

Mesh Parameterization Methods and their Applications

Alla Sheffer
University of British Columbia

Emil Praun
Google

Kenneth Rose
University of British Columbia

Abstract

We present a survey of recent methods for creating piecewise linear mappings between triangulations in 3D and simpler domains such as planar regions, simplicial complexes, and spheres. We also discuss emerging tools such as global parameterization, inter-surface mapping and parameterization with constraints. We start by describing the wide range of applications where parameterization tools have been used in recent years. We then briefly review the pertinent mathematical background and terminology, before proceeding to survey the existing parameterization techniques. Our survey summarizes the main ideas of each technique and discusses its main properties, comparing it to other methods available. Thus it aims to provide guidance to researchers and developers when assessing the suitability of different methods for various applications. This survey focuses on the practical aspects of the methods available, such as time complexity and robustness and shows multiple examples of parameterizations generated using different methods, allowing the reader to visually evaluate and compare the results.

1 Introduction

Given any two surfaces with similar topology it is possible to compute a one-to-one and onto mapping between them. If one of these surfaces is represented by a triangular mesh, the problem of computing such a mapping is referred to as mesh parameterization [Bennis et al. 1991; Floater and Hormann 2005]. The surface that the mesh is mapped to is typically referred to as the parameter domain. Parameterizations between surface meshes and a variety of domains have numerous applications in computer graphics and geometry processing as described below. In recent years numerous methods for parameterizing meshes were developed, targeting diverse parameter domains and focusing on different parameterization properties. This survey reviews the various parameterization methods, summarizing the main ideas of each technique and focusing on the practical aspects of the methods. It also provides examples of the results generated by many of the more popular methods. When several methods address the same parameterization problem, the survey strives to provide an objective comparison between them based on criteria such as parameterization quality, efficiency and robustness.

We start by surveying the applications which can benefit from parameterization in Section 1.1 and then in Section 2 briefly review the terminology commonly used in the parameterization literature. The rest of the survey describes the different techniques available, classifying them based on the parameter domain used. Section 3 describes techniques for planar parameterization. Section 4 reviews methods for pre-processing meshes for planar parameterization by cutting them into one or more charts. Section 5 examines parameterization methods for alternative domains such as a sphere or a base mesh as well as methods for cross-parameterization between mesh surfaces. Section 6 discusses methods to introduce constraints into a parameterization.

Finally, Section 7 summarizes the paper and discusses potential open problems in mesh parameterization

1.1 Applications

Surface parameterization was introduced to computer graphics as a method for mapping textures onto surfaces [Bennis et al. 1991; Maillot et al. 1993]. Over the last decade, it has gradually become a ubiquitous tool, useful for many mesh processing applications, discussed below (Figure 1).


Detail Mapping: Detailed objects can be efficiently represented by a coarse geometric shape (polygonal mesh or subdivision surface) with the details corresponding to each triangle stored in a separate 2D array. In traditional texture mapping the detail is the local albedo of a Lambertian surface. Texture maps alone can enrich the appearance of a surface in a static picture, but since neighboring pixels will have similar shadowing, objects may still look flat in animations with varying lighting conditions. Bump mapping stores small deviations of the point-wise normal from that of the smooth underlying surface and uses the perturbed version during shading [Blinn 1978]. Normal mapping [Wikipedia, Sheffer et al. 2005] is a similar technique that replaces the normals directly rather than storing a perturbation. As the light direction changes, the shading variations produced by the normal perturbations simulate the shadows caused by small pits and dimples in the surface. Since the actual geometry of the object is not modified, the silhouettes still look polygonal or smooth. Displacement mapping addresses this problem by storing small local deformations of the surface, typically in the direction of the normal. Recent techniques [Lengyel et al. 2001, Peng et al. 2004, Porumbescu et al. 2005] model a thick region of space in the neighborhood of the surface by using a volumetric texture, rather than a 2D one. Such techniques are needed in order to model detail with complicated topology or detail that cannot be easily approximated locally by a height field, such as sparsely interwoven structures or animal fur. The natural way to map details to surfaces is using planar parameterization (Section 3).

Detail Synthesis: While the goal of texture mapping is to *represent* the complicated appearance of 3D objects, several methods make use of mesh parameterization to *create* the local detail necessary for a rich appearance. Such techniques can use as input flat patches with sample detail [Pederson 1995, Praun et al. 2000, Wei and Levoy 2001, Turk 2001, Ying et al. 2001, Soler et al. 2002]; parametric or procedural models; or direct user input and editing [Igarashi and Cosgrove 2001, Carr and Hart 2004]. The type of detail can be quite varied and the intermediate representations used to create it parallel the final representations used to store it.

Morphing and Detail Transfer: A map between the surfaces of two objects allows the transfer of detail from one object to another [Praun et al. 2001, Sorkine et al. 2004, Litke et al., 2005], or the interpolation between the shape and appearance of several objects [Lee et al. 1999, Alexa 2000, Alexa 2002, Kraevoy and Sheffer 2004; Schreiner et al. 2004]. By varying the interpolation ratios over time, one can produce morphing animations. In spatially-varying and frequency-varying morphs, the rate of change can be different for different parts of the objects, or different frequency bands (coarseness of the features being transformed) [Lee et al. 1999, Allen et al. 2003; Kraevoy and Sheffer, 2004]. Such a map can either be computed directly or, as more commonly done, computed by mapping both object surfaces to a common domain (Sections 5 and 6).




Texture Mapping [Lévy 2001]




Normal Mapping [Sheffer et al., 2005]




Detail Transfer [Biermann et al. 2002]



Morphing [Lee et al. 1999]




Mesh Completion [Kraevoy and Sheffer 2005]




Editing [Lévy 2003]



Databases [Allen et al. 2003]



Remeshing [Praun and Hoppe 2003]



Surface Fitting [Li et al. 2006]

Figure 1. Parameterization Applications.

In addition to transferring the static appearance of surfaces, inter-surface parameterizations allow the transfer of animation data between shapes, either by transferring the local surface influence from bones of an animation rig, or by directly transferring the local affine transformation of each triangle in the mesh [Sumner and Popovic 2004].

Mesh Completion: Meshes from range scans often contain holes and multiple components. Lévy [2003] uses planar parameterization to obtain the natural shape for hole boundaries and to triangulate those. In many cases, prior knowledge about the overall shape of the scanned models exists. For instance, for human scans, templates of a generic human shape are readily available. Allen et al. [2003], and Anguelov et al. [2005] use this prior knowledge to facilitate completion of scans by computing a mapping between the scan and a template human model. Kraevoy and Sheffer [2005] develop a more generic and robust template-based approach for completion of any type of scans. The techniques typically use an inter-surface parameterization between the template and the scan (Sections 5 and 6).

Mesh Editing: Editing operations often benefit from a local parameterization between pairs of models. Biermann et al. [2002] use local parameterization to facilitate cut-and-paste transfer of

details between models. They locally parameterize the regions of interest on the two models in 2D and overlap the two parameterizations. They use the parameterization to transfer shape properties from one model to the other. Sorkine et al. [2004] and Lévy [2003] use local parameterization for mesh composition in a similar manner. They compute an overlapping planar parameterization of the regions near the composition boundary on the input models and use it to extract and smoothly blend shape information from the two models.

Creation of Object Databases: Once a large number of models are parameterized on a common domain (Sections 5 and 6) one can perform an analysis determining the common factors between objects and their distinguishing traits. For example on a database of human shapes [Allen et al. 2003] the distinguishing traits may be gender, height, and weight, while a database of human faces may add facial expressions [Blanz and Vetter 1999, Marschner et al. 2000, Blanz et al. 2003, Blanz et al. 2004]. Objects can be compared against the database and scored against each of these dimensions, and the database can be used to create new plausible object instances by interpolation or extrapolation of existing ones.

Remeshing: There are many possible triangulations that represent the same shape with similar levels of accuracy. Some triangulation may be more desirable than others for different applications. For example, for numerical simulations on surfaces, triangles with a good aspect ratio (that are not too small or too “skinny”) are important for convergence and numerical accuracy. One common way to remesh surfaces, or to replace one triangulation by another, is to parameterize the surface, then map a desirable, well-understood, and easy to create triangulation of the domain back to the original surface. For example, Gu et al. [2002] use a regular grid sampling of a planar square domain, while subdivision based methods [Guskov et al. 2000, Lee et al. 2000, Khodakovsky et al. 2003] use regular subdivision (usually 1-to-4 triangle splits) on the faces of a simplicial domain. Such locally regular meshes can usually support the creation of smooth surfaces as the limit process of applying subdivision rules. To generate high quality triangulations Desbrun et al. [2002] parameterize the input mesh in the plane and then use planar Delaunay triangulation to obtain a high quality remeshing of the surface. One problem these methods face is the appearance of visible discontinuities along the cuts created to facilitate the parameterization.

Surazhsky and Gotsman [2003] avoid global parameterization, and instead use local parameterization to move vertices along the mesh as part of an explicit remeshing scheme. Ray et al. [2005; 2006] use global periodic parameterization to generate a predominantly quadrilateral mesh directly on the 3D surface. Dong et al. [2006] use a parameterization induced by the Morse complex to generate a quad only mesh of the surface.

More details on the use of parameterization for remeshing can be found in a recent survey by Alliez et al. [2005].

Mesh Compression: Mesh compression is used to compactly store or transmit geometric models [Alliez and Gotsman 2005]. As with other data, compression rates are inversely proportional to the data entropy. Thus higher compression rates can be obtained when models are represented by meshes that are as regular as possible, both topologically and geometrically. Topological regularity refers to meshes where almost all vertices have the same degree. Geometric regularity implies that triangles are similar to each other in terms of shape and size and vertices are close to the centroid of their neighbors. Such meshes can be obtained by parameterizing the original objects and then remeshing with regular sampling patterns [Gu et al. 2002, Hoppe and Praun 2005]. The quality of the parameterization directly impacts the compression efficiency.

