

Parameterization of Triangulations and Unorganized Points

Michael Floater¹ and Kai Hormann²

¹ SINTEF Applied Mathematics, N-0314 Oslo, Norway

² University of Erlangen, D-91058 Erlangen, Germany

Abstract. A parameterization of a surface is a one-to-one mapping from a parameter domain to the surface. This paper studies the construction of piecewise linear parameterizations of triangulations and of discrete parameterizations of sets of unorganized points. Such parameterizations are useful tools for the approximation of triangulations by smoother surfaces as well as for the construction of triangulations from unorganized points. Parameterization can also be used as the starting point for multiresolution analysis of triangulations, through remeshing by triangulations with subdivision connectivity.

1 Introduction

A parametric surface is defined by a one-to-one mapping $\phi : \Omega \rightarrow \mathbb{R}^3$, with $\Omega \subset \mathbb{R}^2$ the parameter domain, and we call ϕ a *parameterization* of the surface. This paper deals with the construction of parameterizations both of triangulated surfaces and of discrete sets of points in \mathbb{R}^3 .

One of the main uses of parameterization is in the approximation of scattered data. In its simplest form, we are given a set of distinct (scattered) points (x_i, y_i) , $1 \leq i \leq N$, in some region Ω of \mathbb{R}^2 , and an associated set of values z_i . The task is to find a function $s : \Omega \rightarrow \mathbb{R}$, with some order of smoothness, which approximates the data in the sense that

$$s(x_i, y_i) \approx z_i.$$

Many representations for s have been proposed, such as piecewise polynomials over a triangulation of the points (x_i, y_i) , radial basis functions, and tensor-product splines, and various methods such as interpolation and least square approximation have been applied to find s ; see [29]. If the data are sampled from some (unknown) function $f(x, y)$, i.e. $f(x_i, y_i) = z_i$, then s will be an approximation to f .

In computer graphics and reverse engineering, however, we frequently encounter a set of points $\mathbf{x}_i = (x_i, y_i, z_i)$ sampled from a surface whose geometry is complex in the sense that it cannot be represented simply as the graph of a bivariate function $f(x, y)$. In this case the function s will not be a suitable approximation to the underlying surface. If, however, the underlying surface is homeomorphic to a simply connected planar region then we can approximate it by a parametric surface. If we can first choose a suitable parameter

domain Ω and suitable parameter points $\mathbf{u}_i = (u_i, v_i)$ in Ω , then a simple solution is to use one of the usual scattered data methods to independently construct three functions $s_1, s_2, s_3 : \Omega \rightarrow \mathbb{R}$ such that

$$s_1(\mathbf{u}_i) \approx x_i, \quad s_2(\mathbf{u}_i) \approx y_i, \quad s_3(\mathbf{u}_i) \approx z_i.$$

Then letting

$$\mathbf{s}(u, v) = (s_1(u, v), s_2(u, v), s_3(u, v)), \quad (1)$$

we have

$$\mathbf{s}(u_i, v_i) \approx (x_i, y_i, z_i).$$

Thus the problem can be reduced to that of finding suitable parameter points \mathbf{u}_i from given data points \mathbf{x}_i . Analogous to the parameterization of surfaces, we view the set of points \mathbf{u}_i as a (discrete) parameterization of the set of points \mathbf{x}_i . Various methods of constructing such parameterizations have been proposed and it has been found that the choice of parameterization usually has a large bearing on the quality of any resulting surface approximation \mathbf{s} .

The construction of such discrete parameterizations is made easier when we have some neighbourhood information about the points \mathbf{x}_i . In fact, if the points \mathbf{x}_i are organized in a triangulation, with the points \mathbf{x}_i the vertices, then any continuous, piecewise linear parameterization of the triangulation contains a discrete parameterization of the points \mathbf{x}_i . We will study the construction of such piecewise linear parameterizations in the first few sections of the paper. Later we will treat the situation in which the points \mathbf{x}_i are truly 'unorganized'. In this case we no longer have the neighbourhood information supplied by a triangulation, and must instead build up neighbourhoods by looking locally for the k nearest points, or using some similar technique.

Parameterizations have several other applications; notably *remeshing* and *texture mapping*. We will briefly describe the use of parameterization for remeshing and how it can be used to set up a multiresolution analysis for triangulations.

2 Parameterization of Simple Triangulations

We begin by studying parameterizations of what, in this paper, we will call *simple* triangulations, namely triangulations which are homeomorphic to (simply connected) planar triangulations. The more complex case of manifold triangulations will be dealt with in a later section.

First we define a planar triangulation. By a triangle we understand the convex hull $T = [v_1, v_2, v_3]$ of three non-collinear points $v_1, v_2, v_3 \in \mathbb{R}^k$ with $k = 2, 3$.

Definition 1. Let $\mathcal{T} = \{T_1, \dots, T_M\}$ be a set of triangles in \mathbb{R}^2 . We call \mathcal{T} a *planar triangulation* if

- (i) $T_i \cap T_j, i \neq j$, is either empty, a common vertex or a common edge,
- (ii) the edges in \mathcal{T} which belong to only one triangle form a simple polygon.

Thus the union of the triangles,

$$\Omega_{\mathcal{T}} = \bigcup_{i=1}^M T_i,$$

is simply connected and its boundary is the polygon in (ii).

In Figure 1, the triangle sets (a) and (b) do not fulfill property (i) of Definition 1 and (ii) is not true for (c). Only the triangle set (d) is a planar triangulation.

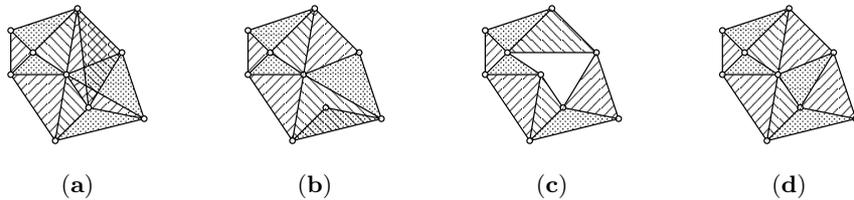


Fig. 1. Of these triangle sets only (d) is a planar triangulation

Next we will define a simple triangulation as the image of a one-to-one mapping from a planar triangulation to \mathbb{R}^3 (or as an embedding of a planar triangulation in \mathbb{R}^3). It will be useful to define $S_k(\mathcal{T})$ to be the linear space of all continuous functions $s : \Omega_{\mathcal{T}} \rightarrow \mathbb{R}^k$ which are linear over each triangle in \mathcal{T} . Thus each of the k components of s belongs to $S_1^0(\mathcal{T})$, where $S_d^r(\mathcal{T})$ is the usual notation for the spline space over $\Omega_{\mathcal{T}}$ of piecewise polynomials of degree d and of smoothness C^r (see [32]).

Definition 2. Let $\mathcal{T} = \{T_1, \dots, T_M\}$ be a planar triangulation. For any injective function $\phi \in S_3(\mathcal{T})$ we call the set of triangles

$$\mathcal{S} = \phi(\mathcal{T}) = \{\phi(T_1), \dots, \phi(T_M)\}$$

a *simple triangulation* and we let

$$\Omega_{\mathcal{S}} = \phi(\Omega_{\mathcal{T}}) = \bigcup_{i=1}^M \phi(T_i).$$

We call $\phi : \Omega_{\mathcal{T}} \rightarrow \Omega_{\mathcal{S}}$ a *parameterization* of $\Omega_{\mathcal{S}}$.

Thus Ω_S is a parametric surface with piecewise linear parameterization ϕ . The following observation shows that Ω_T and Ω_S are in fact homeomorphic.

Proposition 1. *The mapping $\phi : \Omega_T \rightarrow \Omega_S$ is a homeomorphism.*

Proof. By definition ϕ is injective and thus a bijective mapping from Ω_T to Ω_S . Therefore ϕ has an inverse $\psi = \phi^{-1} : \Omega_S \rightarrow \Omega_T$. Due to the compactness of Ω_T and Ω_S and the fact that ϕ is continuous, a well-known result from real analysis (see [19], pp. 94) shows that ψ is also continuous and thus ϕ is a homeomorphism. \square

Figure 2 shows an example of a planar triangulation \mathcal{T} and a corresponding simple triangulation \mathcal{S} in \mathbb{R}^3 .

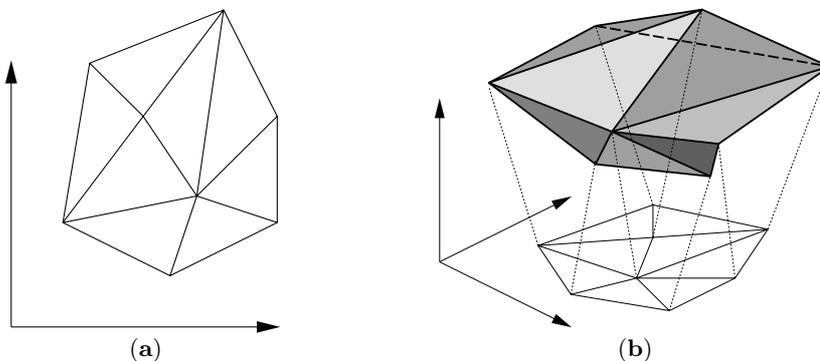


Fig. 2. Planar triangulation (a) and corresponding simple triangulation in \mathbb{R}^3 (b)

Example 1. Let Ω_T be a planar triangulation and let $f \in S_1(\mathcal{T})$. Then the mapping $\phi \in S_3(\mathcal{T})$ defined by $\phi(x, y) = (x, y, f(x, y))$ for $(x, y) \in \Omega_T$ is clearly injective and so $\mathcal{S} = \phi(\mathcal{T})$ is a simple triangulation with parameterization ϕ .

In this example, the surface Ω_S is simply the graph of the piecewise linear bivariate function f . In the next example, it is no longer possible to represent Ω_S as the graph of a bivariate function (in general).

Example 2. Let Ω_T be a planar triangulation contained in the open rectangle $(0, 1) \times (0, 2\pi)$ and let $\phi \in S_3(\mathcal{T})$ be given by $\phi(x, y) = (x, \cos y, \sin y)$ for $(x, y) \in \Omega_T$. Then $\mathcal{S} = \phi(\mathcal{T})$ is a simple triangulation with parameterization ϕ .

Here, the vertices of \mathcal{T} are mapped onto the cylinder of radius 1 and whose axis is the x -axis.

Example 3. A further example of a simple triangulation is given in Figure 3.

