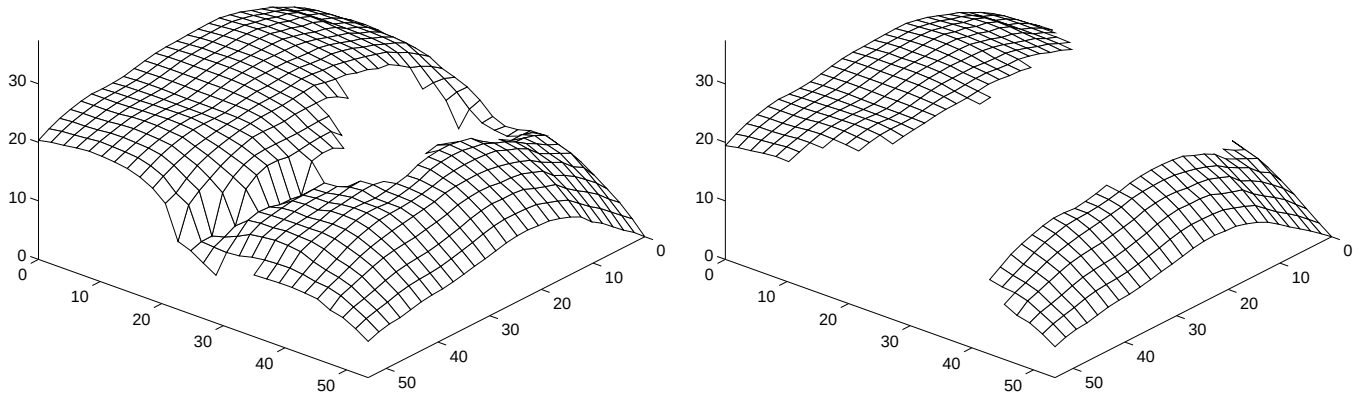

(c) (d)

Fig. 1. Extraction of a depth-map of the skull from CT data. (a) Rendered view of the bone surface from a CT data set, (b) detailed view of the defect with a user-defined region of defect-free bone highlighted, (c) depth-map of the rendered view of the defect in (b), (d) partial depth-map corresponding to the highlighted region in (b).

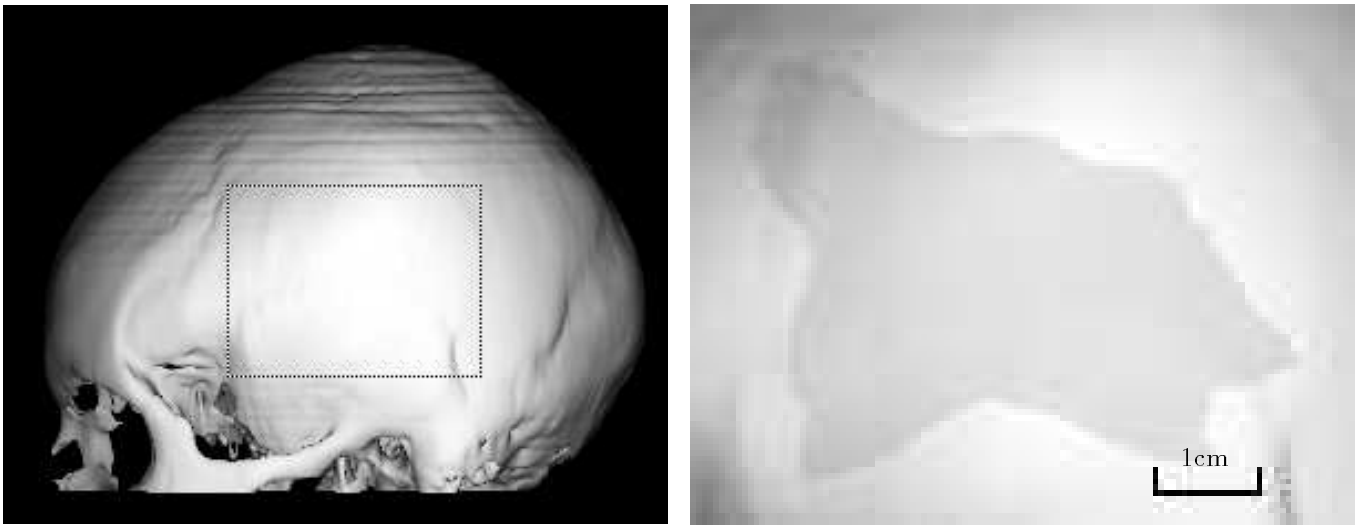

(a) (b)