Surface Fitting: One of the earlier applications of mesh parameterization is surface fitting [Floater 1997; Hormann and Greiner 2000; Losasso et al. 2003; He et al. 2005; Li et al., 2006]. Many applications in geometry processing require a smooth analytical surface to be constructed from an input mesh. A parameterization of the mesh over a base domain significantly simplifies this task. Earlier methods either parameterized the entire mesh in the plane [Floater 1997] or segmented it and parameterized each patch independently (Sections 3 and 4). More recent methods [Losasso et al. 2003; Gu et al. 2005; Li et al., 2006] focus on constructing smooth global parameterizations (Section 5.1) and use those for fitting, achieving global continuity of the constructed surfaces.

Modeling from Material Sheets: While computer graphics focuses on virtual models, geometry processing has numerous real-world engineering applications. Particularly, planar mesh parameterization is an important tool when modeling 3D objects from sheets of material, ranging from garment modeling to metal forming or forging [Bennis et al. 1991; Mitani and Suzuki 2004; McCartney et al. 1999; Julius et al. 2005]. All of these applications require the computation of planar patterns to form the desired 3D shapes. Typically, models are first segmented into nearly developable charts (Section 4), and these charts are then parameterized in the plane (Section 3).

Medical Visualization: Complex geometric structures are often better visualized and analyzed by mapping the surface normal-map, color, and other properties to a simpler, canonical domain. One of the structures for which such mapping is particularly useful is the human brain [Hurdal et al. 1999; Haker et al. 2000; Gu et al. 2004b]. Most methods for brain mapping use the fact that the brain has genus zero, and visualize it through spherical [Haker et al. 2000; Gu et al. 2004b] (Section 5.2) or planar [Hurdal et al. 1999] (Section 3) parameterization.

Given the vast range of processing techniques that have benefited from parameterization, we expect that many more applications can utilize it as a powerful processing tool.

2 Terminology

Before proceeding to describe various parameterization techniques in the next section, we first briefly establish some terminology. We are concerned with the parameterization of triangle meshes. The topology of such meshes is typically represented as a simplicial complex: a set of 1, 2 and 3-element subsets of a set V of labels, corresponding respectively to the vertices, edges, and triangles of the mesh. The geometry of the mesh is represented as 3D coordinates associated with each of the vertices $c:V \rightarrow \mathbb{R}^3$, making edges correspond to (open) line segments in 3D, and mesh triangles to (open) triangles in 3D.

The purpose of mesh parameterization is to obtain a map between such a mesh and a triangulation of a domain. The map is piecewise linear, associating each triangle of the original mesh with a triangle in the domain. An important goal of parameterization is to obtain *bijective* (invertible) maps, where each point on the domain corresponds to exactly one point of the mesh. Many applications of planar parameterization can, with some modifications, handle global domain overlaps (Figure 2 (a)). For such applications the bijectivity requirement can be weakened, requiring only *local* rather than *global* bijectivity. Local bijectivity [Sheffer et al. 2005] requires a map of any sufficiently small region of the mesh to be bijective. This condition is violated when the mappings of adjacent mesh triangles intersect, in this case the parameterization is said to “fold over” or contain “triangle flips” (Figure 2 (b)).




Figure 2 Non-bijective parameterizations: (a) planar embedding with a global overlap; (b) planar embedding with a local overlap. The normal of the highlighted flipped triangle is inverted with respect to the other triangle normals.

The geometric shape of the domain triangles will typically be slightly different than the shape of the original triangles, resulting in angle and area distortion. Applications typically try to minimize the distortion for the whole mesh in a least squares sense. Very few meshes admit *isometric*, namely zero distortion, parameterizations. For example, only developable surfaces (such as cylindrical or conical sheets) admit planar isometric parameterizations. Maps that minimize the angular distortion, or *shear*, are called *conformal*, and maps that minimize area distortion are called *authalic*. Often conformal maps are also called *harmonic*, though, as shown by Floater and Hormann [2005], the two are not equivalent. These terms are borrowed from differential geometry of smooth surfaces, where Riemann's theorem guarantees that conformal maps (with zero angular distortion) always exist, mapping any infinitesimally small circle on the surface to a circle on the domain, and thus preserving angles, but allowing the scale factor of the transformation to vary across the map. The Riemann theorem does not hold for meshes. For instance, when considering planar parameterization, the sum of angles around an interior mesh vertex in 3D can vary, while the sum of angles around a vertex in the plane is always 2π . Some conformal parameterization methods applied to a series of progressively denser meshes of the same object (with smaller and smaller triangles, maybe obtained through subdivision) will converge in the limit to a smooth conformal map. The conformality of a mesh can be measured in multiple different ways [Hormann and Greiner 2000; Sheffer and de Sturler 2000; Lévy et al. 2002; Desbrun et al. 2002], resulting in different functionals to be optimized. For instance, Hormann and Greiner [2000] consider the minimal and maximal eigenvalues γ and Γ of the first fundamental form of the mapping. Alternatively, Sheffer and de Sturler [2000] directly measure the difference between the corresponding angles in the mesh and domain triangles.

As pointed out by Floater and Hormann [2005], though authalic parameterizations are achievable, they are not very useful by themselves, as they allow extreme angular and linear distortion. Thus, methods that consider area preservation [Desbrun et al. 2002; Degener et al. 2003] typically balance it with angle preservation.

Other metrics of parameterization distortion measure the preservation of distances across the mesh [Sander et al. 2001; Zigelman et al. 2002; Zhou et al. 2004]. The *stretch* metrics proposed by Sander et al. [2001] are now commonly used in the graphics community as standard measures of distance preservation. Sander et al. [2001] observe that the linear map over each triangle can be decomposed into a translation, a rotation and a non-uniform scale along two orthogonal axes. The two scale factors $0 \leq \gamma \leq \Gamma$ are the singular values of the transformation matrix, or the square roots of the eigenvalues of the integrated metric tensor (the transpose of the matrix times the matrix itself). Intuitively, the linear transformation map will stretch a unit circle to an ellipse

with axes γ and I . The L_∞ stretch for a triangle is defined by Sander et al. [2001] as $\max(\gamma, I) = I$, while the L_2 stretch is defined as $\sqrt{(\gamma^2 + I^2)/2}$. The name for the stretch metric comes from applications that map signals with regular sampling in the domain to 3D surfaces; these applications want to minimize the stretch of the signal over the surface, or the space between the locations of mapped samples. In other words, stretch penalizes undersampling the mesh, but not oversampling it.

When parameterizations are used to resample or compress 3D objects, the quality of the reconstruction can be measured by the symmetric Hausdorff metric between the original and reconstructed meshes. This metric measures the maximum distance between any point on either mesh and its projection on the other mesh. In practice most people use an RMS-average of these point distances. For compression schemes that allow trade-offs between the bit rate (file size) and approximation accuracy, the performance is measured using rate-distortion curves, plotting one against the other.

Keeping these definitions in mind we now proceed to survey the parameterization techniques available.

3 Parameterization of topological discs

The early papers to address parameterization for computer graphics applications were interested in planar parameterization of meshes with disk-like topology. The first application for such parameterizations was texture mapping. More recent applications include mapping of other surface properties such as normals or BRDFs, and mesh processing operations such as remeshing and compression. Two recent surveys [Floater and Hormann 2002; Floater and Hormann 2005] list more than twenty different planar parameterization techniques. Both surveys focus on the mathematical aspects of these techniques. To avoid unnecessary overlaps, we will address the more practical considerations, such as the suitability of the techniques for computer graphics applications in terms of distortion (type and amount), robustness, and efficiency. In our discussion we classify the methods based on the type of parametric distortion minimized.

Planar parameterization of 3D surfaces inevitably creates distortion in all but special cases. A well-known result from differential geometry is that for a general surface patch there is no distance-preserving (isometric) parameterization in the plane [do Carmo 1976]. Distance-preserving parameterizations exist only for developable surfaces: that is, surfaces with zero Gaussian curvature. Cutting the surface into charts or introducing seams (as discussed in Section 4), can reduce this distortion.

We classify planar parameterizations into roughly four groups: methods that ignore distortion altogether (Section 3.1), methods that minimize angular distortion (Section 3.2), methods that minimize stretch (Section 3.3), and methods that minimize area distortion (Section 3.4). There are also several techniques providing tools for achieving a tradeoff between different types of distortion (Section 3.5).

Ideally, most parameterization applications work best on zero distortion parameterizations, though most are tolerant to some amount of distortion, some being more tolerant to shear and others to stretch. In general, applications that depend on regular grids for sampling, such as different types of detail mapping and synthesis, as well as compression and regular resampling schemes (e.g. geometry images [Gu et al. 2002]), tend to perform better on stretch minimizing parameterizations, since stretch is directly related to under-sampling. In contrast, applications based on irregular sampling, such as remeshing [Desbrun et al. 2002], are very sensitive to

shearing, but can handle quite significant stretch. When acceptable levels of shear or stretch are not attainable because a surface is too complex, the surface needs to be cut prior to parameterization (Section 4) in order to achieve acceptable distortion.