The situation we want to consider now is that we are given a simple triangulation \mathcal{S} only. Our task is to find a planar triangulation \mathcal{T} and associated parameterization $\phi \in S_3(\mathcal{T})$ such that $\phi(\mathcal{T}) = \mathcal{S}$. We solve this by constructing a one-to-one piecewise linear mapping $\psi : \Omega_{\mathcal{S}} \rightarrow \mathbb{R}^2$ and setting $\phi = \psi^{-1}$. As for planar triangulations, we can use the linear space notation and we define $S_k(\mathcal{S})$ as the linear space of all continuous functions $s : \Omega_{\mathcal{S}} \rightarrow \mathbb{R}^k$ which are linear over each triangle of \mathcal{S} . We can then express the problem as that of constructing an element ψ of $S_2(\mathcal{S})$ which is injective.

Example 4. If $\Omega_{\mathcal{S}}$ is the graph of the function $f \in S_1(\mathcal{T})$, as in Example 1, then the projection $\psi \in S_2(\mathcal{S})$ given by $\psi(x, y, z) = (x, y)$ is injective and its inverse ϕ is as in Example 1.

Projection from \mathbb{R}^3 into the (x, y) plane is a valid form of parameterization whenever the triangulated surface $\Omega_{\mathcal{S}}$ is the graph of a function of x and y . However, such a projective parameterization can lead to very poor results when used for surface fitting and it may well be better (especially when f has steep gradients for example) to fit a parametric surface based on one of the kinds of parameterization we will discuss later.

3 Parameterization by Linear Methods

We now describe a method for constructing an injective mapping $\psi \in S_2(\mathcal{S})$. Since such a mapping ψ is piecewise linear, it is completely determined by the (parameter) points $\psi(\mathbf{v}) \in \mathbb{R}^2$ for vertices \mathbf{v} in the vertex set $V = V(\mathcal{S})$ of \mathcal{S} .

Let V_I denote the interior vertices of \mathcal{S} and V_B the boundary ones. Due to Definition 1, the boundary vertices of \mathcal{S} form a polygon $\partial\mathcal{S}$ in \mathbb{R}^3 which we call the *boundary polygon* of \mathcal{S} . Two distinct vertices \mathbf{v} and \mathbf{w} in \mathcal{S} are *neighbours* if they are the end points of some edge in \mathcal{S} . For each $\mathbf{v} \in V$, let

$$N_{\mathbf{v}} = \{\mathbf{w} \in V : [\mathbf{w}, \mathbf{v}] \in E\},$$

the set of neighbours of \mathbf{v} , where $E = E(\mathcal{S})$ is the set of edges in \mathcal{S} .

The first step of the method is to choose any points $\psi(\mathbf{v}) \in \mathbb{R}^2$, for $\mathbf{v} \in V_B$, such that the boundary polygon $\partial\mathcal{S}$ of \mathcal{S} is mapped into a polygon $\psi(\partial\mathcal{S})$ in the plane. In the second step, for $\mathbf{v} \in V_I$, we choose a set of strictly positive values $\lambda_{\mathbf{vw}}$, for $\mathbf{w} \in N_{\mathbf{v}}$, such that

$$\sum_{\mathbf{w} \in N_{\mathbf{v}}} \lambda_{\mathbf{vw}} = 1. \quad (2)$$

Then we let the points $\psi(\mathbf{v})$ in \mathbb{R}^2 , for $\mathbf{v} \in V_I$, be the unique solutions of the linear system of equations

$$\psi(\mathbf{v}) = \sum_{\mathbf{w} \in N_{\mathbf{v}}} \lambda_{\mathbf{vw}} \psi(\mathbf{w}), \quad \mathbf{v} \in V_I. \quad (3)$$

Since these equations force each point $\psi(\mathbf{v})$ to be a convex combination of its neighbouring points $\psi(\mathbf{w})$, we call ψ a *convex combination mapping*. Figure 3 shows an example of a simple triangulation (in \mathbb{R}^3). Figure 4 shows a convex combination mapping of \mathcal{S} into a planar triangulation \mathcal{T} , whose boundary was chosen to be a rectangle. Figure 4 also shows a tensor-product spline approximation (in fact a least square approximation, taken from [9]) to the vertices of \mathcal{S} based on their parameter points, the vertices of \mathcal{T} .

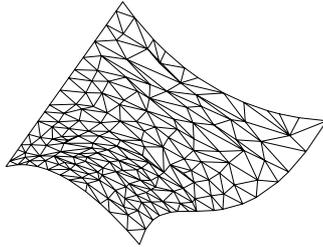


Fig. 3. Simple triangulation

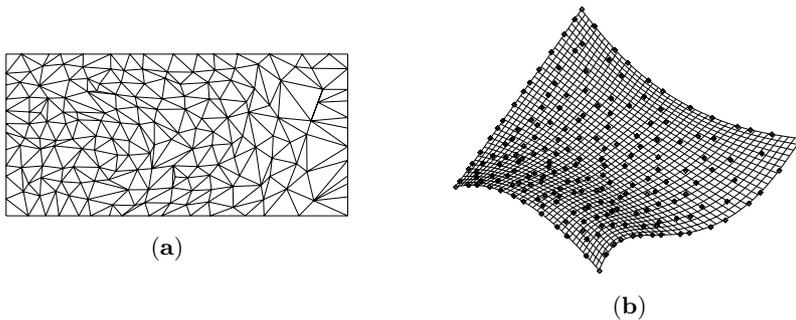


Fig. 4. Parameterization (a) and resulting tensor-product spline approximation (b)

Let us take a closer look at the linear system. We must show that it has a unique solution. To this end, note that it can be rewritten in the form

$$\psi(\mathbf{v}) - \sum_{\mathbf{w} \in N_{\mathbf{v}} \cap V_I} \lambda_{\mathbf{vw}} \psi(\mathbf{w}) = \sum_{\mathbf{w} \in N_{\mathbf{v}} \cap V_B} \lambda_{\mathbf{vw}} \psi(\mathbf{w}), \quad \mathbf{v} \in V_I. \quad (4)$$

This can be written as the matrix equation

$$A\mathbf{x} = \mathbf{b},$$

where $\mathbf{x} = (\psi(\mathbf{w}))_{\mathbf{w} \in V_I}$ is the column vector of unknowns in some arbitrary ordering, \mathbf{b} is the column vector whose elements are the right hand sides

of (4), and the matrix $A = (a_{vw})_{v,w \in V_I}$ has dimension $n \times n$, with $n = |V_I|$, and elements

$$a_{vw} = \begin{cases} 1, & w = v, \\ -\lambda_{vw}, & w \in N_v, \\ 0, & \text{otherwise.} \end{cases}$$

The existence and uniqueness of the solution to (3) follows from the structure of the matrix A , namely that its off diagonal elements are either zero or negative and each row of A is diagonally dominant. Moreover every row corresponding to a vertex $v \in V_I$ which has at least one neighbour in V_B is strictly diagonally dominant and every interior vertex can be connected to the boundary by a path of vertices. A standard result in linear algebra shows then that A is non-singular (in fact A is a so-called M-matrix, and such matrices frequently occur in numerical approximations to elliptic partial differential equations; see [31]).

The interesting question is whether ψ is one-to-one. It will not be one-to-one in general but the following result from [10] gives a sufficient condition. We say that an interior edge of \mathcal{S} is a *dividing edge* of \mathcal{S} if both its end points are boundary vertices of \mathcal{S} .

Theorem 1. *Suppose $\psi \in S_2(\mathcal{S})$ is a convex combination map which maps the boundary $\partial\mathcal{S}$ homeomorphically to the boundary $\partial\Omega$ of a convex planar region Ω . Then ψ is one-to-one if and only if no dividing edge $[v, w]$ is mapped by ψ into $\partial\Omega$.*

4 Choosing the Weights

A simple choice of weights λ_{vw} is to take them to be constant for each vertex v , i.e. $\lambda_{vw} = 1/d(v)$, $w \in N_v$ where $d(v)$ is the degree $|N_v|$ of v . Then every interior vertex $\psi(v)$ of the solution to the linear system will be the barycentre of its neighbours. We then say that $\phi = \psi^{-1} : \Omega_{\mathcal{T}} \rightarrow \Omega_{\mathcal{S}}$ is a *uniform parameterization* of $\Omega_{\mathcal{S}}$. However, numerical examples show that this parameterization usually leads to poor spline surfaces (1) when used for approximation. Look for example at Figures 5 and 6. Figure 5 shows a simple triangulation \mathcal{S} and Figure 6 shows the planar triangulation \mathcal{T} of a uniform parameterization of \mathcal{S} together with a Clough-Tocher (C^1 piecewise cubic) interpolant to \mathcal{S} over the triangulation \mathcal{T} (see [29]). Clearly the iso-curves are badly behaved.

One reason for the bad behaviour of the surface approximation is that the weights λ_{vw} are independent of the geometry of the vertices v of \mathcal{S} . In practice it is much better to choose weights for which the Euclidean distance

$$\|v - \sum_{w \in N_v} \lambda_{vw} w\| \tag{5}$$

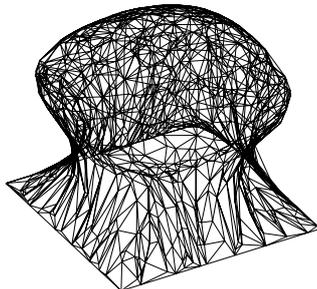
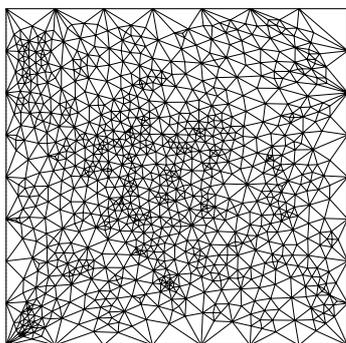
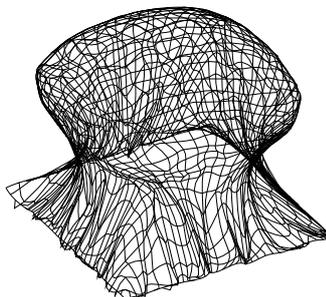


Fig. 5. Simple triangulation



(a)



(b)

Fig. 6. Uniform parameterization (a) and Clough-Tocher interpolant (b)

between \mathbf{v} and $\sum_{\mathbf{w} \in N_{\mathbf{v}}} \lambda_{\mathbf{v}\mathbf{w}} \mathbf{w}$ is as small as possible and in particular, when \mathbf{v} and its neighbours lie in a plane, we should have

$$\mathbf{v} = \sum_{\mathbf{w} \in N_{\mathbf{v}}} \lambda_{\mathbf{v}\mathbf{w}} \mathbf{w}. \quad (6)$$

This latter condition implies linear precision: if the whole triangulation \mathcal{S} lies in a plane and $\psi|_{\partial\mathcal{S}}$ is an affine mapping then the whole convex combination mapping ψ is an affine mapping.