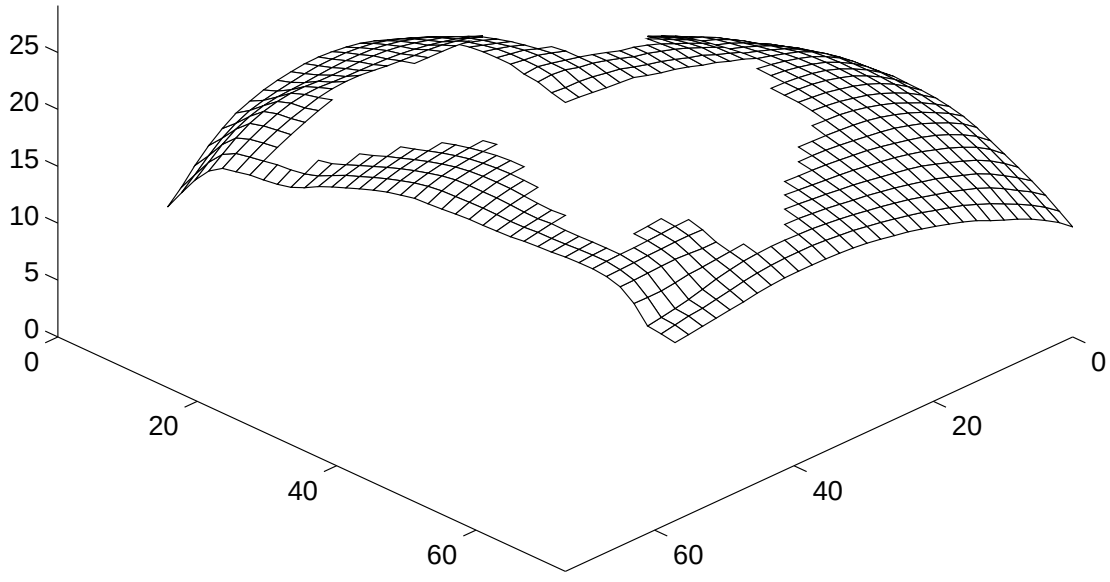

(c)

Fig. 2. Test data for surface interpolation. (a) Original CT data set. (b) Detail of the test surface with an artificial defect superimposed. (c) Depth-map corresponding to (b).