In addition to distortion, several other factors should be considered when choosing a parameterization method for an application at hand:

- **Free versus fixed boundary.** Many methods assume the boundary of the planar domain is pre-defined and convex. Fixed boundary methods typically use very simple formulations and are very fast. Such methods are well suited for some applications, for instance those that utilize a base mesh parameterization, see Section 5.1. Free-boundary techniques, which determine the boundary as part of the solution, are often slower, but typically introduce significantly less distortion.
- **Robustness.** Most applications of parameterization require it to be bijective. For some applications local bijectivity (no triangle flips) is sufficient while others require global bijectivity conditions (the boundary does not self-intersect). Only a subset of the parameterization methods can guarantee local or global bijectivity. Some of the others can guarantee bijectivity if the input meshes satisfy specific conditions.
- **Numerical Complexity.** The existing methods can be roughly classified according to the optimization mechanism they use into linear and non-linear methods. Linear methods are typically significantly faster and simpler to implement. However, as expected the simplicity usually comes at the cost of increased distortion.

Method	Distortion minimized	Boundary	Bijectivity	Complexity
Uniform [Tutte 1963]	None	Fixed, convex	Yes	Linear
Harmonic [Eck 1995]	Angles	Fixed, convex	No	Linear
Shape preserving [Floater 1997]	Angles	Fixed, convex	Yes	Linear
Mean-value [Floater 2003]	Angles	Fixed, convex	Yes	Linear
LSCM/DCP [Lévy et al. 2002, Desbrun et al. 2002]	Angles (&Area)	Free	No	Linear
ABF/ABF++ [Sheffer and de Sturler 2000, Sheffer et al. 2005]	Angles	Free	Local (no flips)	Non-linear
MIPS [Hormann and Greiner, 2000]	Angles	Free	Global	Non-linear
Circle Patterns [Kharevych et al. 2006]	Angles	Free	Local (no flips)	Non-linear (unique minimum)
Stretch minimizing [Sander et al. 2001]	Distances	Free	Global	Non-linear
MDS [Zigelman et al. 2002]	Distances	Free	No	Non-linear
[Degener et al. 2003]	Areas	Free	Yes	Non-linear

Table 1: Planar method summary.

Table 1 provides a summary of a number of popular recent methods with respect to these four criteria. Figure 6 to Figure 9 and Table 2 provide distortion and runtime comparison between some of the more popular recent methods.

3.1 Planar mesh embedding

One of the oldest methods referenced in the context of mesh parameterization is the graph embedding method of Tutte [1963]. Tutte's formulation of graph embedding directly applies to triangular meshes. It also provides a general framework employed by many more-recent methods. The graph embedding formulation uses a two stage procedure. First, the boundary vertices of the mesh are mapped to the boundary of a convex region in 2D. Then the positions of the rest of the vertices are obtained by solving a linear system of the form

$$\begin{aligned}
 Lu = 0, Lv = 0 \\
 L_{i,j} = \begin{cases} -\sum_{k \neq i} L_{i,k} & i = j \\ w_{ij} & (i, j) \in E \\ 0 & \text{otherwise} \end{cases} \quad (1)
 \end{aligned}$$

The system is solved independently for the u and v coordinates. The parameter values of the fixed boundary vertices are used to specify the boundary conditions for the system. The weights w_{ij} are defined for each edge of the mesh. It was proven [Tutte 1963, Floater 1997] that if the weights are positive and the matrix is symmetric then the obtained parameterization is guaranteed to be bijective. Tutte [1963] used uniform unit weights, setting $w_{ij} = 1$ iff (i, j) is an edge in the mesh. Note that this setting defines L as the classical graph Laplacian matrix [Chung 1997]. While the resulting parameterization is provably bijective, it does not preserve any shape properties of the mesh (Figure 6 and Figure 8, top row).

Another well known, but less frequently used, mesh embedding technique is based on the circle packing theorem [Rodin and Sullivan 1987]. Given a triangular mesh, a circle packing is a collection of circles, one for each mesh vertex, such that two circles are tangential if there is a mesh edge between their associated vertices. By connecting the centers of these circles we obtain a planar embedding. The circle packing theorem states that for any given triangular mesh with disk topology and for any selection of radii r associated with the boundary vertices of the mesh, there exists a unique (up to symmetries) circle packing of the mesh with boundary circles having these radii. Collins and Stephenson [2003] provide a constructive technique to use the theorem to obtain mesh embeddings based only on combinatorial mesh information. Thus the method is guaranteed to generate bijective embeddings but does not preserve shape properties.

3.2 Angle preserving parameterization

Many of the disk-parameterization methods are focused on minimizing the angular distortion, or shear, of the parameterization. Several applications require angle-preserving parameterization; for instance, remeshing replaces one triangulation with another of better quality, typically by mapping a regular triangulation of the domain onto the mesh. The quality is often defined in

terms of angles: numerical simulations are adversely affected by small angles, while good compression depends on regular structures with triangles similar to their neighbors.

Angle-preserving parameterizations are efficient to obtain and are also often suitable for other applications, such as texture mapping, as long as the stretch of the parameterization is relatively low. This is often the case when the parameterized surfaces are not too far from developable. Riemann's theorem [do Carmo 1976] states that for smooth surfaces a conformal planar parameterization exists for any planar domain. Thus, since meshes are often viewed as approximations of smooth surfaces, we can argue that it is possible to map them to the plane with very little angular distortion.

Several parameterization techniques [Floater 1997, Eck et al. 1995, Floater 2003] use an approach similar to Tutte's [1963]. They first map the boundary of the surface to the boundary of a convex domain in 2D and then obtain the parameterization for the rest of the vertices solving a system of equations (Equation 1), independently for the u and v coordinates. The main difference between these methods is in the values w_{ij} assigned as edge weights. The choice of weights influences both the distortion and the bijectivity of the parameterization.

The harmonic or cotangent weights [Eck et al. 1995, Pinkall and Polthier 1993] are perhaps the most widely known in the graphics community. They define the matrix elements as:

$$w_{ij} = (\cot \alpha_{ij} + \cot \beta_{ij}) / 2 \quad (2)$$

where α_{ij} and β_{ij} are the opposite angles in the two triangles shared by the edge (i,j) as shown in Figure 3. These weights are derived from a finite-element description of harmonic energy and therefore aim at reducing the angular distortion of the parameterization. The greatest drawback of harmonic parameterization is that if the mesh contains obtuse angles, the weights can be negative, and, as a result, the parameterization can be non-bijective. Hence this formulation is not suitable, as is, for applications where bijectivity is a major concern. In practice many methods suggest first improving the mesh quality by bisecting obtuse angles or flipping edges. It has been proven that if the mesh satisfies the Delaunay criterion, even if it contains obtuse triangles, the parameterization obtained using the cotangent weights (Equation 2) will always be bijective. Recently Kharevych et al [2006] suggested using "intrinsic" Delaunay triangulation of the surface as an input to the harmonic mapping to guarantee bijectivity. Figure 6 and Figure 8 show a few examples of parameterizations generated using harmonic parameterization on a circular domain.

The shape preserving parameterization [Floater 1997] is also based on discrete harmonic mapping. It uses weights which are positive and symmetric thus ensuring that the parameterization is always bijective. However, the weights vary non-smoothly, i.e. they have derivative discontinuities, as the vertex i moves inside its one-ring polygon. Guskov [2002] modifies the shape preserving parameterization [Floater 1997] by introducing an extra term that allows an anisotropic parameterization stretched in the direction of a given surface direction field.

The mean-value weights also proposed by Floater [Floater 2003] are much simpler and vary smoothly

$$w_{ij} = (\tan(\gamma_{ij} / 2) + \tan(\delta_{ij} / 2)) / \|c(i) - c(j)\|,$$




Figure 3. Angles used for harmonic and mean-value weights.

where $c(i)$ and $c(j)$ are the 3D positions of vertices i and j , and γ_{ij} and δ_{ij} are the angles in the two triangles shared by the edge (i,j) shown in Figure 3. The mean-value weights are always positive. The resulting matrix is not symmetric ($w_{ij} \neq w_{ji}$) thus the bijectivity proof of Tutte [1963] does not directly apply to this formulation. Nevertheless, Floater [2003] proves that mean-value parameterization is guaranteed to be bijective. As shown in Figure 6 and Figure 8 (third row) the mean-value weights lead to parameterizations with small angular distortion. However, it was demonstrated [Zayer et al. 2005b] that they can in some cases introduce larger angular distortion than the more commonly used harmonic weights [Eck et al. 1995].

All of these techniques are very efficient and simple to implement as they only require solving a single linear system. However, since the parameterization distortion depends on how close the actual boundary shape matches the 2D domain shape, the fixed-boundary techniques perform poorly when the 3D meshes have non-convex boundaries (Figure 8), or boundaries that differ significantly from the specified boundary of the planar domain. Thus, these techniques work best when the inputs have well-shaped nearly convex boundaries. For instance, in a base-mesh setting (Section 5.1) the input mesh is first segmented into nearly triangular patches. In this setting the fixed boundary methods introduce very little distortion when mapping the patches to the corresponding triangular base mesh faces.

To reduce the distortion, the boundary of the 2D domain can be computed as part of the solution. Lee et al. [2002] free the domain boundary by introducing one or more layers of triangles around the original boundary in 3D, creating a *virtual* boundary (Figure 4 (a)). The virtual boundary is fixed onto a given convex polygon and the mapping for the rest of the vertices is computed using Floater's method [1997]. Since the original boundary vertices are free to move, the method obtains a parameterization with less distortion than with the original convex combination approach (Figure 4). A similar idea was used by Kos and Varady [2003]. Zhang et al. [2005] use *scaffold triangles* in the virtual boundary that only contribute a flip penalty to a stretch-like optimization metric.