We will now describe a certain choice of (positive) weights $\lambda_{\mathbf{v}\mathbf{w}}$ in (3) which has these properties (see [7]). For each $\mathbf{v} \in V_I$ we compute, in two steps, all the $\lambda_{\mathbf{v}\mathbf{w}}$ for $\mathbf{w} \in N_{\mathbf{v}}$ according to \mathbf{v} and its neighbours in \mathcal{S} . Let us label these neighbours $\mathbf{v}_1, \dots, \mathbf{v}_d$, where $d = d(\mathbf{v})$, in some anticlockwise order around \mathbf{v} relative to \mathcal{S} .

Let $\mathcal{C}_{\mathbf{v}}$ be the set of all triangles in \mathcal{S} containing \mathbf{v} . We call the union of these triangles, $\Omega_{\mathcal{C}_{\mathbf{v}}}$, the *cell* of \mathbf{v} . The first step is to use a local injective mapping $\psi_{\mathbf{v}} \in S_2(\mathcal{C}_{\mathbf{v}})$ in order to “flatten out” the cell $\Omega_{\mathcal{C}_{\mathbf{v}}}$ into the plane, yielding local (temporary) parameter points $\mathbf{p} = \psi_{\mathbf{v}}(\mathbf{v})$ and $\mathbf{p}_j = \psi_{\mathbf{v}}(\mathbf{v}_j)$; see

Figure 7. We use an approximation of the geodesic polar map, adapted to triangulations. We let \mathbf{p} be arbitrary and choose the neighbours \mathbf{p}_j such that for each $j = 1, \dots, d$,

$$\|\mathbf{p}_j - \mathbf{p}\| = \|\mathbf{v}_j - \mathbf{v}\|$$

and for each triangle $[\mathbf{v}, \mathbf{v}_j, \mathbf{v}_k]$ in \mathcal{C}_v ,

$$\text{ang}(\mathbf{p}_k, \mathbf{p}, \mathbf{p}_j) = \rho \text{ang}(\mathbf{v}_k, \mathbf{v}, \mathbf{v}_j),$$

where ρ is a constant. The scaling factor ρ is chosen to ensure that the interior angles in the mapped cell $\psi_v(\Omega_{\mathcal{C}_v})$ sum to 2π . Note here that in the event that the cell $\Omega_{\mathcal{C}_v}$ is planar, we have $\rho = 1$ and the local mapping ψ_v is an affine mapping.

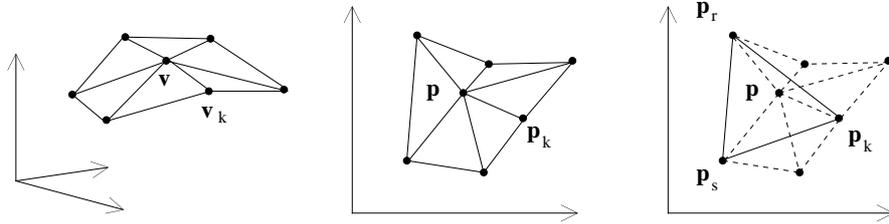


Fig. 7. Calculating the shape-preserving weights

The second step is to express \mathbf{p} as a convex combination of the neighbouring mapped points $\mathbf{p}_1, \dots, \mathbf{p}_d$ in order to obtain linear precision. For each $k = 1, \dots, d$, we locate an edge $[\mathbf{p}_r, \mathbf{p}_s]$ in $\psi_v(\mathcal{C}_v)$ for which

$$\mathbf{p} \in [\mathbf{p}_k, \mathbf{p}_r, \mathbf{p}_s],$$

and with $\tau_k^k, \tau_r^k, \tau_s^k$ the barycentric coordinates of \mathbf{p} in this latter triangle, we have

$$\mathbf{p} = \tau_k^k \mathbf{p}_k + \tau_r^k \mathbf{p}_r + \tau_s^k \mathbf{p}_s.$$

Letting $\tau_j^k = 0$ for all $j \neq k, r, s$ we then have

$$\mathbf{p} = \sum_{j=1}^d \tau_j^k \mathbf{p}_j.$$

Finally, we take the weights to be averages of the local weights τ_j^k over all $k = 1, \dots, d$,

$$\lambda_{\mathbf{v}\mathbf{v}_j} = \frac{1}{d} \sum_{k=1}^d \tau_j^k, \quad (7)$$

and we have

$$\mathbf{p} = \sum_{j=1}^d \lambda_{\mathbf{v}\mathbf{v}_j} \mathbf{p}_j, \quad \text{and} \quad \sum_{j=1}^d \lambda_{\mathbf{v}\mathbf{v}_j} = 1,$$

and $\lambda_{\mathbf{v}\mathbf{v}_j} > 0$ for $j = 1, \dots, d$.

Using the fact that each polar map is an affine map when the associated cell $\Omega_{\mathcal{C}_v}$ is planar, equation (6) follows. The inverse $\phi = \psi^{-1} : \Omega_{\mathcal{T}} \rightarrow \Omega_{\mathcal{S}}$ is called a *shape-preserving parameterization* of $\Omega_{\mathcal{S}}$.

Figure 8 shows the result of interpolating \mathcal{S} of Figure 5 with a Clough-Tocher interpolant over the planar triangulation \mathcal{T} of a shape-preserving parameterization of \mathcal{S} . The surface approximation is clearly better than that of Figure 6, using uniform parameterization.

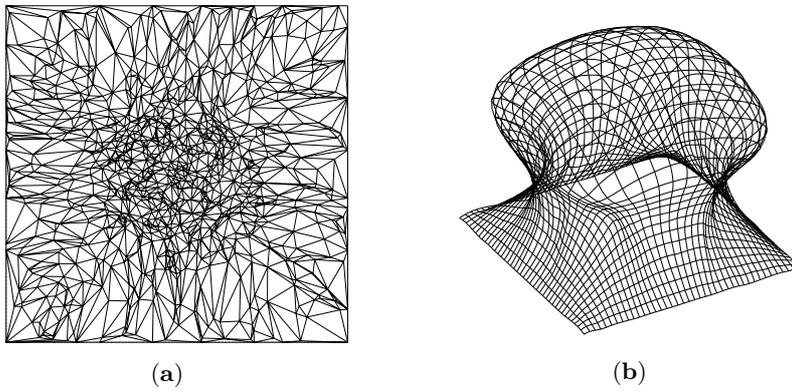


Fig. 8. Shape-preserving parameterization (a) and Clough-Tocher interpolant (b)

The effect of choosing different boundary polygons $\psi(\partial\mathcal{S})$ is shown in Figure 10. In (a) the parameter points of the boundary vertices were distributed by chord length on a rectangle and in (b) on a circle.

5 Parameterization by Weighted Least Squares

A special case of convex combination parameterizations arises from minimizing so-called spring energy. First, choose as before a convex polygon $\psi(V_B)$. Secondly, for each interior edge $[\mathbf{v}, \mathbf{w}]$ in \mathcal{S} choose some value $\mu_{\mathbf{v}\mathbf{w}} = \mu_{\mathbf{w}\mathbf{v}} > 0$. Then let the points $\psi(\mathbf{v})$, $\mathbf{v} \in V_I$, minimize the function

$$F = \sum_{[\mathbf{v}, \mathbf{w}] \in E(\mathcal{S})} \mu_{\mathbf{v}\mathbf{w}} \|\psi(\mathbf{v}) - \psi(\mathbf{w})\|^2. \quad (8)$$

The normal equations for (8) are

$$\psi(\mathbf{v}) = \frac{\sum_{\mathbf{w} \in N_{\mathbf{v}}} \mu_{\mathbf{v}\mathbf{w}} \psi(\mathbf{w})}{\sum_{\mathbf{w} \in N_{\mathbf{v}}} \mu_{\mathbf{v}\mathbf{w}}}, \quad \mathbf{v} \in V_I$$

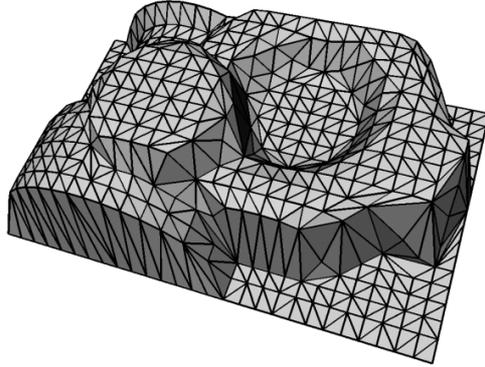


Fig. 9. Simple triangulation

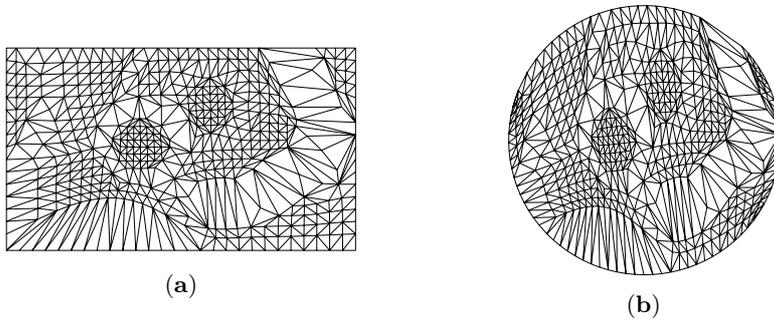


Fig. 10. Shape-preserving parameterizations of the triangulation in Figure 9

and so minimizing F is equivalent to solving (3) where

$$\lambda_{vw} = \frac{\mu_{vw}}{\sum_{w \in N_v} \mu_{vw}}, \quad v \in V_I, \quad w \in N_v.$$

Notice that in general $\lambda_{vw} \neq \lambda_{wv}$ even though $\mu_{vw} = \mu_{wv}$. We call $\phi = \psi^{-1} : \Omega_{\mathcal{T}} \rightarrow \Omega_{\mathcal{S}}$ a *least squares parameterization* of $\Omega_{\mathcal{S}}$. As an example one might choose $\mu_{vw} = 1/\|v - w\|$, but this method will not have linear precision. Currently it is not known whether it is possible to find coefficients μ_{vw} which are at the same time positive and yield linear precision. A choice of coefficients which does have linear precision is that of the so-called *discrete harmonic map* ψ , first proposed in [25] in the context of differential geometry and later in [5] from the point of view of computer graphics. It is based on the fact that harmonic maps minimize the *Dirichlet energy*, which for a function $f : \Omega \rightarrow \mathbb{R}^k$ is defined as

$$E_D(f) = \frac{1}{2} \int_{\Omega} \|\nabla f\|^2.$$

For our piecewise linear functions ψ we find $E_D(\psi)$ to be of the form (8) with

$$\mu_{vw} = \frac{1}{8}(\cot \alpha + \cot \beta),$$

where α and β are the angles opposite to $[v, w]$ in the triangles adjacent to this edge. Unfortunately, these weights will rarely all be positive and so Theorem 1 is no longer applicable. In fact, as observed in [10],

$$\cot \alpha + \cot \beta = \frac{\sin(\alpha + \beta)}{\sin \alpha \sin \beta},$$

which means that $\mu_{vw} \geq 0$ if and only if $\alpha + \beta \leq \pi$. It follows [10] that the condition that $\mu_{vw} \geq 0$ for every interior edge $[v, w]$ of \mathcal{S} is precisely the condition that \mathcal{S} is a Delaunay triangulation (see Section 6.1 of [17]) in the case that \mathcal{S} is planar. Through a simple counterexample, it was shown in [8] that when some weights μ_{vw} are negative, the discrete harmonic mapping ψ may not be one-to-one.