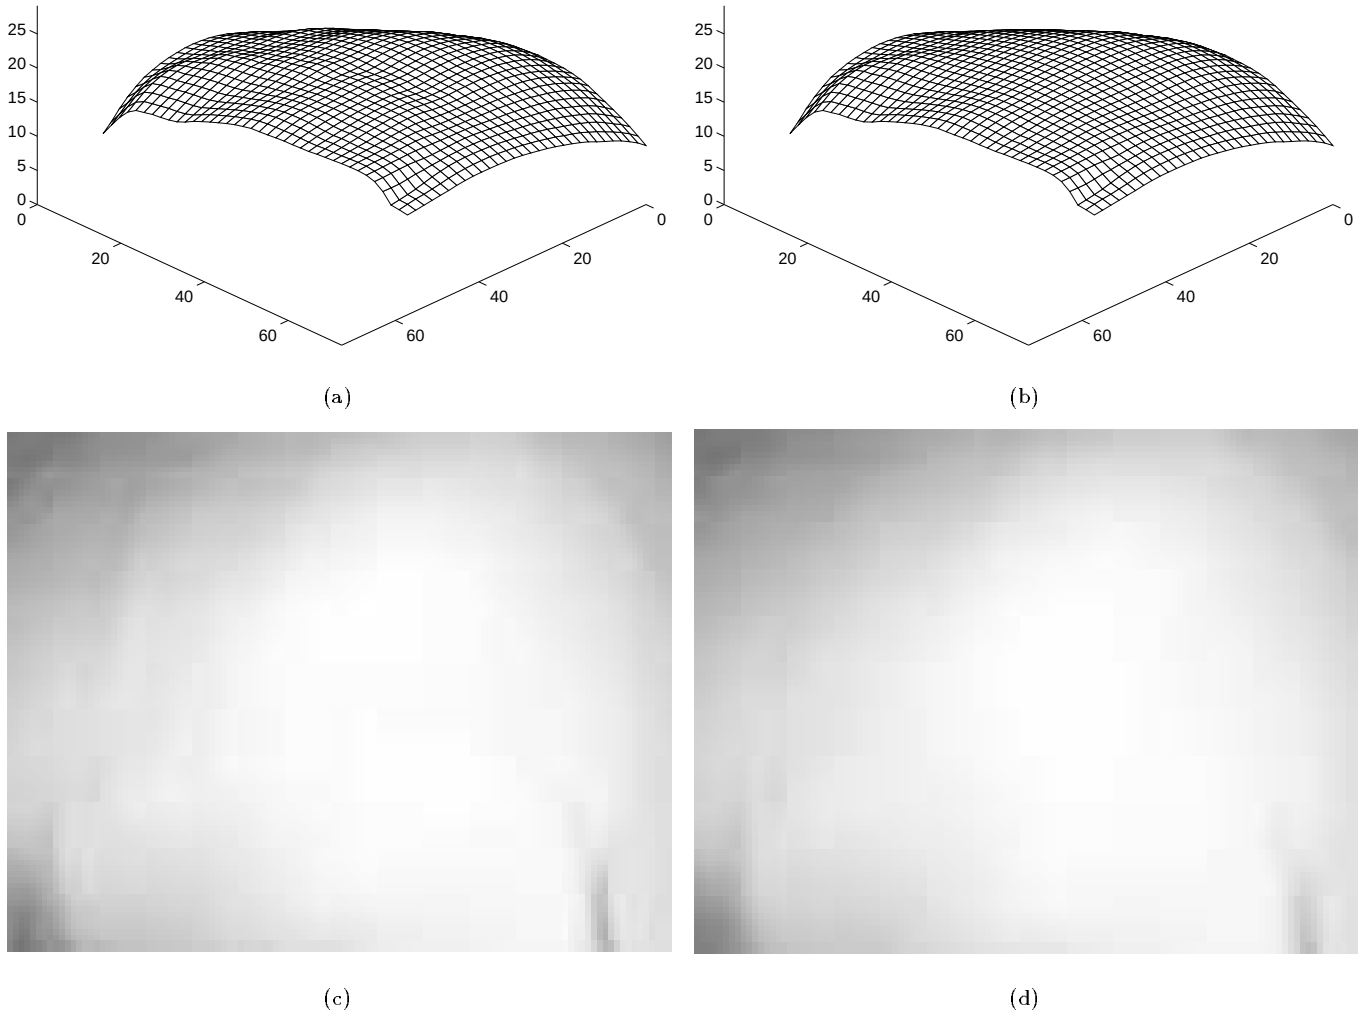

Fig. 3. Comparison of the interpolated surface with the original from Figure 2. (a) Depth-map for original data. (b) Depth-map for thin-plate spline surface. (c) Rendered view of original data. (d) Rendered view of thin-plate spline surface.

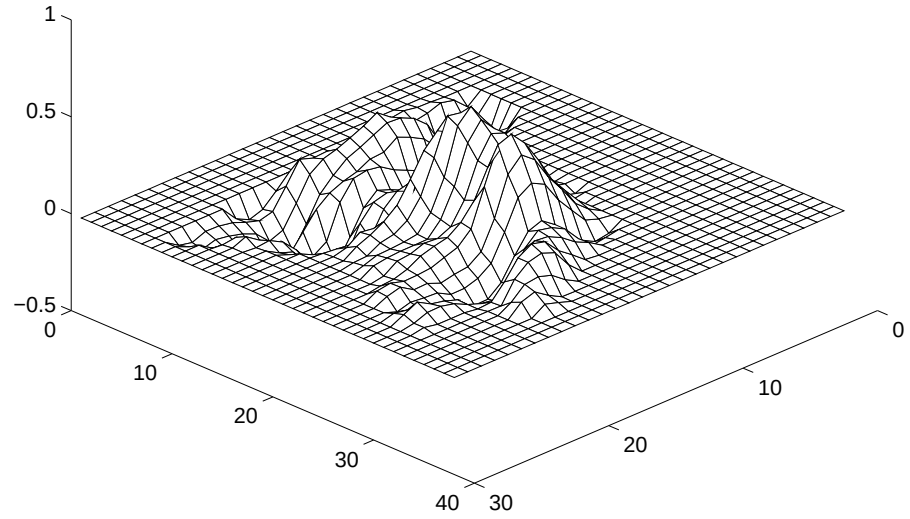

(a)

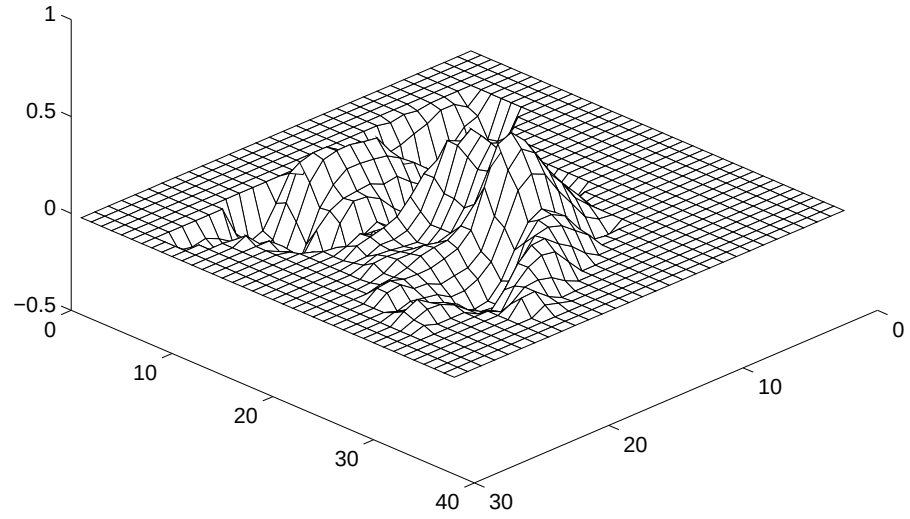

(b)

Fig. 4. Comparison of surface variation between the thin-plate spline and linear radial bases; the variation is in millimetres. (a) Difference between the thin-plate spline surface and original data. (b) Difference between the linear surface and original data.