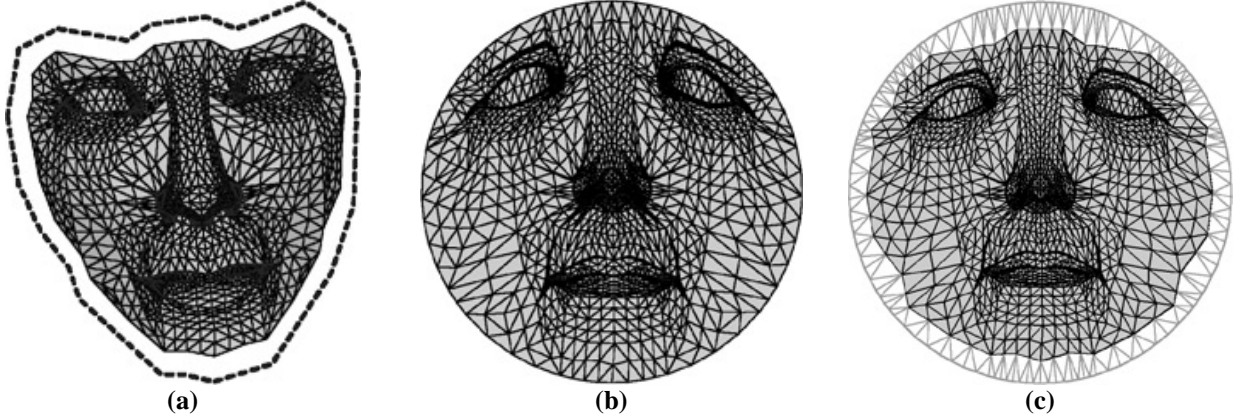


Figure 4. (a) Adding a virtual boundary to the original mesh. (b) Shape Preserving [Floater 1997] parameterization of the original mesh. (c) Parameterization of the original mesh and its virtual boundary. The virtual boundary vertices are fixed, allowing the real boundary vertices to move.

Recently, two explicit formulations of free-boundary, linear parameterization, LSCM and DCP, were independently proposed by Lévy et al. [2002] and Desbrun et al. [2002]. Using different formulations of harmonic energy, both sets of authors derived equivalent formulations for free-boundary parameterization that aim to minimize angular distortion. The formulation of Lévy et al. [2002] is based on the observation that given the angles α_1 , α_2 , α_3 of a planar triangle (Figure 5) the following holds:

$$(u_3, v_3) - (u_1, v_1) = \frac{\sin \alpha_2}{\sin \alpha_3} R^{\alpha_1} [(u_2, v_2) - (u_1, v_1)], \quad (3)$$

where u_i, v_i are the planar coordinates of the triangle vertices and R^α a rotation matrix with angle α . To minimize angular distortion the authors plug the original 3D angles into the formula and minimize the sum of square distances between the left and right sides of Equation (3). They fix two vertices to avoid the degenerate solution where all the vertices are at one point. Interestingly, the conditions specified by these two methods for interior mesh vertices are equivalent to the discrete harmonic energy formulation (Equation (2)). Thus the difference in the results can be viewed as a difference in the imposed boundary conditions.




Figure 5. LSCM [Lévy et al. 2002] notations.

Figure 6 and Figure 8 (bottom row) show some models parameterized with LSCM. As demonstrated LSCM/DCP introduce significantly less distortion than fixed boundary approaches, particularly near the domain boundaries. The methods do not guarantee local or global bijectivity, and can theoretically result in both flipped triangles and global overlaps. The

HLSCM method [Ray and Levy, 2003] describes a mechanism for speeding up the solution process for LSCM using a hierarchical solver.

On meshes with high curvature, non-linear conformal methods, discussed below, introduce significantly less stretch than DCP/LSCM (Figure 9). The MIPS method [Hormann et al., 1999; Hormann and Greiner 2000] optimizes a non-linear functional that measures mesh conformality (Figure 7 and Figure 9, top row). To obtain a solution it starts with a harmonic fixed boundary parameterization [Floater 1997] as an initial guess. It then proceeds to move vertices one at a time to reduce the distortion metric. To prevent flips, vertices are only permitted to move inside the kernel of neighbouring vertices. When moving boundary vertices, the method also checks for boundary overlaps and prevents those. Thus, it guarantees that the resulting parameterization remains globally bijective throughout the procedure. The authors advocate the use of conformal mapping for remeshing [Labsik et al., 2000] and surface fitting [Hormann 2000].

Instead of defining a planar parameterization in terms of vertex coordinates, the ABF method [Sheffer and de Sturler 2000; Sheffer and de Sturler 2001] defines it in terms of the angles of the planar triangles. Sheffer and de Sturler [2000] specify a set of constraints that angle values have to satisfy to define a planar triangular mesh. They search for angles that are as close as possible to the original 3D mesh angles and which satisfy those constraints. They then convert the solution angles into actual (u,v) vertex coordinates. The resulting parameterizations are guaranteed to have no flipped triangles (local bijectivity) but can contain global overlaps. The authors provided a mechanism for resolving such overlaps, but it has no guarantees of convergence. The original ABF method is relatively slow and suffers from stability problems in the angle-to- uv conversion stage for large meshes. ABF was augmented to yield ABF++ [Sheffer et al. 2005], a technique addressing both problems (Figure 7 and Figure 9, second row). ABF++ introduces a stable angle-to- uv conversion and drastically speeds up the solution by introducing both direct and hierarchical solution approaches.

Zayer et al. [2003] proposed to extend ABF by using additional constraints on the angles enforcing the parameter domain to have convex boundaries, thus guaranteeing global bijectivity. Zayer et al. [2005b] introduce an iterative free-boundary conformal method. They start with a fixed boundary parameterization and then relax it. In their comparisons the method seems to be faster but introduce more distortion than ABF and MIPS (Figure 10).

Hurdall et. al [1999] and Bowers and Hurdal [2003] compute angle-preserving parameterizations using a variation of the circle-packing formulation described later by Collins and Stephenson [2003]. They observe that, in the circle-packing setting, angle preservation corresponds to preservation of specified distances between circles. They construct patterns of non-intersecting circles, where each circle corresponds to a vertex of a triangulation. To preserve angles they optimize distances between circles corresponding to neighboring vertices. The authors use the conformal mapping technique for medical visualization, focusing on visualizing the structure of the human brain.

Kharevych et al. [2006] use a circle patterns approach where each circle corresponds to a mesh face. In contrast to classical circle packing, they use intersecting circles, with prescribed intersection angles. Given these angles, the circle radii follow as the unique minimizer of a


convex energy. The method first computes the intersection angles using non-linear constrained optimization and then finds the energy minimizer. Since the solution for the intersection angles is conformal only for a Delaunay triangulation, the authors employ a pre-processing stage that involves “intrinsic” Delaunay triangulation. At the final stage of the method, the computed angles are converted to actual uv -coordinates. The method supports equality and inequality constraints on the angles along the boundary of the planar parameter domain. Similar to ABF, the parameterization is locally bijective, but can contain global overlaps. As shown in Figure 7 and Figure 9 the amount of distortion introduced by the method is comparable with that of other non-linear conformal techniques. Kharevych et al. propose an extension of the method to global parameterization of meshes by introducing *cone singularities*, as described in Section 6.

Karni et al. [2005] propose an interesting iterative method for resolving triangle flips in existing planar triangulations. The input to their method is a planar triangulation generated by one of the methods above. If the triangulation contains flipped triangles they employ an iterative vertex relocation procedure to eliminate the flips. Their method can also be used to support soft positional constraints in planar triangulations. The method reduces the number of flipped triangles does not guarantee that the final triangulation will be flip-free.


Table 2 lists the runtimes for the different conformal methods. We used the Graphite [Graphite, 2003] 3D modeling system to time the fixed boundary methods, LSCM and ABF++. For the other methods the timings were provided by the authors. As expected, linear techniques are about one order of magnitude faster than the non-linear ones. Nevertheless, even the non-linear methods are fairly fast taking less than two minutes to process average size models.

Method	Cat Head (257 Δ)	Nefertiti (8071 Δ)	Cow (5804 Δ)	Camel (78K Δ)
Uniform [Tutte 1963] [*]	0.02 s.	0.23 s.	0.14 s.	2.91 s.
Harmonic [Eck et al. 1995] [*]	0.02 s.	0.26 s.	0.17 s.	3.21 s.
Mean-value [Floater 2003] [*]	0.02 s.	0.25 s.	0.16 s.	3.19 s.
LSCM [Lévy et al. 2002] [*]	0.03 s.	0.38 s.	0.20 s.	5.28 s.
ABF++ [Sheffer et al. 2005] [*]	0.06 s.	1.87 s.	0.77 s.	36.31 s.
MIPS [Hormann and Greiner, 2000] [‡]	1 s.	8 s.	5 s.	83 s.
Circle Patterns [Kharevych et al. 2006] [*]	0.1 s.	3.6 s.	2.1 s.	76.7 s.
Stretch minimizing [Sander et al. 2001]	5.17 s. [†]	12 s.	8.5 s.	127 s.
[†] 1.2 GHz Pentium M [*] 3.0 GHz Pentium 4 [‡] 3.2 GHz Pentium 4				


Table 2. Timings of various parameterization techniques.




Parameterization with uniform weights [Tutte 1963] on a circular domain.



Parameterization with harmonic weights [Eck et al. 1995] on a circular domain.






Parameterization with mean value weights [Floater 2003] on a circular domain.





Parameterization with LSCM [Lévy et al. 2002].




Figure 6. Linear parameterization methods






Parameterization with MIPS [Hormann and Greiner 2000].



Parameterization with ABF++ [Sheffer et al. 2005].




Parameterization with Circle Patterns [Kharevych et al. 2006].