6 Parameterization by Non-Linear Methods

The main advantage of convex combination parameterizations is (i) that they only require the solution of a linear system, and (ii) that they guarantee a one-to-one mapping in the case that the boundary of the parameter domain is taken to be convex.

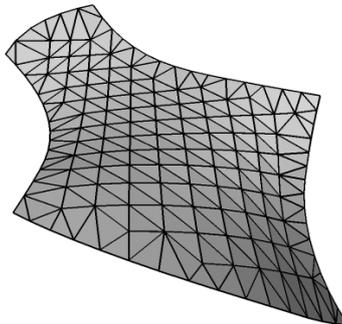


Fig. 11. Simple triangulation

However, for some triangulations we may want to parameterize \mathcal{S} over a domain with a non-convex boundary $\psi(\partial\mathcal{S})$, typically one that reflects the shape of the boundary of the triangulation $\partial\mathcal{S}$. If we choose a circle or a rectangle as boundary polygon for the triangulation in Figure 11 we obtain a parameterization with quite a few distorted triangles near the boundary (Figure 12 (a) and (b)), while projecting $\partial\mathcal{S}$ into the plane that fits the

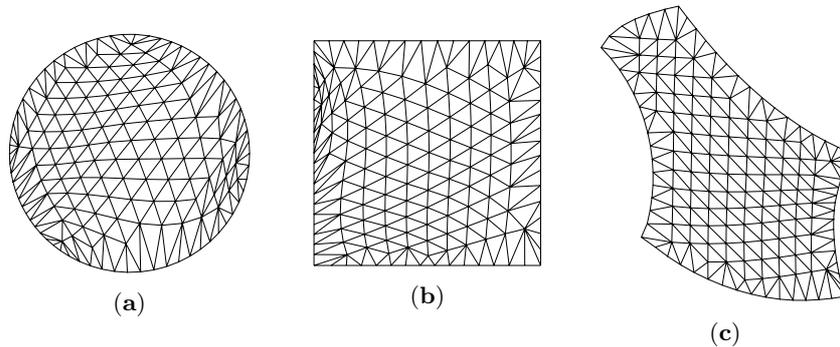


Fig. 12. Shape-preserving parameterizations of the triangulation in Figure 11

boundary vertices best in a least square sense gives a much better result (Figure 12 (c)).

We will now describe an alternative to convex combination maps which simultaneously generates both a parameterization and the shape of the parameter domain. Although, in contrast to Theorem 1, the injectivity of ψ in this method can no longer be guaranteed, ψ usually turns out to be injective in practice. However, since the method is non-linear, it is not surprisingly slower than the method of convex combination maps. Moreover, it requires an initial one-to-one mapping ψ (such as a convex combination map) as the start point of an iteration towards the solution. We therefore view this method as a useful enhancement to a convex combination map through iterative improvement when it is essential to have a boundary which is far from convex.

The method is based on the observation that any mapping $\psi \in S_2(\mathcal{S})$ will normally deform the shape of the triangles of \mathcal{S} and aims at minimizing these deformations. Consider e.g. the configuration of the cell $\Omega_{\mathcal{C}_v}$ in Figure 7 which can only be parameterized without distorting the triangles if the interior angles sum to 2π , e.g. if $\Omega_{\mathcal{C}_v}$ is contained in a plane. This reflects a well-known fact from differential geometry [3] that only developable surfaces (e.g. planar, cylindrical and conical surfaces) can be parameterized without any distortion. Such surfaces are *isometric* to the plane.

Definition 3. Let $\Omega \subset \mathbb{R}^2$ and $f : \Omega \rightarrow \mathbb{R}^3$ be a differentiable and injective mapping and thus a parameterization of the surface $f(\Omega) \subset \mathbb{R}^3$. The symmetric 2×2 matrix

$$\mathbf{I}_f = \nabla f^T \cdot \nabla f$$

is called the *first fundamental form* of f . The parameterization f is called *isometric*, if

$$\mathbf{I}_f = \text{Id}$$

and a surface $S \subset \mathbb{R}^3$ is called *isometric* to Ω if there exists an isometric parameterization f of S with $f(\Omega) = S$.

Example 5. For the mapping $f : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ with

$$f \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} 4u^3 + 1 \\ 3u^3 + 3v \\ v - 2 \end{pmatrix}$$

we have

$$\nabla f = \begin{pmatrix} 12u^2 & 0 \\ 9u^2 & 3 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad \mathbf{I}_f = \begin{pmatrix} 225u^4 & 27u^2 \\ 27u^2 & 10 \end{pmatrix},$$

thus f is not isometric. But still the surface $S = f(\mathbb{R}^2)$ is isometric to \mathbb{R}^2 which can be seen by applying the two reparameterizations $\pi_1, \pi_2 : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ with

$$\pi_1 \begin{pmatrix} u \\ v \end{pmatrix} = \frac{1}{65} \begin{pmatrix} 13 & -9 \\ 0 & 25 \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} \quad \text{and} \quad \pi_2 \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} \sqrt[3]{u} \\ v \end{pmatrix}$$

and verifying that the alternative parameterization

$$g = f \circ \pi_2 \circ \pi_1$$

of $S = f(\mathbb{R}^2) = g(\mathbb{R}^2)$ is isometric.

Example 6. The vertical cylinder $C = \{(x, y, z) \in \mathbb{R}^3 : x^2 + y^2 = 1\}$ can be parameterized over $\Omega = \{(u, v) \in \mathbb{R}^2 : 0 \leq u < 2\pi\}$ by $f : \Omega \rightarrow C$ with

$$f \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} \cos u \\ \sin u \\ v \end{pmatrix} \quad \implies \quad \nabla f = \begin{pmatrix} -\sin u & 0 \\ \cos u & 0 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad \mathbf{I}_f = \text{Id},$$

and therefore C is isometric to Ω .

The first fundamental form \mathbf{I}_f expresses how a surface inherits the metric from the parameter domain by f . The further \mathbf{I}_f deviates from the identity matrix, the more the metric quantities (length, angle, and area) are distorted. Surfaces that are parameterized isometrically will therefore behave (locally) like planes which is very desirable in many applications. If we use f and g of Example 5 to paste texture information onto the surface S (Figure 13), we can see how f distorts the image while the isometric parameterization g perfectly preserves the planar metric.

Let us now return to the problem of finding for a given simple triangulation \mathcal{S} an injective mapping $\psi \in S_2(\mathcal{S})$ such that the distortion between the triangles $T \in \mathcal{S}$ and their images $\psi(T)$ is small. Since ψ is piecewise linear over $\Omega_{\mathcal{S}}$, we can write the restriction of ψ to each $T \in \mathcal{S}$ as an affine mapping

$$\psi(\mathbf{x}) = A\mathbf{x} + \mathbf{p}, \quad \mathbf{x} \in T$$

with $\mathbf{p} \in \mathbb{R}^2$ and $A \in \mathbb{R}^{2 \times 3}$. Note that by introducing a local orthonormal coordinate system over T with the third axis perpendicular to T , we can assume that T lies in the xy -plane, in which case $A \in \mathbb{R}^{2 \times 2}$ and

$$\psi|_T = f|_T,$$

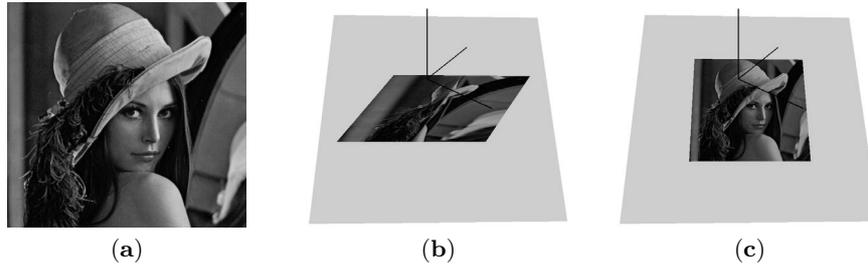


Fig. 13. Mapping a texture map (a) to a surface using a non-isometric (b) and an isometric parameterization (c)

for some affine function $f : \mathbb{R}^2 \rightarrow \mathbb{R}^2$, i.e.,

$$f(\mathbf{x}) = A\mathbf{x} + \mathbf{p}, \quad \mathbf{x} \in \mathbb{R}^2. \quad (9)$$

The question now is how to measure the shape deformation of $\psi|_T$. We do this by measuring the shape deformation of f in terms of the vector \mathbf{p} and the matrix A which define it. Let Π_1 denote the linear space of all linear functions $f : \mathbb{R}^2 \rightarrow \mathbb{R}^2$. We will say that a functional $E : \Pi_1 \rightarrow \mathbb{R}$ is *invariant* with respect to a linear function $\pi \in \Pi_1$ if $E(\pi \circ f) = E(f)$ for all $f \in \Pi_1$. We call any functional $E : \Pi_1 \rightarrow \mathbb{R}$ a *deformation functional* if E is

- invariant w.r.t. translations, (*1)
- invariant w.r.t. orthogonal transformations, (*2)
- invariant w.r.t. uniform scalings, (*3)

since these are the linear operations that do not modify the shape of a triangle (cf. Figure 14). They are also commonly known as *similarity transformations*. Note that the *Green-Lagrange deformation tensor* $\|\mathbf{I}_f - \text{Id}\|^2$ which Maillot et al. [24] suggested in order to minimize distortion is not a deformation functional.