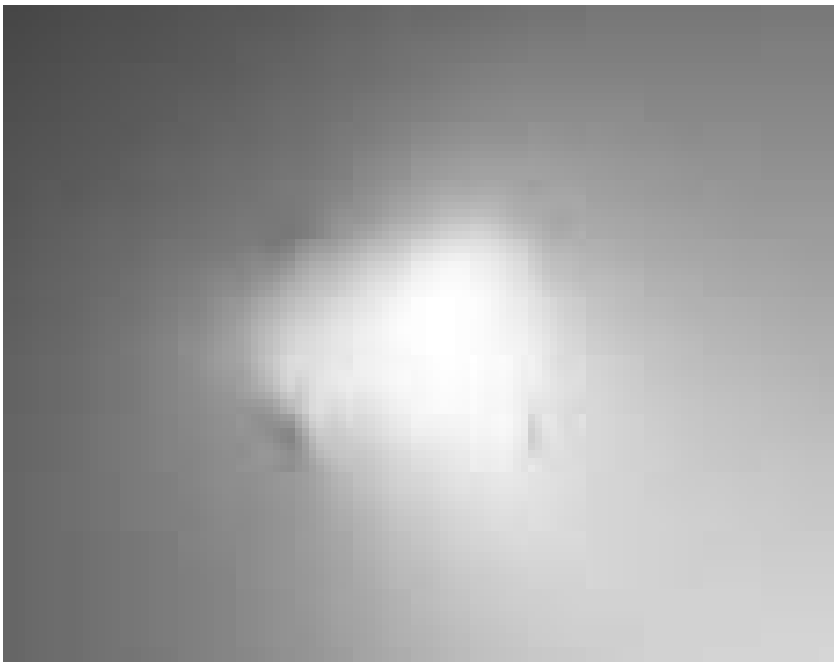

(a)

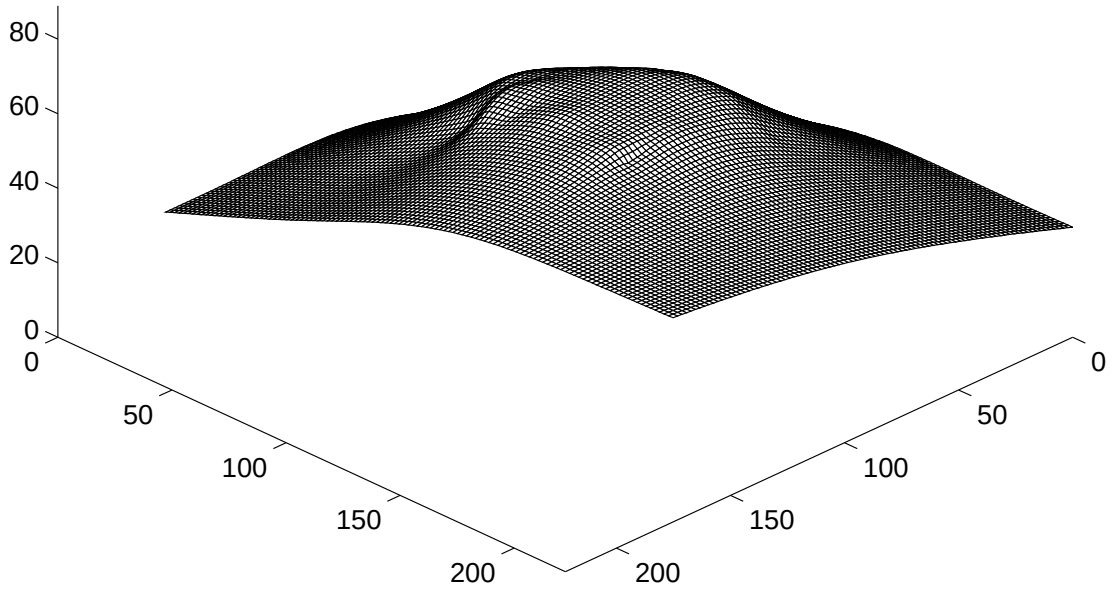

(b)

Fig. 5. Behavior of interpolant in Figure 3(b) far from the data points. (a) Rendered view of fitted thin-plate surface. (b) Mesh plot of thin-plate surface.

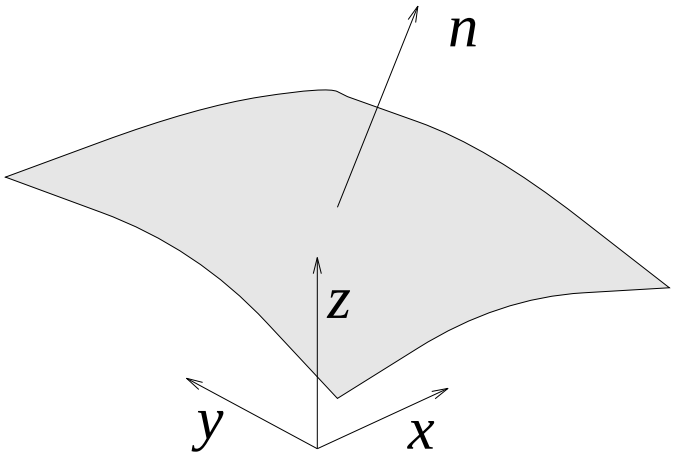

Fig. 6. Depth-map coordinate system.