Stretch Minimizing Parameterization [Sander et al. 2001].


Figure 7. Non-linear parameterization methods




Parameterization with uniform weights [Tutte 1963] on a circular domain.



Parameterization with harmonic weights [Eck et al. 1995] on a circular domain.




Parameterization with mean value weights [Floater 2003] on a circular domain.




Parameterization with LSCM [Lévy et al. 2002].


Figure 8. Linear parameterization methods on closed objects with a seam.




Parameterization with MIPS [Hormann and Greiner 2000].



Parameterization with ABF++ [Sheffer et al. 2005].



Parameterization with Circle Patterns [Kharevych et al. 2006].



Stretch Minimizing Parameterization [Sander et al. 2001].
Figure 9. Non-linear parameterizations for closed objects with a seam.




Figure 10. Setting the Boundary Free [Zayer et al. 2005b]

3.3 Distance preserving parameterization

In contrast to conformal parameterization, which always exists for smooth surfaces, stretch or distance preserving parameterization exists only for developable surfaces. Hence existing methods aim at minimizing linear distortion rather than completely eliminating it. Early parameterization techniques [Bennis et al. 1991; Maillot et al. 1993; Lévy and Mallet, 1998] introduced distance minimizing formulations which were numerically quite complex and thus hard to minimize.

Sander et al. [2001] introduce two metrics of parameterization stretch (L_2 and L_∞), described in Section 2. These two stretch metrics are now commonly used for comparing linear distortion between parameterizations [Zhou et al. 2004; Sheffer et al. 2005; Kharevych et al. 2006]. They then proceed to minimize the L_2 metric using a solution mechanism similar to that of Hormann and Greiner [2000], starting with a shape preserving parameterization [Floater 1997] and moving one vertex at a time (Figure 7 and Figure 9, bottom row). To speed up the parameterization the authors use a hierarchical top-down parameterization approach. Similar to MIPS, the stretch minimization method starts with non-self-intersecting parameter domains, therefore it can effectively prevent both flips and global overlaps, thus guaranteeing globally bijective parameterization. Since the method's goal is to minimize an L_2 norm, the method sometimes compensates for shrinkage in one direction by stretching in another, introducing shearing (see camel's neck in Figure 9, bottom row). Example runtimes for the method are listed in Table 2. The runtimes are slower but within the same order of magnitude as those of the non-linear conformal methods. Sander et al. [2002] introduced an extension of the method to signal specialized parameterization, where the user can influence the distribution of the distortion across the surface. This method assumes that the reconstructed signal is piecewise constant. Tewari et al. [2004] assume that the reconstruction is piecewise linear, which provides greater sensitivity to signal detail.

Zigelman et al. [2002] introduce a mesh parameterization method that aims to preserve the distances between all pairs of vertices on the surface. They first compute those distance using geodesics, and then use multi-dimensional scaling (MDS) to embed the mesh in the plane. The method performs quite well for surfaces that are close to developable. However for complex surfaces it often tends to fold the surface in the plane generating a non-bijective parameterization, since the folded solution provides better distance preservation [Zhou et al. 2004].

Zhou et al. [2004] also use an MDS or *isomap* based approach, introducing several mechanisms to speedup the computation process. If the parameterized planar mesh contains folds, they segment the surface, and parameterize each chart independently (Section 4). Figure 11 (b) shows the charts and the parameterization produced by the method.

In some settings it is possible to avoid solving a linear system and simply compute the parameterization one triangle at a time [McCartney et al. 1999, Sorkine et al. 2002]. However, for most meshes this technique must be combined with simultaneous generation of cuts in the mesh surface. We therefore choose to discuss it together with other cutting techniques in Section 4.

3.4 Area preserving parameterization

As pointed out by Floater and Hormann [2005], from a theoretical point of view, there are two independent qualities that can be minimized in a mapping: angular distortion and area distortion. For continuous surfaces, whereas conformal mappings are nearly unique, area-preserving mappings are not. Thus methods that search for *authalic*, namely area preserving, mesh parameterizations typically introduce additional optimization terms or constraints on the parameterization. Desbrun et al. [2002] use a similar derivation of their conformal DCP embedding to obtain a linear formulation for local area preservation. Degener et al. [2003] develop a non-linear formulation of per-triangle area preservation. Both papers then proceed to combine the area optimization metric with a conformality metric to allow for parameterizations that mediate between the two as explained in the next section.

3.5 Trade-off between metrics

Many applications of parameterizations fail when either the shear or the stretch are extreme. However, they can tolerate a small amount of either type of distortion. Thus it is sometimes advisable to provide a parameterization that provides an adequate degree of tradeoff between angle preservation and area, or length preservation.

Desbrun et al. [2002] provide a linear formulation which supports a tradeoff between angular and area distortion. It is not clear if the combined formulation can be used in a free-boundary setting. Degener et al. [2003] extended the MIPS method [Hormann and Greiner, 2000] to find parameterizations that mediate between angle and area deformation by modifying the minimized functional. They define the minimized energy to be a sum of per-triangle term. Each term is a product of the MIPS harmonic energy per triangle and an energy functional that measures area distortion. The powers of the components in the product provide the mediation between the two measures. The method maintains the bijectivity guarantees of the original MIPS algorithm. Clarenz et al. [2004] provide a different set of non-linear energy metrics allowing for trade-off between angle, area and distance distortion measures.

In recent years a number of researchers suggested using a two stage approach. First a mesh is parameterized using one of the existing techniques listed above, and then a second, post-processing stage is applied to improve the parameterization. The improvement typically focuses on reducing some measure of distortion, or providing a tradeoff between several distortion

measures. The methods measure the distortion of parameterization produced in the first stage of the algorithm, and then adopt the parameters of the second stage to reduce this distortion. Sheffer and de Sturler [2002] first computed an angle preserving parameterization using ABF [Sheffer and de Sturler 2000]; they then used an overlay grid to compute the stretch of the resulting parameterization; finally, they smoothed the overlay grid to mitigate the stretch. Yoshizawa and Belyaev [2004] and Zayer et al. [2005a] also start with a parameterization that minimizes angular distortion. They first use linear, fixed-boundary methods to compute one. They then compute the stretch of the parameterization; scale the weights used in the linear method by a mitigating factor and solve the linear system again. To scale the weights, Yoshizawa and Belyaev [2004] used the stretch metric of Sander et al. [2001], while Zayer et al. [2005a] use a simpler construction based on inverse edge lengths.

4 Cutting / chart generation

Planar parameterization is only applicable to surfaces with disk topology. Hence, closed surfaces and surfaces with genus greater than zero have to be cut prior to planar parameterization. As previously noted, greater surface complexity usually increases parameterization distortion, independent of the parameterization technique used. To allow parameterizations with low distortion, the surfaces must be cut to reduce the complexity. Since cuts introduce discontinuities into the parameterization, a delicate balance between the conflicting goals of small distortion and short cuts has to be achieved. It is possible to use constrained parameterization techniques to reduce cross-cut discontinuities [Kraevoy et al. 2003; Zhou et al. 2005] as discussed in Section 6. Cutting and chart generation are most commonly used when computing parameterizations for mapping of textures and other signals onto the surface. They are also used for applications such as compression [Gu et al. 2002] and remeshing.

The techniques for cutting surfaces can be roughly divided into two categories: segmentation techniques which partition the surface into multiple charts (Section 4.1), and seam generation techniques which introduce cuts into the surface but keep it as a single chart (Section 4.2). Multiple charts created by segmentation typically have longer boundaries than those created by seam cutting. However, they can often be more efficiently packed (Section 4.3) into a compact planar domain.

4.1 Mesh segmentation

Depending on the application, mesh segmentation techniques use different criteria for creating charts [Shamir 2004]. Segmentation techniques, used for parameterization, break the surface into several charts such that the parametric distortion when parameterizing each chart is sufficiently low, while the number of charts remains small and their boundaries are kept as short as possible. Since planes are developable by definition, one possible approach is to segment the surface into nearly planar charts [Maillot et al. 1993, Garland et al. 2001, Sander et al. 2001, Sander et al. 2003, Cohen-Steiner et al. 2004].

Earlier approaches are based on incremental clustering of mesh faces into charts. They typically start by selecting several initial *seed* mesh faces, and then grow the charts around those seeds, adding one mesh face at a time, until all the faces are classified as belonging to one chart or the other.

Recently, Sander et al. [2003] introduced a segmentation method inspired by Lloyd quantization (Figure 11 (d)). Their algorithm iterates between chart growing and reseeding stages. After each chart growing iteration the method selects the best seeds with respect to each chart and repeats the process. The authors demonstrate that by iterating, they obtain better results than single pass methods such as that of Sander et al. [2001].

Planes are a special type of developable surfaces. Thus for parameterization purposes planar segmentation is over-restrictive and usually generates more charts than necessary. Several recent approaches focused on developable segmentation instead [Lévy et al. 2002; Zhu et al. 2004; Julius et al. 2005]. Lévy et al. [2002] proposed to detect high mean-curvature regions on the mesh and then generate charts starting from seeds which are farthest from those regions (Figure 11 (b)). This approach tends to capture many developable regions, but can also introduce charts which are far from developable.

Zhou et al. [2004] propose a segmentation method based on spectral analysis of the surface (Figure 11 (a)). They compute a matrix of geodesic mesh distances and then perform face clustering by growing charts around the n farthest points in a space defined by the dominant eigenvectors of the matrix.

Zhang et al. [2005] use a Reeb graph connecting the critical points of the average geodesic distance to cut the mesh, reducing its genus to zero. The surface is further segmented into features by finding field iso-contours where the feature area increases significantly; then the features are classified as linear ellipsoids, flat ellipsoids, or spheres, and cut using a lengthwise, circular, or respectively a "baseball seam" cut. The resulting pieces are close to being developable and can be parameterized with little distortion (Figure 11 (e)).