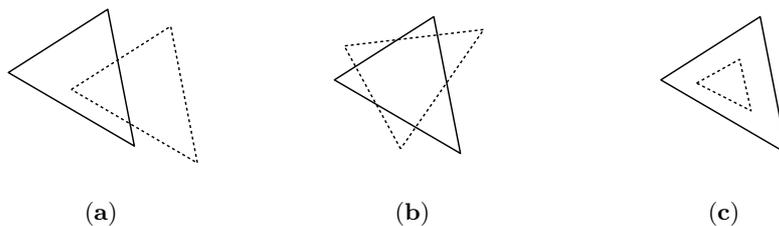


Fig. 14. Translations (a), orthogonal transformations (b) and uniform scalings (c) do not modify the shape of a triangle

Theorem 2. *Let $E : \Pi_1 \rightarrow \mathbb{R}$ be a deformation functional. and let $f \in \Pi_1$ be as in (9). Then $E(f)$ only depends on the ratio $\rho = \sigma_1/\sigma_2$ of the singular values $\sigma_1 \geq \sigma_2 \geq 0$ of A . Any such deformation functional can thus be expressed as a function $E : \{\rho \in \mathbb{R} : \rho \geq 1\} \rightarrow \mathbb{R}$.*

Proof. Because of (*1), $E(f)$ does not depend on the constant part \mathbf{p} of f . If we then consider the *singular value decomposition* [14] of A which guarantees the existence of orthogonal matrices U and V such that

$$U^t A V = \Sigma = \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix},$$

where $\sigma_1 \geq \sigma_2 \geq 0$ are the singular values of A , it is clear that any E fulfilling (*1) and (*2) depends on σ_1 and σ_2 only. Due to (*3), E depends only on the ratio σ_1/σ_2 . \square

Since the ratio $\rho = \sigma_1/\sigma_2$ is minimal ($\rho = 1$) for all similarity transformations and maximal ($\rho = \infty$) for singular mappings ($\sigma_2 = 0$) that collapse triangles to lines, we are specifically interested in deformation functionals which are increasing as a function of ρ . This motivates

Definition 4. A *proper deformation functional* is a deformation functional E which is monotonically increasing as a function of ρ .

We can use any such deformation functional E to find a mapping ψ^* in $S_2(\mathcal{S})$ which minimizes

$$F(\psi) = \sum_{T \in \mathcal{S}} E(\psi|_T). \quad (10)$$

Since $F(\pi \circ \psi) = F(\psi)$ for any similarity transformation π , we need to restrict the minimization in order to ensure uniqueness of the minimum, at least locally. This can be achieved in practice by iterating towards a minimum from an initial starting point ψ_0 .

We will now look for a choice of local deformation functional E for which the global functional F is relatively simple to minimize. The concept of *condition numbers* of matrices [14] provides an adequate tool. The condition number $\kappa(X) = \|X\| \|X^{-1}\|$ measures the distance of a matrix X to the set of singular matrices w.r.t. a certain matrix norm $\|\cdot\|$ [14].

Corollary 1. *The condition numbers κ_2 and κ_F of A and of the first fundamental form $\mathbf{I}_f = A^T A$ are proper shape deformation functionals.*

Proof. From linear algebra [14] we know that the 2-norm and the Frobenius norm of a 2×2 matrix are

$$\|A\|_2 = \sigma_1 \quad \text{and} \quad \|A\|_F^2 = \sigma_1^2 + \sigma_2^2$$

and that the singular values of \mathbf{I}_f are σ_1^2 and σ_2^2 . Therefore

$$E_1(f) = \kappa_2(A) = \|A\|_2 \|A^{-1}\|_2 = \frac{\sigma_1}{\sigma_2},$$

$$E_2(f) = \kappa_F(A) = \|A\|_F \|A^{-1}\|_F = \sqrt{\sigma_1^2 + \sigma_2^2} \sqrt{\left(\frac{1}{\sigma_1}\right)^2 + \left(\frac{1}{\sigma_2}\right)^2} = \frac{\sigma_1^2 + \sigma_2^2}{\sigma_1 \sigma_2},$$

$$E_3(f) = \kappa_2(\mathbf{I}_f) = \kappa_2(A^t A) = \frac{\sigma_1^2}{\sigma_2^2},$$

$$E_4(f) = \kappa_F(\mathbf{I}_f) = \kappa_F(A^t A) = \frac{\sigma_1^4 + \sigma_2^4}{\sigma_1^2 \sigma_2^2},$$

or, if seen as functions of $\rho = \sigma_1/\sigma_2$,

$$E_1(\rho) = \rho, \quad E_2(\rho) = \rho + \frac{1}{\rho}, \quad E_3(\rho) = \rho^2, \quad E_4(\rho) = \left(\rho + \frac{1}{\rho}\right)^2 - 2. \quad \square$$

In order to make the minimization of (10) as simple as possible, we would further like to choose a deformation functional E that is a simple function of the unknown parameter points $\psi(\mathbf{v})$, $\mathbf{v} \in T$. We cannot expect to find a deformation functional which is a quadratic polynomial in the coordinates of the three points $\psi(\mathbf{v}_1)$, $\psi(\mathbf{v}_2)$, $\psi(\mathbf{v}_3)$, as is the case for the discretization of the harmonic map of Section 5. However, it can be shown ([16]) that E_2 is a rational quadratic polynomial in these coordinates.

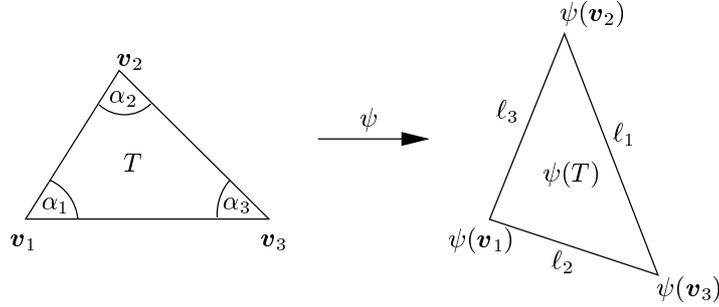


Fig. 15. Two-dimensional linear mapping ψ

Proposition 2. *Given the notations of Figure 15, the proper deformation functional E_2 can be written as*

$$E_2(\psi|_T) = \frac{\cot \alpha_1 \|\ell_1\|^2 + \cot \alpha_2 \|\ell_2\|^2 + \cot \alpha_3 \|\ell_3\|^2}{\ell_2 \times \ell_3}$$

with $\ell_1 = \psi(\mathbf{v}_3) - \psi(\mathbf{v}_2)$, $\ell_2 = \psi(\mathbf{v}_3) - \psi(\mathbf{v}_1)$, and $\ell_3 = \psi(\mathbf{v}_2) - \psi(\mathbf{v}_1)$.

With this choice of deformation functional the inverse $\phi = \psi^{-1} : \Omega_{\mathcal{T}} \rightarrow \Omega_{\mathcal{S}}$ was called a *most isometric parameterization* [16]. Figure 16 shows the most isometric parameterization of the triangulation in Figure 9.

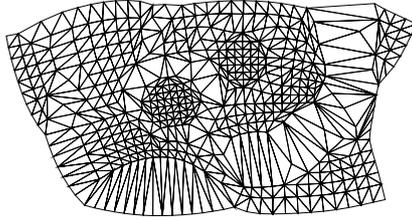


Fig. 16. Most isometric parameterization of the triangulation in Figure 9

Finally, we remark that another characterization of a deformation functional E as defined above is simply that it should depend only on the angles $\theta_1, \theta_2, \theta_3$ of the image $\psi(T)$ of the triangle T , for these angles are unchanged under translations and uniform and orthogonal scalings of $\psi(T)$. This implies that any of the deformation functionals derived above can be expressed purely as functions of these three angles. To see this in the case of E_2 , we can use the sine rule $\|\ell_i\|/\|\ell_j\| = \sin \alpha_i/\sin \alpha_j$ to eliminate the l_i , and we get

$$E_2(\psi|_T) = \frac{\cot \alpha_1 \sin^2 \theta_1 + \cot \alpha_2 \sin^2 \theta_2 + \cot \alpha_3 \sin^2 \theta_3}{\sin \theta_1 \sin \theta_2 \sin \theta_3},$$

which can be simplified to

$$E_2(\psi|_T) = \sum_{\substack{i,j=1 \\ i \neq j}}^3 \cot \alpha_i \cot \theta_j,$$

which clearly depends only on the angles θ_i of $\psi(T)$ (in addition to the angles α_i of T).

Sheffer and de Sturler [28] have also exploited the fact that linear maps from one triangle to another are uniquely determined by the angles of their triangles if we ignore translations and uniform and orthogonal scalings. They therefore propose finding a mapping ψ in $S_2(\mathcal{S})$ which minimizes $F(\psi)$ in (10) where

$$E(\psi|_T) = \sum_{i=1}^3 (\theta_i - \alpha_i)^2.$$

Though this minimization problem is linear in the unknowns α_i , it becomes non-linear as a number of constraints (some of which are non-linear) have to be taken into account to guarantee the validity of the solution.

Other non-linear parameterization methods have been developed especially for the application of texture mapping in computer graphics. Bennis, Vézien, and Iglésias [1] proposed a method based on differential geometry. They map isoparametric curves of the surface onto curves in the parameter domain so that the geodesic curvature at each point is preserved. The parameterization is then extended to both sides of that initial curve until some distortion threshold is reached. Lévy and Mallet [22] presented a functional that attempts to preserve perpendicularity and constant spacing of the isoparametric curves traced on the surface. Zigelman, Kimmel, and Kiryati [33] use the method of multidimensional scaling to find parameter points such that the distances among them best approximate the geodesic distances between the corresponding data points in a least square sense. Sander et al. [27] minimize a non-linear ‘stretch’ metric when making parameterizations for texture mapping. The problem of constrained parameterization where a number of (interior) parameter points are fixed in advance, has been addressed in [6] and [21].

7 Remeshing, Subdivision Connectivity, and Multiresolution

Before looking at further aspects of parameterization, we will next look at how parameterization can be applied to multiresolution analysis of spatial triangulations. Many of the following ideas are taken from the work of Lounsbery, DeRose, and Warren [23], one of the first papers to deal with this topic.

We begin by explaining the idea of *remeshing*. Remeshing is simply the process of approximating a given triangulation by a new triangulation with certain desirable properties. Let us see how we might remesh a given simple triangulation \mathcal{S} . We begin by finding a parameterization for \mathcal{S} , i.e., we construct a one-to-one mapping $\psi \in S_2(\mathcal{S})$ (such as a convex combination mapping) so that $\psi(\mathcal{S})$ is the planar triangulation \mathcal{T} . We then retriangulate the domain $\Omega_{\mathcal{T}}$ by a new triangulation \mathcal{T}' in such a way that $\Omega_{\mathcal{T}'} = \Omega_{\mathcal{T}}$ (see Figure 17 (a) where \mathcal{T} is the triangulation with normal and \mathcal{T}' the triangulation with dotted lines). Now we define a mapping $\phi \in S_3(\mathcal{T}')$ by setting

$$\phi(v) = \psi^{-1}(v) \tag{11}$$

for all vertices v in \mathcal{T}' . Then the triangulation $\mathcal{S}' = \phi(\mathcal{T}')$ will be a piecewise linear approximation to the original triangulation \mathcal{S} : the vertices of \mathcal{S}' are sampled from $\Omega_{\mathcal{S}}$, and we have $\phi \approx \psi^{-1}$ (see Figure 17 (b)).