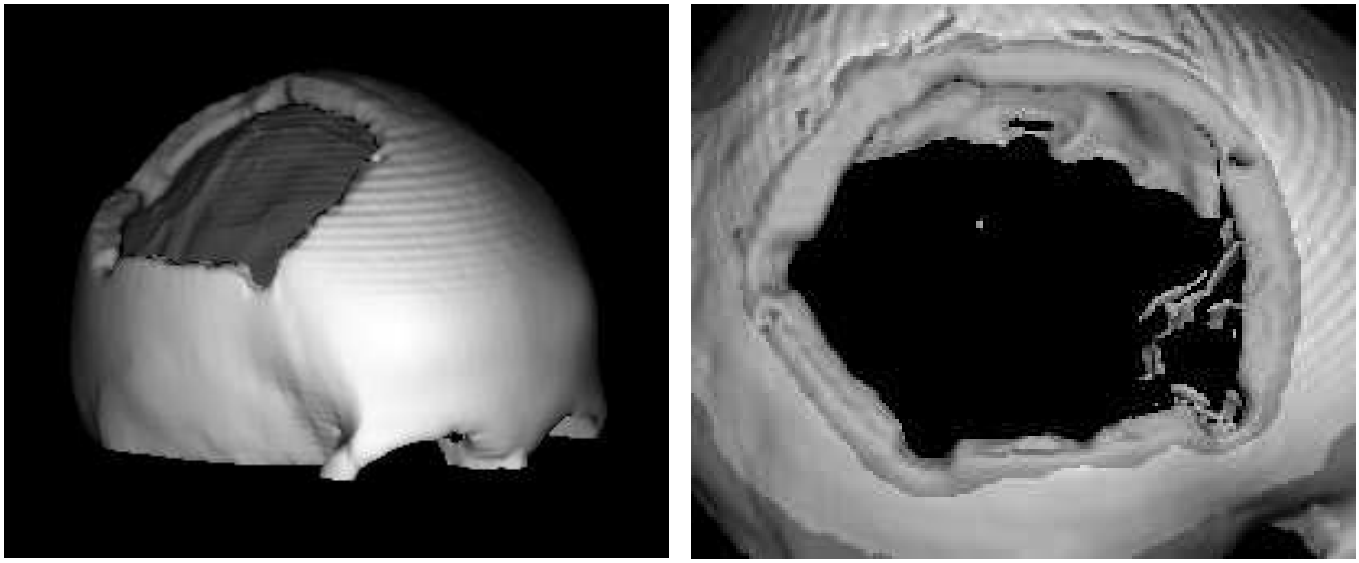

(a) (b)

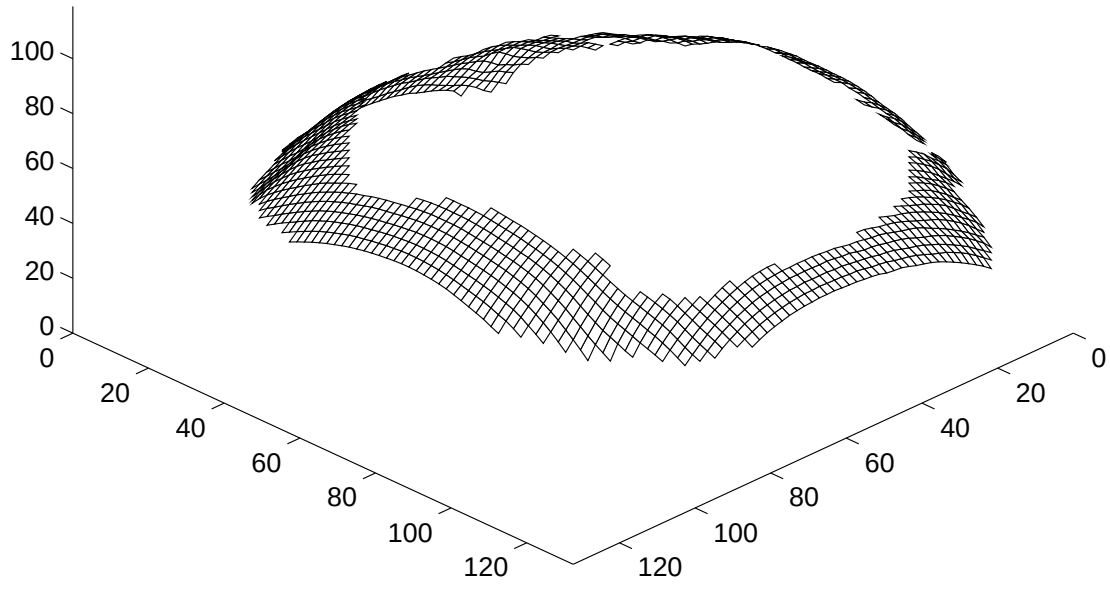

(c)

Fig. 7. Example of fitting a surface to a cranial defect. (a) Rendered view of full CT data set. (b) Detail of the defect region with support region for surface highlighted. (c) Depth-map corresponding to (b).