Similar to Sander et al. [2003], Julius et al. [2005] use Lloyd iterations of growing and reseeding. But instead of looking for planar regions they search for a larger subset of developable surfaces, the so called "developable surfaces of constant slope", which are characterized by having a constant angle between the normal to the surface and some axis vector. They provide a metric to measure if a chart closely approximates such a surface, and use this metric in the chart growing and reseeding stages (Figure 11 (c)).

4.1.1 Chart packing

Chartification techniques raise an additional post-processing challenge. Following the parameterization of each individual chart, those charts need to be placed, or *packed*, in a common parameter domain. For efficient storage of the parameterized meshes, the packing has to be as compact as possible. The optimal packing problem is NP-hard [Milenkovic 1998], thus only heuristic or approximate packing algorithms exist. The Tetris algorithm [Lévy et al. 2002] introduces charts one by one, searching for the best fit along the active-front of the charts packed so far (Figure 11 (b)). Sander et al. [2003] extend this algorithm, testing for more options (Figure 11 (d)). Zhou et al. [2004] use a similar approach but consider also non-square parameter domains (Figure 11 (a)).




Figure 11. Segmentation and texture atlases for (a) Isocharts [Zhou et al. 2004] – 13 charts; (b) [Lévy et al. 2002] – 44 charts; (c) D-charts [Julius et al. 2005] – 12 charts; (d) multi-chart geometry images [Sander et al. 2003] – 15 charts; and (e) [Zhang et al 2005] – 24 charts.

If discontinuities along chart boundaries are not a concern, one way to obtain a parameterization with very low distortion is to view each triangle as a single chart. In this case, the challenge of chartification shifts to finding compact packing [Carr and Hart 2002].

When the shape of the charts is fixed, more efficient packing is possible. Tarini et al. [2004] parameterize a 3D mesh over a set of square charts as described in Section 5.1. They then store texture as a collection of small square images, and are able support triangles spanning multiple charts and perform correct filtering across tile boundaries by storing a secondary index structure. Mesh vertices have 3D, rather than 2D, texture coordinates, mapping them to the space in the vicinity of a base domain formed by several cubes glued together. These 3D coordinates are interpolated for each pixel drawn for the mesh, and then a fragment program on the graphics card projects them onto a nearby cube face (identified using a small 3D texture) and reads the associated texture for the face.

4.2 Seam Cutting

It is possible to reduce the parameterization distortion without cutting the surface into separate patches by introducing multiple partial cuts or seams inside a single patch. This typically leads to shorter cuts than those created by segmentation. Piponi and Borshukov [2000] generated such cuts manually using a network of edges.

McCartney et al. [1999] and Sorkine et al. [2002] perform parameterization and cutting simultaneously. They unfold the mesh vertices onto the plane one after the other, optimizing the local mapping. Whenever the distortion of the mapping reaches a threshold, Sorkine et al. [2002] cut the mesh to reduce it. As a result they have a hard bound on the distortion, but can end up with long and complicated boundaries (Figure 12 (a)). To measure distortion, Sorkine et al. [2002] use the singular values γ and Γ of the per-triangle mapping (Section 2).

Gu et al. [2002] use parameterization results to facilitate the cutting process. The authors first parameterize the surface using shape preserving parameterization [Floater 1997]. They then find the point of maximal parametric distortion on the mapping, and generate the shortest cut from the surface boundary to that point. They repeat the process until the distortion falls below a certain threshold (Figure 12 (b)).




Figure 12. Generating seams: (a) Sorkine et al [2002]; (b) Geometry Images [Gu et al., 2002]; (c) Seamster [Sheffer and Hart, 2002]. The color in (c) shows visibility- green is more visible, red is less. The Seamster cuts go across the palm which is considered more occluded than the back of the hand.

The *Seamster* algorithm [Sheffer 2002; Sheffer and Hart 2002] considers the differential geometry properties of the surface, independent of a particular parameterization technique. It first finds regions of high Gaussian curvature on the surface and then uses a minimal spanning

tree of the mesh edges to connect those. Finally it cuts the mesh along the tree edges. While Sheffer [2002] uses only the edge length criterion when generating the cuts, Sheffer and Hart [2002] use a visibility metric as an edge weight when computing the minimal spanning tree. This way they are able to trace the cuts through the less visible parts of the surface hiding the potential cross-cut discontinuities in texture or other maps on the surface (Figure 12 (c)). This algorithm was also used to cut the cow and camel in Figure 8 and Figure 9.

While methods that segment meshes into multiple charts implicitly generate charts with disk topology, seam cutting methods, such as [Sheffer and Hart 2002, Gu et al. 2002] require an explicit preprocessing stage to convert surfaces with high genus into topological disks. The generation of minimal length cuts that convert a high genus surface into a topological disk is NP-hard [Erickson and Har-Peled 2004]. Lazarus et al. [2001] introduce a method for extraction of canonical schema, converting a high genus surface into a topological disk which in practice produces relatively short cuts. Erickson and Har-Peled [2004] propose a handle cutting method which has some elegant theoretical guarantees but is complex to implement. A more practical approach is taken by Gu et al. [2002], who trace a spanning graph of all the faces in the mesh and then prune this graph, obtaining a genus reducing cut. Ni et al. [2004] smooth scalar functions over a mesh obtaining "fair" Morse functions that have few critical points and whose Morse complexes generate a small number of cuts for genus reduction.

5 Alternate base domains

Some applications are quite sensitive to discontinuities in the parameterization, or cannot tolerate them at all. In such cases, when the object to be parameterized is not a topological disc, it is worthwhile to use a different base domain for the parameterization. Examples of such domains that have been investigated include simplicial complexes (Section 5.1), spheres (Section 5.2), and periodic planar regions with transition curves (Section 5.3).

In addition, numerous applications of parameterization require cross-parameterization or inter-surface mapping between multiple models [Schreiner et al. 2004; Kraevoy et al. 2004]. Pair-wise mapping between models can be used for the transfer of different properties between the models, including straightforward ones, such as texture, and less obvious ones such as deformation and animation [Sumner and Popovic, 2004]. It can also be used for blending and morphing [Schreiner et al. 2004; Kraevoy et al. 2004; Allen et al. 2003], as well as mesh completion and repair [Kraevoy and Sheffer, 2005]. The most common approach for pair-wise mapping is to parameterize both models on a common base domain. Free-boundary planar parameterization is clearly unsuitable for this purpose. Instead alternate domains such as a simplicial complex or a sphere are commonly used.

5.1 Simplicial and quadrilateral complexes

Historically, the most popular non-planar base domain has been a simplicial complex [Lee et al. 1998; Lee et al. 1999; Guskov et al. 2000; Guskov et al., 2002; Lee et al. 2000; Praun et al. 2001; Khodakovsky et al. 2003; Purnomo et al. 2004; Schreiner et al. 2004; Kraevoy and Sheffer 2004]. A simplicial complex can be considered as just the connectivity part of a traditional triangle mesh: the sets of vertices, edges, and faces. Most applications typically use simplicial complexes representing 2-manifolds with a boundary (an edge can only be adjacent to 1 or 2 faces) with a small number of elements. One method for obtaining such complexes is to simplify

an original mesh. Once a suitable base mesh has been chosen, the original mesh is parameterized by assigning each of its vertices to a simplex of the base domain (vertex, edge, or face), along with barycentric coordinates inside it.

Early methods took a two-step approach to computing a parameterization; in the first step, elements of the fine mesh were assigned to faces of the base simplicial complex, while the second step would compute barycentric coordinates for these elements, usually using one of the fixed boundary parameterization methods discussed in Section 3. These steps could be repeated, but typically not mixed. More recent methods, such as [Khodakovsky et al. 2003], try to perform both steps at the same time.

5.1.1 Computing base complexes

To obtain the simplicial complex, Eck et al. [1995] grow Voronoi regions of faces from seed points and then use the dual triangulation (Figure 13 (a)). The seed points are initially linked using shortest paths across mesh edges that provide the initial boundaries of the patches corresponding to base domain faces. To straighten each of these paths, the two adjacent patches are parameterized to a square. The path in question is then replaced with the diagonal of the square mapped onto the mesh surface.




Figure 13. Base mesh construction.

Lee et al. [1998] simplify the original mesh, keeping track of correspondences between the original vertices and the faces of the simplified mesh (Figure 13 (b)). Others, like Guskov et al. [2000] (Figure 13 (c)) and Khodakovsky et al. [2003] use clustering techniques to generate the patch connectivity and derive the base-mesh from it.

The construction becomes more challenging when multiple models need to be parameterized on the same complex [Praun et al. 2001; Schreiner et al. 2004; Kraevoy and Sheffer 2004]. Praun et al. [2001] partition a mesh into triangular patches, which correspond to the faces of a user given simplicial complex, by drawing a network of paths between user-supplied feature vertices that correspond to the vertices of the base mesh (Figure 14 (a)).

Kraevoy et al. [2003] address a similar problem in the context of constrained planar parameterization (Section 6). Schreiner et al. [2004], and Kraevoy and Sheffer [2004] extend the methods of Praun et al. [2001] and Kraevoy et al. [2003] to construct the simplicial complex automatically, in parallel to the patch formation. The input to both methods includes a set of correspondences between feature vertices on the two input models. The methods use those as the vertices of the base complex. They simultaneously trace paths on the input meshes between corresponding pairs of vertices, splitting existing mesh edges if necessary (Figure 14 (b) and (c)).