One desirable property we might demand of \mathcal{T}' is *subdivision connectivity* as it is the start point for a multiresolution analysis (MRA, see [26,2]), applications of which are data compression and multiresolution editing, among others. To understand subdivision connectivity, let \mathcal{T}^0 be a planar triangulation, and consider its dyadic refinement \mathcal{T}^1 . By dyadic refinement we mean

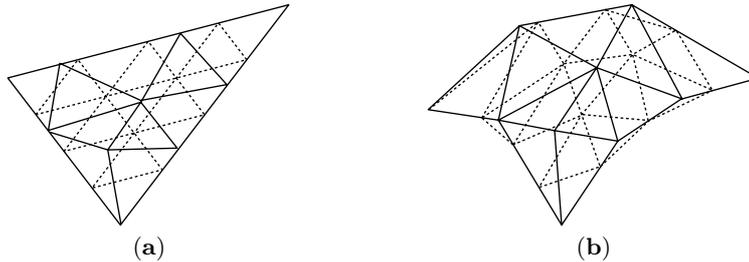


Fig. 17. Retriangulation of a planar (a) and a simple triangulation (b)

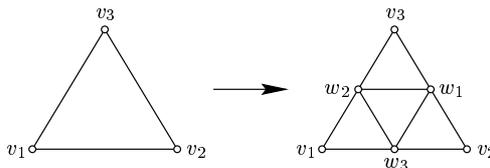


Fig. 18. Dyadic refinement.

that we divide each triangle $T = [v_1, v_2, v_3]$ in \mathcal{T}^0 into the four congruent subtriangles,

$$[v_1, w_2, w_3], \quad [w_1, v_2, w_3], \quad [w_1, w_2, v_3], \quad [w_1, w_2, w_3],$$

where

$$w_1 = \frac{v_2 + v_3}{2}, \quad w_2 = \frac{v_3 + v_1}{2}, \quad w_3 = \frac{v_1 + v_2}{2},$$

are the midpoints of the edges of T . This refinement is also referred to as the one-to-four split, and the set of all such subtriangles forms a triangulation \mathcal{T}^1 , see Figure 18 for an example.

Note that the triangles of \mathcal{T}^0 and \mathcal{T}^1 cover the same domain, i.e., $\Omega_{\mathcal{T}^0} = \Omega_{\mathcal{T}^1}$. Similarly, we can refine \mathcal{T}^1 to form \mathcal{T}^2 , and so on. We let V^j be the set of vertices in \mathcal{T}^j , $j = 0, 1, 2, \dots$, and E^j the set of edges. Thus the V^j are nested,

$$V^0 \subset V^1 \subset V^2 \subset \dots$$

We also have $\Omega_{\mathcal{T}^j} = \Omega_{\mathcal{T}^0}$ for all $j = 0, 1, \dots$

Definition 5. We will say that the triangulation \mathcal{T}^j above has **subdivision connectivity** of order j . Similarly, we will say that any simple triangulation $\mathcal{S} = \phi(\mathcal{T}^j)$, induced by a one-to-one mapping $\phi \in S_3(\mathcal{T}^j)$ has subdivision connectivity of order j .

Thus in order to remesh (approximate) a given simple triangulation \mathcal{S} with a triangulation \mathcal{S}^j with subdivision connectivity of order j , we apply the remeshing procedure given by (11) and take \mathcal{T}' to be a triangulation \mathcal{T}^j with subdivision connectivity of order j . In practice we would first retriangulate

$\Omega_{\mathcal{T}}$ by a triangulation \mathcal{T}^0 with relatively few triangles and then let \mathcal{T}^j be its j -fold dyadic refinement.

To briefly see how an MRA works in this setting, let $S^j = S_3(\mathcal{T}^j)$ for $j \geq 0$. These linear spaces are clearly nested,

$$S^0 \subset S^1 \subset S^2 \subset \dots,$$

and this allows us to decompose the triangulation $\mathcal{S}^j = \phi(\mathcal{T}^j)$ by simply decomposing the piecewise linear mapping $\phi \in S_3(\mathcal{T}^j)$ which defines it.

For each $j = 1, 2, \dots$, we choose W^{j-1} to be some subspace of S^j with the property that

$$S^{j-1} \oplus W^{j-1} = S^j,$$

where \oplus denotes in general a direct sum. We call W^{j-1} a *wavelet space* of S^{j-1} in S^j , and we can decompose any space S^j into its coarsest subspace S^0 and a sequence of wavelet spaces,

$$\begin{aligned} S^j &= S^{j-1} \oplus W^{j-1} \\ &= S^{j-2} \oplus W^{j-2} \oplus W^{j-1} \\ &\vdots \\ &= S^0 \oplus W^0 \oplus \dots \oplus W^{j-1}. \end{aligned} \tag{12}$$

Within this framework, we decompose a given $f \in S_3(\mathcal{T}^j)$ into *levels of detail*. Equation (12) implies that there exist unique functions $f^{j-1} \in S^{j-1}$ and $g^{j-1} \in W^{j-1}$ such that

$$f^j = f^{j-1} + g^{j-1},$$

and we regard f^{j-1} as an approximation to f^j at a lower resolution or level of detail. The function g^{j-1} is the error introduced when replacing the original function f^j by its approximation f^{j-1} . We can continue this decomposition until $j = 0$. Then we will have $f^0 \in S^0$ and $g^i \in W^i$, $i = 0, 1, \dots, j-1$ such that

$$f^j = f^0 + g^0 + g^1 + \dots + g^{j-1},$$

and f^0 will be the coarsest possible approximation to f^j . Multiresolution editing can now be carried out by altering one or more of the functions g^k . Data compression algorithms can be designed by introducing bases for the spaces S^k and W^k and throwing away small wavelet coefficients. Various approaches to suitable wavelet bases have been proposed, such as those of [23], [30], and [12], and in the references therein.

A numerical example of remeshing and wavelet decomposition applied (in an analogous way) to a manifold triangulation will be given in Section 9.

8 Parameterization of Unorganized Points

So far we have assumed that the given data points \mathbf{x}_i are organized in a triangulation, so that neighbourhood information is given. We will suppose now that no such triangulation is given and we wish to find suitable parameter points \mathbf{u}_i without the need for given neighbourhood information. For simplicity we will only treat the case that the points \mathbf{x}_i are sampled from a simple surface, i.e. a surface which can be parameterized over a simply-connected domain in \mathbb{R}^2 . Thus we want to construct a discrete one-to-one mapping ψ_N from the vector $X_N = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ of points in \mathbb{R}^3 , to another vector $U_N = \{\mathbf{u}_1, \dots, \mathbf{u}_N\}$ of points in \mathbb{R}^2 . Then the inverse $\phi_N = \psi_N^{-1} : U_N \rightarrow X_N$ will be a discrete parameterization of the point set X_N . Since the discrete mapping ψ_N is independent of any given topological structure amongst the \mathbf{x}_i , we call the construction of ψ_N *meshless parameterization*.

However, in order to guide the choice of ψ_N we have in mind that if the points \mathbf{x}_i were the vertices of a triangulation \mathcal{S} then ψ_N should be the restriction to the vertices of \mathcal{S} of some one-to-one piecewise linear map $\psi \in S_2(\mathcal{S})$. This motivates taking ψ_N to be a convex combination map, just as for triangulated data. The only change is that we now have to construct neighbourhoods from data points which are near to each given data point \mathbf{x}_i .

Note that in addition to the application of meshless parameterization to scattered data approximation as in (1), it can also be used to triangulate the data set \mathbf{x}_i with a simple triangulation. One can simply triangulate the parameter points \mathbf{u}_i with a triangulation \mathcal{T} say, which could for example be a Delaunay triangulation (see Section 6.1 of [17]), and this immediately gives a corresponding triangulation \mathcal{S} of the points \mathbf{x}_i . In other words, we take \mathcal{S} to be the set of all triangles $[\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k]$ for which $[\mathbf{u}_i, \mathbf{u}_j, \mathbf{u}_k]$ is a triangle in \mathcal{T} . Thus meshless parameterization can be used for surface reconstruction.

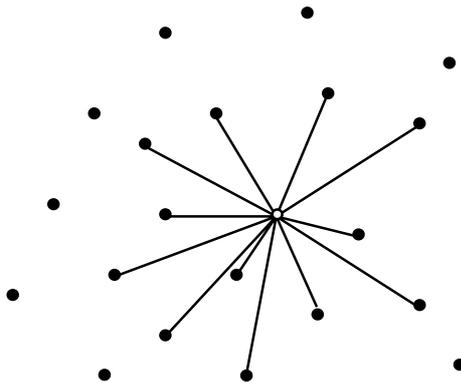


Fig. 19. 12 nearest neighbours

The method, described in [13], begins by first dividing the set X_N into two disjoint subsets: X_I , the set of interior points, and X_B , the set of boundary points. Moreover the boundary points must be ordered. A method was outlined in [13] for first identifying boundary points and subsequently ordering them using a univariate analog of meshless parameterization. After that the method is identical to the convex combination mapping method of Section 3, the only difference being that we now have to find suitable neighbourhoods $N_{\mathbf{x}_i}$ for the points $\mathbf{x}_i \in X_I$. Several choices of such neighbourhoods were proposed in [13] but a simple and effective choice is to take the ‘ d nearest neighbours’, in other words, we let $N_{\mathbf{x}_i}$ be the set of the d points \mathbf{x}_j closest to \mathbf{x}_i , illustrated in Figure 19. Going on several numerical tests, it appears that setting $d = 10$ or $d = 20$ is adequate for all but the most extreme data sets. As regards the choice of weights λ_{ij} , the naive choice of uniform weights $\lambda_{ij} = 1/d_i$, where $d_i = |N_i|$, can lead to the result that two data points \mathbf{x}_i and \mathbf{x}_j end up being mapped to the same points: $\mathbf{u}_i = \mathbf{u}_j$, as shown in [13]. A better choice which has been found to perform very well in practice is to use the weights

$$\lambda_{\mathbf{x}_i, \mathbf{x}_j} = \frac{1}{\|\mathbf{x}_j - \mathbf{x}_i\|} \bigg/ \sum_{\mathbf{x}_k \in N_{\mathbf{x}_i}} \frac{1}{\|\mathbf{x}_k - \mathbf{x}_i\|},$$

which depend on the distances between \mathbf{x}_i and its neighbours. Notice that the neighbourhoods $N_{\mathbf{x}_i}$ induce a directed graph G on the points \mathbf{x}_i , where the points \mathbf{x}_i are the nodes and the pair $(\mathbf{x}_i, \mathbf{x}_j)$ is a directed edge whenever $\mathbf{x}_j \in N_{\mathbf{x}_i}$. This graph will rarely be a planar graph (as is a triangulation) and will usually have many more edges. Even so this method of parameterization has been shown to work extremely well on many data sets. Notice also that the graph G determines the solvability of the linear system (3). Indeed there is a unique solution if for every point $\mathbf{x}_i \in X_I$, there is at least one path of points \mathbf{x}_j in G connecting \mathbf{x}_i to a boundary node in X_B .