(a)

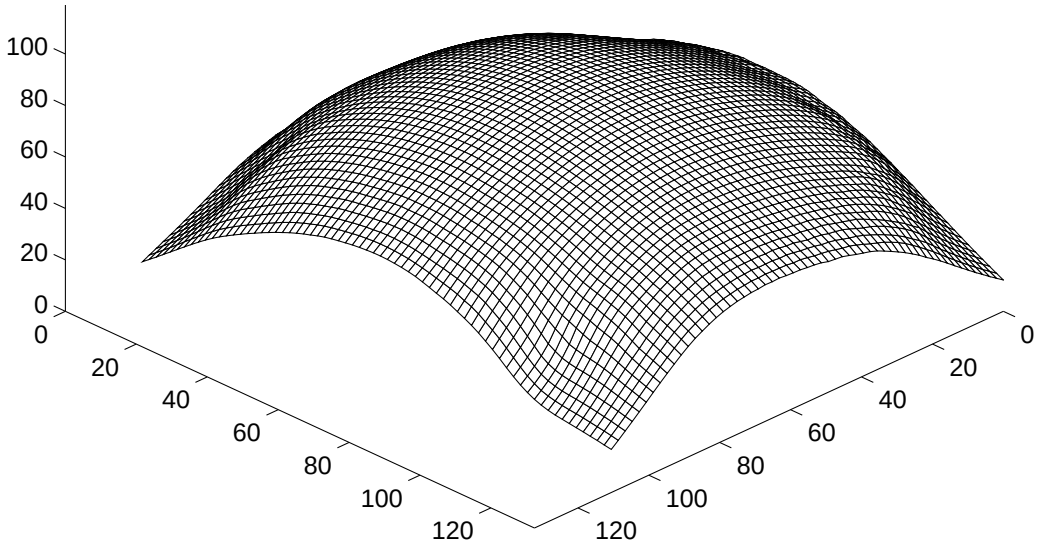

(b)

Fig. 8. (a) Rendered view of surface fitted to depth-map in Figure 7(c). (b) Depth-map of fitted surface.

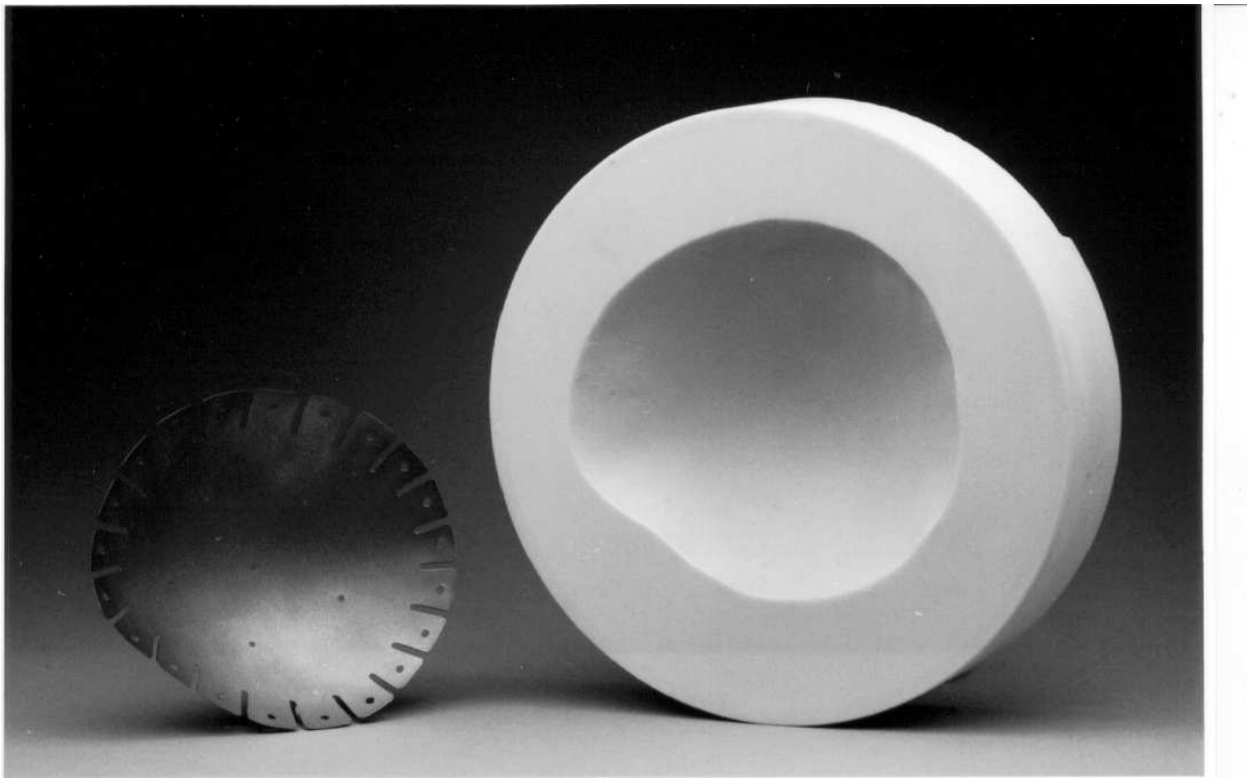

(a)