Tarini et al. [2004] were the first, to our knowledge, to use a quadrilateral base domain (Figure 15 (a) and Figure 17 (a)). Such a domain is much more suitable for quadrilateral remeshing of the input surface and for spline fitting. Tarini et al. [2004] generate the base domain manually.

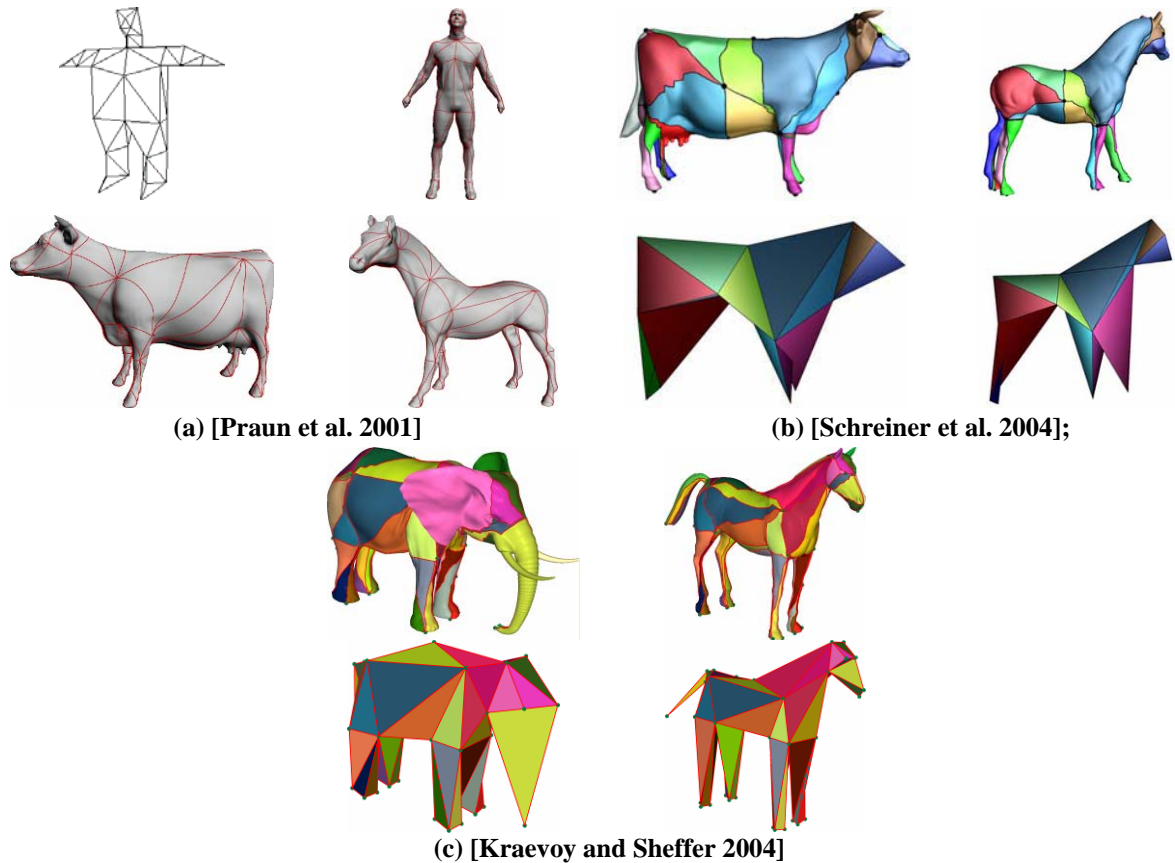


Figure 14. Consistent base mesh creation.

Boier-Martin et al. [2004] first construct a coarse mesh using normal-based clustering of faces followed by spatially-based clustering of the initially generated charts. This coarse mesh is then cleaned up and quadrangulated, yielding the base domain over which the input mesh is parameterized (Figure 15 (b)).

Recently, Tong et al. [2006] introduced a method for semi-automatic construction of curvature aligned quadrilateral base meshes. The method finds the umbilical points of the curvature field and then connects them along paths aligned with the principal curvature directions. They then proceed to find a global parameterization across the base domain as explained in Section 5.3 (Figure 15 (d)).

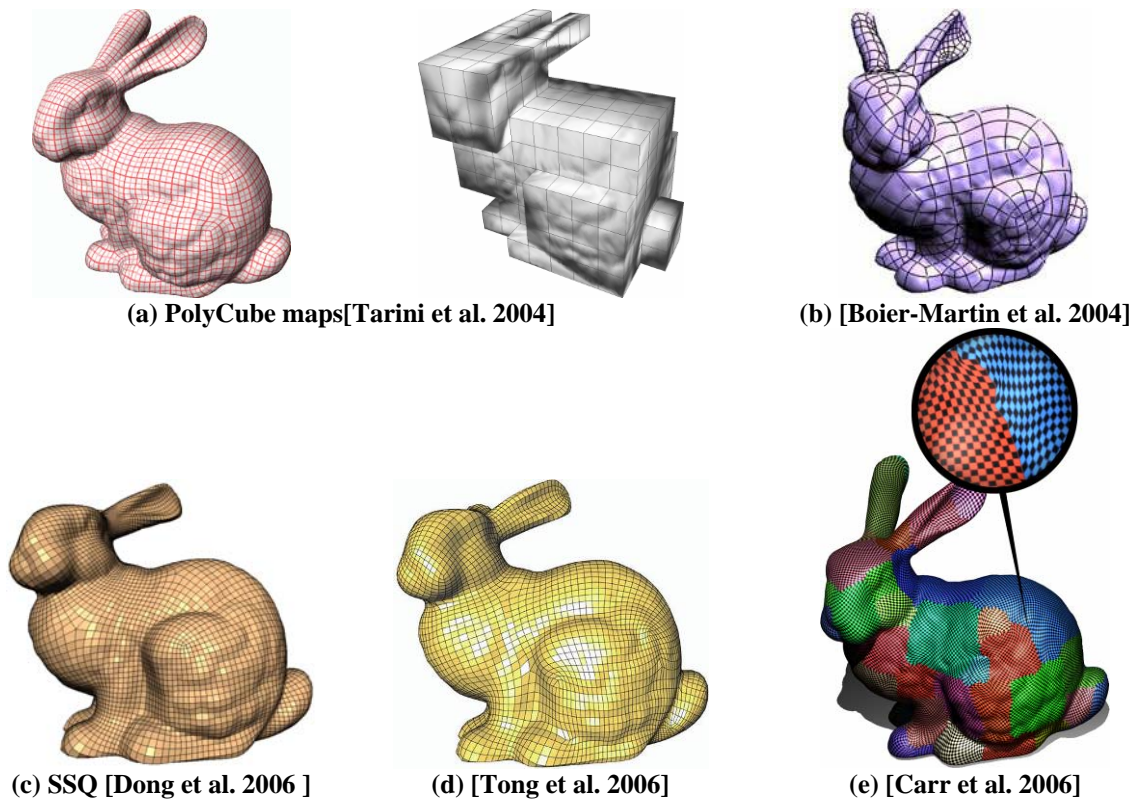



Figure 15. Quadrilateral base domain construction and mapping.

Dong et al. [2006] combine spectral analysis and Morse theory, computing the Morse complex of one of the eigenvectors of the Laplacian matrix (Equation 1) weighted by the cotangent weights (Section 3.2). After smoothing, the Morse complex partitions the mesh into quad patches (Figure 15 (c)). The constructed base domains and the subsequent parameterization are typically not aligned with principal curvatures, an undesirable property when the parameterization is used for surface fitting or remeshing.

Similar to Sorkine et al. [2002], Carr et al. [2006] grow clusters outward from initial seed faces. In order to constrain cluster growth to a rectangular shape, a parameterization distance metric is employed which roughly approximates geodesic distance on the surface. The rectangular charts are then parameterized using shape preserving parameterization [Floater 1997], with the constraint that the parameterization is continuous on chart boundaries (Figure 15 (e)).

5.1.2 Mapping to the base mesh

Once the discrete assignment to base domain faces has been done, the barycentric coordinates can be computed using fixed-boundary planar parameterization with the techniques from Section 3 [Eck et al. 1995; Guskov et al. 2000; Lee et al. 2000]. Earlier methods computed the barycentric coordinates once, based on the initial assignment of the vertices to the base triangles (Figure 16 (a) and (b)). More recent methods [Khodakovsky et al. 2003; Kraevoy and Sheffer 2004; Kraevoy and Sheffer 2005; Tarini et al. 2004] use an iterative process where vertices can be reassigned between base faces.




(a) MAPS [Lee et al. 1998] (b) [Guskov et al. 2000] (c) [Khodakovsky et al. 2003]
Figure 16. Mappings to base, visualizing derivative magnitudes of the tangent field
 (gray: small; red: large). The figures come from [Khodakovsky et al. 2003].

Khodakovsky et al. [2003] perform the vertex-to-patch assignment and coordinate relaxation in a single procedure, by letting vertices cross patch boundaries using transition functions (Figure 16 (c)). A transition function expresses the barycentric coordinates of a vertex with respect to a base domain face as barycentric coordinates for a neighboring base domain face. For this procedure only the images of the base domain vertices needs to be fixed, rather than the edges as well as in the previous methods. The authors relax the base domain vertices separately, prompting a new run of the main relaxation. In practice, this cycle is repeated only very few times. The implementation sometimes needs to discard some relaxation results when mesh vertices moved around base domain vertices end up with barycentric coordinates that are invalid for all the base domain faces around that vertex.