Figure 20 shows an example of unorganized points (in \mathbb{R}^3) and a convex combination mapping of these points to a rectangular parameter domain. Here, the natural boundary of the data points is quite evident from the figure. In practice, finding a suitable boundary can be quite delicate: a method was outlined in [13]. Figure 21 shows the Delaunay triangulation of the parameter points and this leads to a reasonable triangulation of the original data points. Figure 22 illustrates the method on a more realistic example, a foot data set, taken from [13]. On the left is the meshless parameterization, after Delaunay triangulation, and the lifted triangulation is shown on the right.

By combining the parameterization of unorganized points with that of triangulations, we end up with a fast, linear method for approximating unorganized points by spline surfaces: we can go from Figure 20 (a) to Figure 21 (b), which is the same as the triangulation in Figure 3, and from there find Figure 4 (b). Further numerical examples of meshless parameterization can be found in [13] and [9].

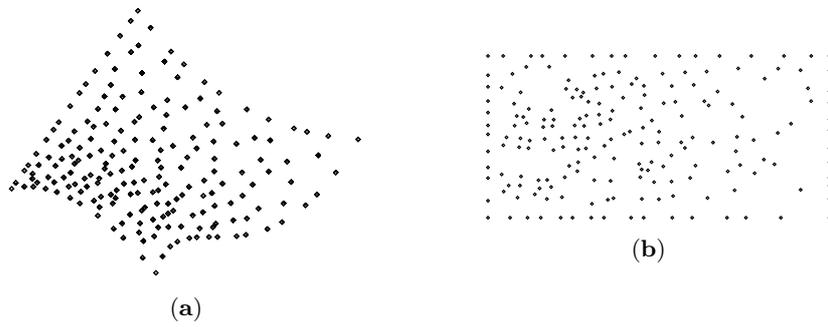


Fig. 20. Unorganized points (a) and resulting parameterization (b)

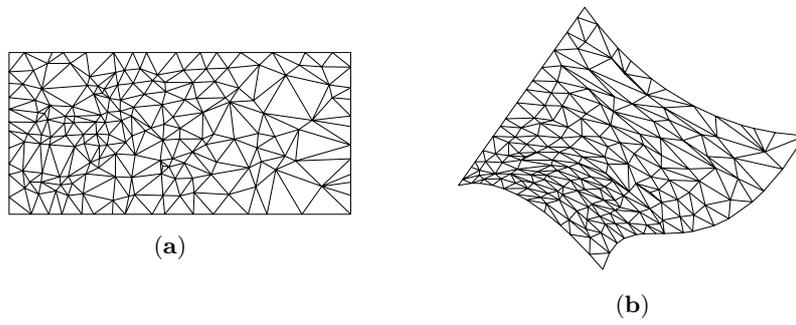


Fig. 21. Delaunay triangulation (a) and resulting spatial triangulation (b)

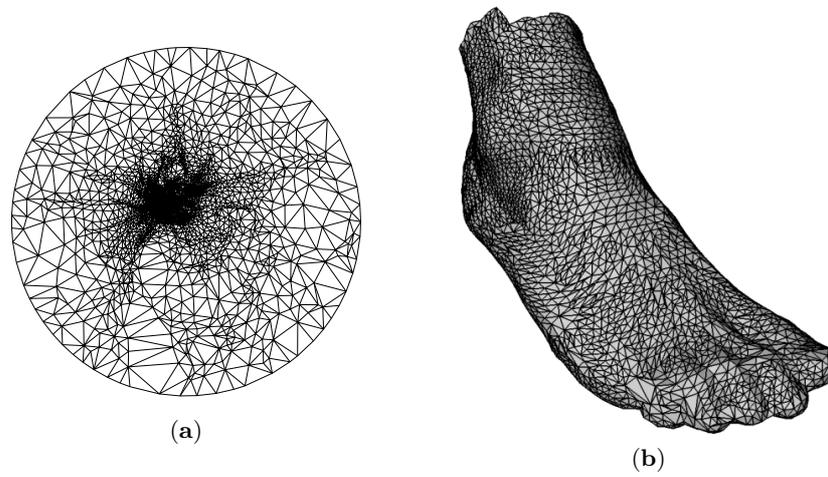


Fig. 22. Meshless parameterization and (a) and resulting triangulation (b)

9 Parameterization of Manifold Triangulations

In previous sections we have considered the parameterization of *simple* triangulations, but what is to be done if the triangulation is more complicated, for example closed and homeomorphic to a sphere or torus? This last section shows how convex combination maps can be used piecewise to parameterize triangulations of arbitrary topology. Though the basic ideas can be extended to triangulations with boundaries, we will restrict the discussion to manifolds (i.e. closed surfaces).

Definition 6. Let $\mathcal{T} = \{T_1, \dots, T_M\}$ be a set of triangles in \mathbb{R}^3 . We call \mathcal{T} a *manifold triangulation* if

- (i) $T_i \cap T_j$ is either empty, a common vertex or a common edge, $i \neq j$, and
- (ii) the union of the triangles $\Omega_{\mathcal{T}} = \bigcup_{i=1}^M T_i$ is an orientable 2-manifold.

A manifold triangulation \mathcal{T} can no longer be parameterized over a planar parameter domain as in Section 2 and we rather need a spatial domain that is topologically equivalent to $\Omega_{\mathcal{T}}$.

One way of generating a suitable spatial parameter domain is by *mesh decimation* or *mesh simplification* (see [15]). The major class of these algorithms modifies a given triangulation by iteratively applying a topological operator such as *vertex removal* or *half-edge collapse* that decreases the number of triangles by two in each step without changing the topology of the triangulation's surface. These algorithms differ mainly by the strategy that is chosen to decide in which order the simplification steps are to be carried out. They all result in a triangulation \mathcal{T}^0 that approximates the shape of the initial triangulation \mathcal{T} in some way, i.e. $\Omega_{\mathcal{T}^0} \approx \Omega_{\mathcal{T}}$. An example is shown in Figure 23 where a triangulation with 768 triangles (**a**) was simplified to one with 100 triangles (**b**).

Another approach for creating \mathcal{T}^0 is based on growing Voronoi tiles around a previously chosen set of *site faces* and constructing the dual Delaunay Triangulation and was introduced in [5].

A *parameterization* of the given manifold triangulation \mathcal{T} over the spatial domain \mathcal{T}^0 is a homeomorphism $\phi : \Omega_{\mathcal{T}^0} \rightarrow \Omega_{\mathcal{T}}$.

Lee et al. [20] find such a parameterization by simultaneously creating a hierarchy $\mathcal{T} = \mathcal{T}^0, \dots, \mathcal{T}^n$ of triangulations, and parameterizations $\phi^i : \mathcal{T}^i \rightarrow \mathcal{T}^{i-1}$. In each decimation step they remove one vertex of \mathcal{T}^{i-1} and the triangles incident to this vertex and retriangulate the hole. Therefore, \mathcal{T}^{i-1} and \mathcal{T}^i differ only locally and ϕ^i is not hard to find. Composition of the individual ϕ^i finally yields a parameterization $\phi : \mathcal{T}^n \rightarrow \mathcal{T}$, $\phi = \phi^1 \circ \dots \circ \phi^n$ which they further smooth in order to obtain better remeshing results.

An alternative and way of solving the manifold parameterization problem is described in [11]. The method partitions the manifold triangulation \mathcal{T} into $M_0 = |\mathcal{T}^0|$ disjoint subsets $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_{M_0}$ such that each \mathcal{S}_i is a

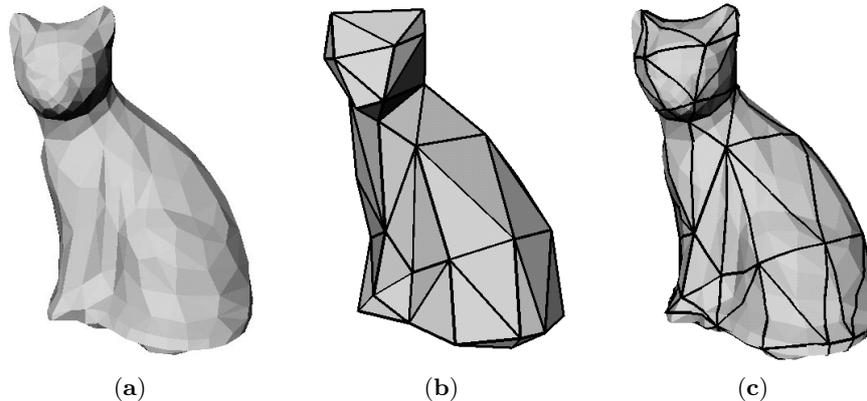


Fig. 23. A manifold triangulation (a), its simplified triangulation (b), and its surface triangulation (c)

simple triangulation and corresponds to the triangle $T_i \in \mathcal{T}^0$ in the spatial parameter domain. We call each \mathcal{S}_i a *surface triangle* and the collection $\mathcal{S} = \{\mathcal{S}_i\}_{i=1, \dots, M_0}$ a *surface triangulation*. Figure 23 (c) shows an example.

We can then use any of the simple parameterization methods to construct a parameterization ϕ_i of \mathcal{S}_i over T_i and combine them in the end to define the parameterization ϕ by letting

$$\phi|_{T_i} = \phi_i, \quad i = 1, \dots, M^0.$$

In order to assure ϕ to be a homeomorphism we need the individual parameterizations ϕ_i to be continuous across the edges of \mathcal{T}^0 . This is achieved by parameterizing the edges first and then fixing these parameter values before parameterizing the interior of each surface triangle.

In practice the whole procedure is performed in three steps. Firstly, ϕ is defined for the vertices V^0 of \mathcal{T}^0 . If \mathcal{T}^0 was constructed by mesh simplification we always have $V^0 \subset V$ and simply set $\phi(\mathbf{v}) = \mathbf{v}$, $\mathbf{v} \in V^0$. If the vertex correspondences are chosen by some other method, e.g. interactively, $\phi(\mathbf{v})$ does not necessarily need to be a vertex of \mathcal{T} . In this case we refine \mathcal{T} by an edge or a one-to-three triangle split to a triangle mesh \mathcal{T}' such that $\phi(\mathbf{v})$ is a vertex of \mathcal{T}' (cf. Figure 24). Note that this does not change the geometry of the given triangulation, i.e. $\Omega_{\mathcal{T}} = \Omega_{\mathcal{T}'}$, but only its connectivity. We therefore assume without loss of generality $\phi(V^0) \subset V$.