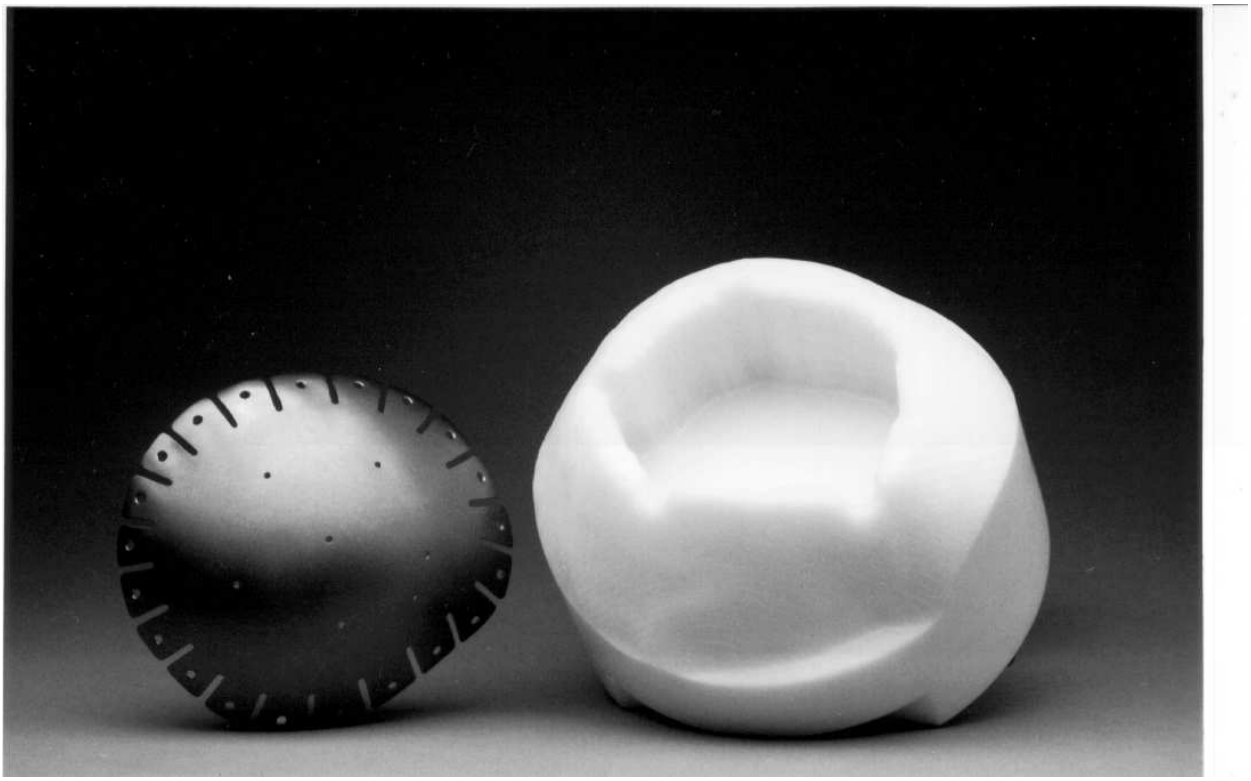

(b)

Fig. 9. (a) Finished plate with the mold used to press the titanium plate. The mold was machined from epoxy-resin using interpolated surface data shown in Figure 8. (b) Finished plate with the plastic model of the defect area machined using depth-map data shown in Figure 7.

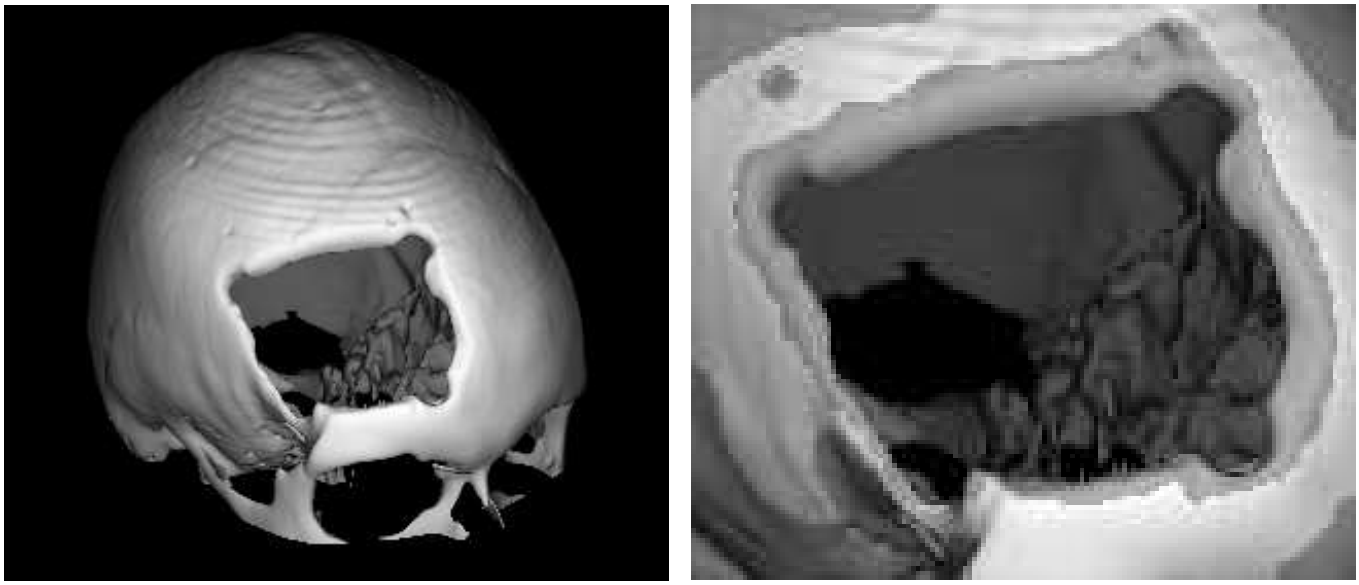

(a) (b)