Secondly, ϕ is extended to the edges E^0 of \mathcal{T}^0 . This is done by finding for each edge $e = [\mathbf{v}, \mathbf{w}] \in E^0$ the shortest path \bar{e} on $\Omega_{\mathcal{T}}$ between $\phi(\mathbf{v})$ and $\phi(\mathbf{w})$, e.g. with the algorithm in [4,18]. We refer to this shortest path as a *surface edge*. Since the shortest path across a triangle is a straight line, shortest paths on triangulations are polygonal curves whose vertices lie on the edges of the triangulation. Thus we may write \bar{e} as a set of line segments

$$\bar{e} = [\mathbf{u}_1, \mathbf{u}_2] \cup [\mathbf{u}_2, \mathbf{u}_3] \cup \dots \cup [\mathbf{u}_{r-1}, \mathbf{u}_r],$$

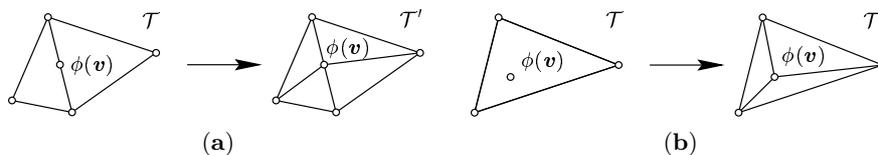


Fig. 24. Refining \mathcal{T} by an edge split (a) or a one-to-three triangle split (b)

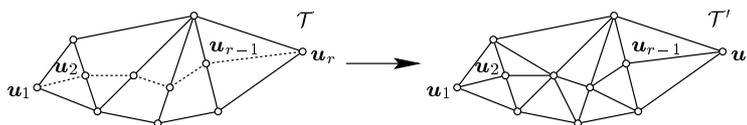


Fig. 25. Refining \mathcal{T} in order to embed a surface edge \bar{e}

where $\mathbf{u}_1 = \phi(\mathbf{v})$, $\mathbf{u}_r = \phi(\mathbf{w})$ and the \mathbf{u}_i are located on the edges of \mathcal{T} . Since we can always refine \mathcal{T} to a triangle mesh \mathcal{T}' as shown in Figure 25 we assume without loss of generality $\mathbf{u}_i \in V$ and $[\mathbf{u}_i, \mathbf{u}_{i+1}] \in E$. We can now parameterize \bar{e} over e by a standard line parameterization technique, e.g. chord length parameterization.

Thirdly, we further extend ϕ to the triangles of \mathcal{T}^0 . For each triangle $T_i = [\mathbf{u}, \mathbf{v}, \mathbf{w}] \in \mathcal{T}^0$ we let the surface triangle \mathcal{S}_i be that region of \mathcal{T} which is bounded by the three surface edges $\phi([\mathbf{u}, \mathbf{v}])$, $\phi([\mathbf{v}, \mathbf{w}])$, $\phi([\mathbf{w}, \mathbf{u}])$. As we have already defined the parameter points of the boundary vertices of \mathcal{S}_i in the previous step we can now apply e.g. shape preserving parameterization to determine the parameter points of the interior vertices. In this way we obtain for each $i = 1, \dots, M^0$ the individual parameterization $\phi_i : T_i \rightarrow \Omega_{\mathcal{S}_i}$ which are then combined to give $\phi : \Omega_{\mathcal{T}^0} \rightarrow \Omega_{\mathcal{T}}$.

Figure 26 finally shows an example of a manifold parameterization and its application to remeshing and wavelet compression. The manifold triangulation \mathcal{T} in (a) has first been decimated to the coarse triangulation \mathcal{T}^0 in (d). Then the parameterization $\phi : \Omega_{\mathcal{T}^0} \rightarrow \mathcal{T}$ has been constructed as explained above (b) and used to define a remesh \mathcal{T}' (c) of \mathcal{T} by mapping the vertices of the dyadic refinement \mathcal{T}^3 of \mathcal{T}^0 (e) as in (11). Finally, the wavelets of [12] were used to decompose and compress \mathcal{T}' to the triangulation in (f) which can be represented by the 59 vertices of a coarse triangulation and 59 wavelet coefficients (vectors) opposed to the 6475 vertices that define the given triangulation \mathcal{T} , i.e., by less than 2% of the initial information.

References

1. C. Bennis, J.-M. Vézien, and G. Iglésias. Piecewise surface flattening for non-distorted texture mapping. *Proceedings of SIGGRAPH*, 1991, 237–246.
2. G.-P. Bonneau. BLaC Wavelets & Non-Nested Wavelets. This volume, Chapter 6.

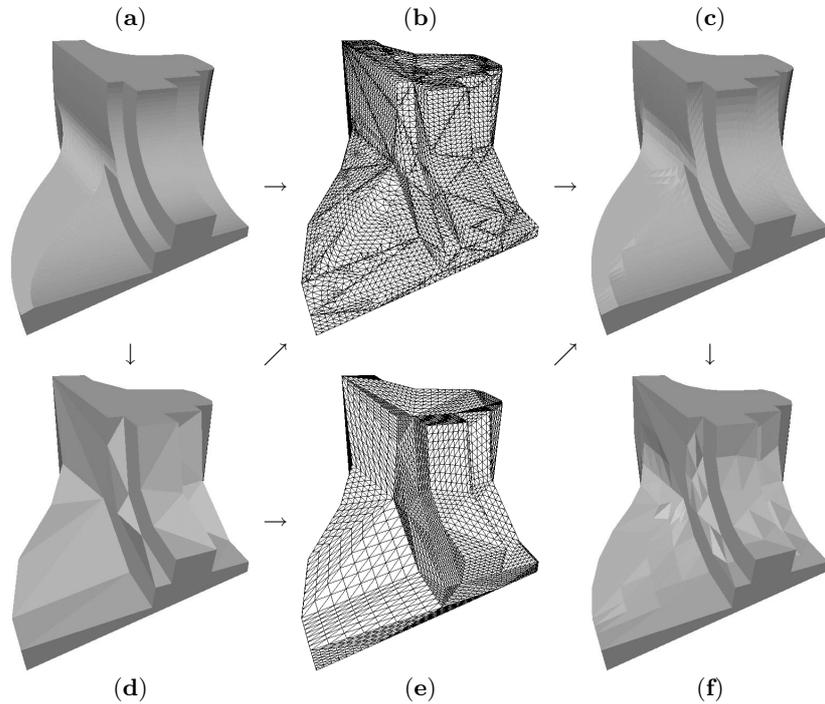


Fig. 26. Decimation (d), parameterization (b), remesh (c), and compression (f) of a given manifold triangulation (a)

3. M. P. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, New Jersey, 1976.
4. J. Chen and Y. Han. Shortest paths on a polyhedron. *Intern. J. Computational Geometry & Applications* **6**, 1996, 127–144.
5. M. Eck, T. DeRose, T. Duchamp, H. Hoppe, M. Lounsbery, and W. Stuetzle. Multiresolution analysis of arbitrary meshes. *Proceedings of SIGGRAPH*, 1995, 173–182.
6. I. Eckstein, V. Surazhsky, and C. Gotsman. Texture mapping with hard constraints. *Proceedings of Eurographics*, 2001.
7. M. S. Floater. Parameterization and smooth approximation of surface triangulations. *Comput. Aided Geom. Design* **14**, 1997, 231–250.
8. M. S. Floater. Parametric tilings and scattered data approximation. *Intern. J. Shape Modeling* **4**, 1998, 165–182.
9. M. S. Floater. Meshless parameterization and B-spline surface approximation. *The Mathematics of Surfaces IX*, R. Cipolla and R. Martin R. (eds.), Springer, 2000, 1–18.
10. M. S. Floater. One-to-one piecewise linear mappings over triangulations. To appear in *Math. Comp.*
11. M. S. Floater, K. Hormann, and M. Reimers. Parameterization of manifold triangulations. To appear in *Approximation Theory X*, Charles K. Chui and Larry L. Schumaker (eds.).

12. M. S. Floater and E. G. Quak. Piecewise linear prewavelets on arbitrary triangulations. *Numer. Math.* **82**, 1999, 221–252.
13. M. S. Floater and M. Reimers. Meshless parameterization and surface reconstruction. *Comput. Aided Geom. Design* **18**, 2001, 77–92.
14. G. H. Golub and C. F. van Loan. *Matrix Computations*. The John Hopkins University Press, Baltimore, 1989.
15. C. Gotsman, S. Gumhold, and L. Kobbelt. Simplification and Compression of 3D Meshes. This volume, Chapter 11.
16. K. Hormann and G. Greiner. MIPS: an efficient global parametrization method. *Curve and Surface Design: Saint-Malo 1999*, Pierre-Jean Laurent, Paul Sablonnière, and Larry L. Schumaker (eds.), Vanderbilt University Press, Nashville, 2000, 153–162.
17. A. Iske. Scattered Data Modelling using Radial Basis Functions. This volume, Chapter 8.
18. B. Kaneva and J. O'Rourke. An implementation of Chen & Han's shortest paths algorithm. *Proceedings of the 12th Canadian Conference on Computational Geometry*, 2000.
19. A. N. Kolmogorov and S. V. Fomin. *Introductory Real Analysis*. Dover Publications, New York, 1975.
20. A. W. F. Lee, W. Sweldens, P. Schröder, L. Cowsar, and D. Dobkin. MAPS: Multiresolution adaptive parameterization of surfaces. *Proceedings of SIGGRAPH*, 1998, 95–104.
21. B. Lévy. Constrained texture mapping for polygonal meshes. *Proceedings of SIGGRAPH*, 2001, 417–424.
22. B. Lévy and J.-L. Mallet. Non-distorted texture mapping for sheared triangulated meshes. *Proceedings of SIGGRAPH*, 1998, 343–352.
23. M. Lounsbery, T. DeRose, and J. Warren. Multiresolution analysis for surfaces of arbitrary topological type. *ACM Trans. on Graphics* **16**, 1997, 34–73.
24. J. Maillot, H. Yahia, and A. Verroust. Interactive texture mapping. *Proceedings of SIGGRAPH*, 1993, 27–34.
25. U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugates. *Exp. Math.* **2**, 1993, 15–36.
26. E. Quak. B-splines and B-wavelets. This volume, Chapter 5.
27. P. Sander, J. Snyder, S. Gortler, and H. Hoppe. Texture mapping progressive meshes. *Proceedings of SIGGRAPH*, 2001, 409–416.
28. A. Sheffer and E. de Sturler. Surface parameterization for meshing by triangulation flattening. *Proceedings of the International Meshing Round Table Conference*, 2000.
29. L. L. Schumaker. Fitting surfaces to scattered data. *Approximation Theory, II*, G. G. Lorentz, C. K. Chui, and L. L. Schumaker (eds.), Academic Press, New York, 1976, 203–268.
30. P. Schröder and W. Sweldens. Spherical wavelets: efficiently representing functions on the sphere. *Proceedings of SIGGRAPH*, 1995, 161–172.
31. R. S. Varga. *Matrix Iterative Analysis*. Springer, 2000.
32. F. Zeilfelder. Scattered Data Fitting with Bivariate Splines. This volume, Chapter 9.
33. G. Zigelman, R. Kimmel, and N. Kiryati. Texture mapping using surface flattening via multi-dimensional scaling. *IEEE Trans. Visualization and Computer Graphics*, 2001.