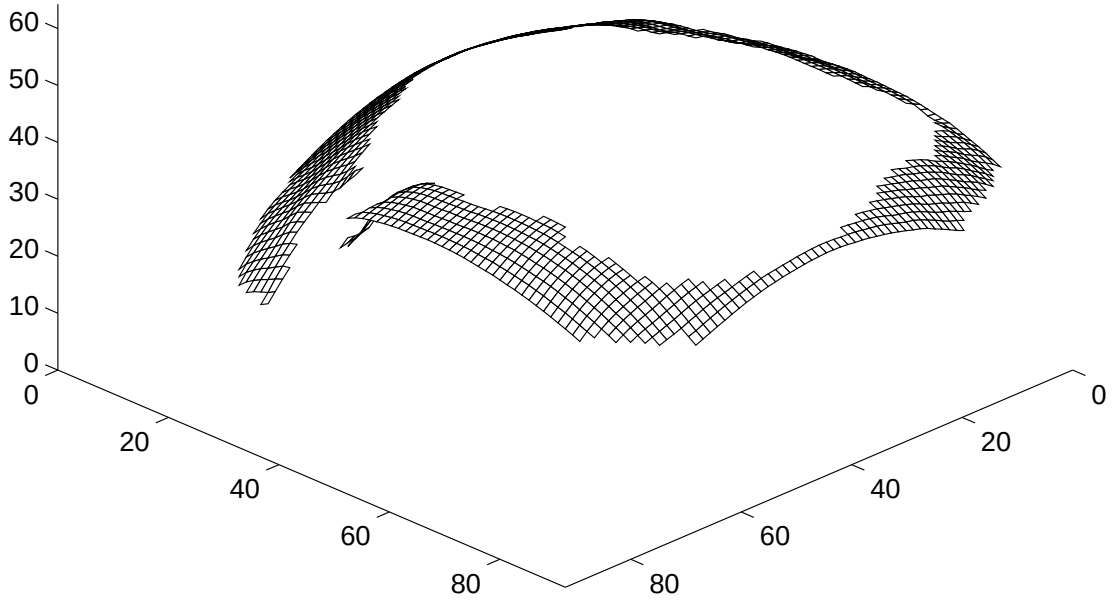

(c)

Fig. 10. Example of fitting a surface to a cranial defect. (a) Rendered view of full CT data set. (b) Detail of the defect region with support region for surface highlighted. (c) Depth-map corresponding to (b).

(a)

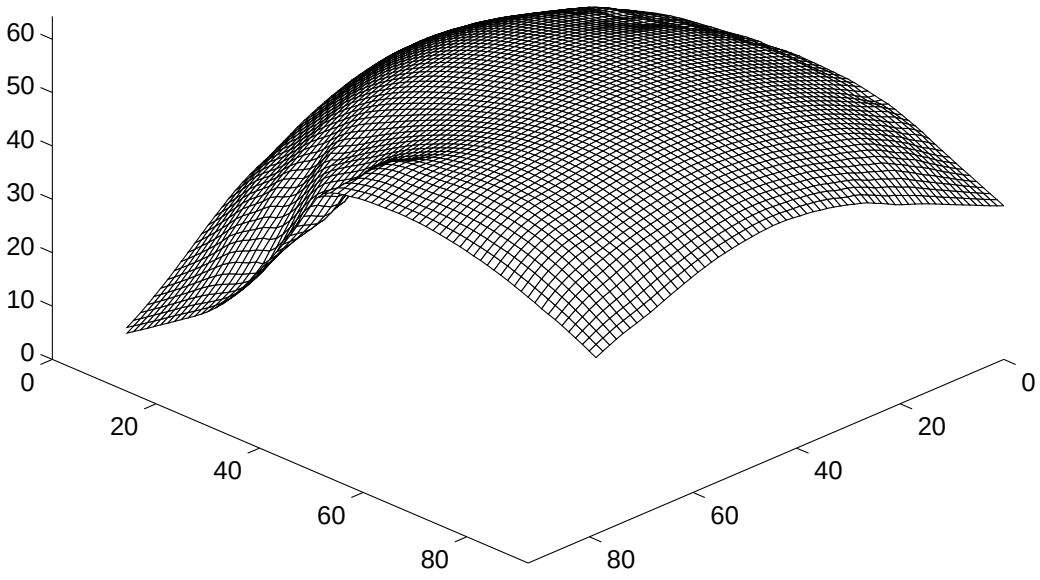

(b)

Fig. 11. (a) Rendered view of fitted surface in Figure 10. (b) Depth-map of fitted surface.