Tarini et al. [2004] and Kraevoy and Sheffer [2004; 2005] fix the boundary of a group of base mesh faces, update the barycentric coordinates in the interior, and then possibly re-assign some vertices to different faces inside the group. The methods differ in the grouping they use and the choice of parameterization technique used for the barycentric coordinates computation (Figure 17).

Schreiner et al. [2004] never compute an explicit map between the full-resolution objects and the base domain. Instead, they alternate the role of base domain between the two meshes, at various complexity levels in a multi-resolution representation. They progressively refine each mesh by adding new vertices and relaxing their location using a stretch-based metric measured on a temporary planar unfolding of their neighborhoods.

Most of the methods listed above assume that the input meshes are closed. Schreiner et al. [2004] handle meshes with boundaries by mapping the boundaries to edges of the base complex and the corresponding holes to a subset of its faces (Figure 18 (a)). Kraevoy and Sheffer [2005] support parameterization of meshes with both holes and multiple components (Figure 18 (b)). They allow free boundary parameterization of the holes and the outer boundaries of the components by introducing a *virtual* triangulation of the holes and gaps and updating this triangulation during the iterative parameterization process. This parameterization method is particularly suitable for mesh completion operations.




(a) [Tarini et al. 2004]





(b) [Kraevoy and Sheffer 2004]

Figure 17. Maps to base domains.



(a) [Schreiner et al. 2004]



(b) [Kraevoy and Sheffer 2005]

Figure 18. (left) Schreiner et al. [2004] parameterize connected models with boundaries. The center model is 50% morph of the side ones. (right) Kraevoy and Sheffer [2005] parameterize meshes with multiple components, such as the teddy mapped here to a sphere.

The modern parameterization methods that use simplicial or quadrilateral complexes, as described above, typically produce parameterizations which are visually smooth everywhere except in the vicinity of irregular vertices of the base meshes.


5.2 The unit sphere

The big advantage of the spherical domain over the planar one is that it allows for seamless, continuous parameterization of genus-0 models, and there are a large number of such models in use. Thus, the spherical domain has received much attention in the last few years, with several papers published about this topic [Haker et al. 2000, Praun and Hoppe 2003, Gotsman et al. 2003, Sheffer et al. 2003, Saba et al. 2005]. Some rigorous theory is being developed, getting close to the level of understanding we have of planar parameterizations.

One attractive approach for spherical parameterization is to extend the barycentric, convex boundary planar methods to the sphere. Several methods [Alexa 2000; Gu and Yau 2002; Gu et al. 2004; Kobbelt et al. 1999] used Gauss-Seidel iterations to obtain such parameterization. They start by computing an initial guess and then moving the vertices one at a time, first computing a 3D position for the vertex using a barycentric formulation [Floater 1997; Eck et al. 1995], and then projecting the vertex to the unit sphere (Figure 19 (a)). Isenburg et al. [2001] split the mesh in two, map the cut onto a great circle and embed each half-mesh onto a hemisphere using a modified Tutte procedure (Figure 19 (c)). Regrettably, as proven by Saba et al. [2005] “projected” Gauss-Seidel iterations decrease the residual for only a finite number of iterations. As the result approaches a bijective solution, the scheme ultimately becomes unstable, the residual increases, and the system collapses to a degenerate solution. Saba et al. [2005] note that this behavior is independent of step size.

Birkholz [2004] uses a hierarchical method which works with the spherical angles of the embedding. The hierarchical approach stabilizes the convergence, but is able to obtain only an approximate solution.

Haker et al. [2000] compute a planar parameterization of the mesh first and then use the stereographic projection to obtain the spherical mapping (Figure 19 (b)). This approach works quite well in practice even for dense meshes, however it doesn't offer any theoretical guarantees since the stereographic projection is bijective only for the continuous case, and can produce triangle flips in the discrete case. A simple proof by example of this statement can be obtained by imagining the great circle supporting the edge AB of a mapped spherical triangle ABC. The (continuous) stereographic projection maps this great circle to a circle in the original plane. The third vertex C can be perturbed in the plane to cross from the interior to the exterior of the circle, without changing the triangle orientation. The spherical triangle ABC will flip however as a result of this perturbation, as the image of C on the sphere will cross from one side to the other of the spherical edge AB.



Gotsman et al. [2003] showed how to correctly generalize the method of barycentric coordinates, with all its advantages, to the sphere. The generalization is based on results from spectral graph theory due to Colin de Verdière [1990] and extensions due to Lovasz and Schrijver [1999]. They provide a quadratic system of equations which is a spherical equivalent of the barycentric formulation. The authors do not provide an efficient way to solve the resulting system, and thus their method is limited to very small meshes. Saba et al. [2005] introduce a method for efficiently solving the system, by providing a good initial guess and using a robust solver (Figure 19 (d)). First, similar to Isenburg et al. [2001], they partition the mesh in two, and embed each half on a hemisphere using a planar parameterization followed by a stereographic projection. They then use a numerical solution mechanism which combines Gauss-Seidel iteration with non-linear minimization to obtain the final solution.

Sheffer et al. [2003] extend the ABF method [Sheffer and de Sturler, 2000] to spherical embeddings, formulating the problem in terms of spherical angles. Regrettably, it appears that the spherical formulation is numerically much less stable than its planar equivalent. Thus the method seems impractical for large meshes.




Figure 19. Spherical parameterizations.

An efficient and bijective alternative is suggested by multi-resolution techniques. These methods obtain an initial guess by simplifying the model until it becomes a tetrahedron (or at least, convex), trivially embed it on the sphere, and then progressively add back the vertices [Shapiro and Tal 1998, Praun and Hoppe 2003]. Shapiro and Tal [1998] compute the embedding using purely topological operations and do not attempt to minimize any type of distortion. Praun and Hoppe [2003] obtain a spherical parameterization by alternating refining steps that add vertices from a multi-resolution decomposition of the object with relaxation of single vertex locations inside their neighborhoods. The relaxation is aimed to minimize the stretch metric of the parameterization and is guaranteed to maintain a valid embedding (Figure 19 (e)).

Zayer et al. [2006] introduce a spherical parameterization method which takes user prescribed poles into account (Figure 19 (f)). They then cut the mesh along a date line connecting these poles. With the mesh then topologically equivalent to a disc, an initial parameterization is

defined over a rectangular domain by solving a Laplace equation in curvilinear coordinates. The parameterization distortion is reduced using a variant of quasi-harmonic maps [Zayer et al. 2005a]. Finally, they reduce the distortion along the date line by performing tangential Laplacian smoothing.

5.3 Discrete differential 1-forms

The classical approach for parameterizing models of arbitrary topology was to first construct a base-complex and then parameterize the mesh onto the complex (Section 5.1). Recently, several methods [Gu and Yau 2002; Ray et al. 2005; Kharevych et al, 2006; Tong et al. 2006] have proposed implicitly creating a parameterization by solving for differential one-forms. Instead of associating (u,v) coordinates with vertices as in typical planar parameterization, they associate planar vectors (du, dv) with the edges of the mesh, computing a gradient field, or *one-form*. These parameterizations can be converted to a special form of planar parameterization by fixing a vertex, then walking around to other vertices and adding edge vectors. The method guarantees that inside a topological disc region, the same coordinates are obtained for each vertex regardless of the path used to walk to it from the source vertex. Genus-1 can be mapped to an infinite plane using a multi-periodic function by tiling the plane using a *modular parallelogram* - a fundamental domain bounded by 4 curves, parallel two by two. For higher genus, applying this method directly will cover the plane multiple times. Such parameterizations are guaranteed to be continuous everywhere, except at a small number of singular points, and are therefore sometimes referred to as globally continuous parameterizations.

The first method for global periodic parameterization was proposed by Gu and Yau [2003] (Figure 20 (a)). After computing a holomorphic one-form, corresponding to a global conformal parameterization, the authors propose constructing a global conformal atlas for the surface, by separating the object handles and mapping each of them to a modular parallelogram. Each parallelogram can have interior transition curves that conformally connect the parameterization to corresponding curves inside other modular parallelograms. The images on the mesh of these transition curves separate the handles.

The theory behind this parameterization comes from an elegant connection to Riemann geometry [Morgan 1998]. In practice, computing the parameterization involves solving linear systems with three types of constraints: harmonicity, closedness, and duality. The harmonicity equations are similar to the ones used by methods in Section 3: the sum of vectors for edges incident to a vertex, weighed by the cotangent harmonic weights, is zero. The closedness equations state that the sum of the 3 vectors for the edges of each face is zero (a “gradient field” has no divergence). The duality conditions replace the boundary conditions in traditional parameterizations; they impose fixed values for the integral of the field on the closed loops forming the homology basis of the surface (the $2g$ curves that *cut open* the handles of the mesh, where g is the genus of the surface). Several systems are solved, in each system the integral across one of the curves is constrained to 1, and all the others to 0.

The same authors present a faster, multi-resolution method [Gu et al., 2004a], applications to mesh compression [Gu et al. 2003], and brain surface mapping [Gu et al. 2004b]. In another recent work, Gu et al. [2004a] find the unique locations of the singular points that satisfy conformality. They also solve for the optimal conformal transformation that minimizes global stretch. Since the number of singular points remains constant, the resulting parameterizations

often still exhibit significant stretch. The methods do not guarantee bijectivity but perform well in practice.

Gortler et al. [2005] provide further theoretical insight into the applications of one-forms to parameterizing meshes. They develop a discrete counterpart of the Hopf-Poincaré index theorem. This theorem generalizes the notion of singular points mentioned above and allows their splitting and merging. It states that for a surface of genus g , the indices of all singular points sum to $2g-2$, where the index of a singular point corresponds to its multiplicity. Their results include a simpler proof of Tutte’s theorem and a proof that the method applied to genus-1 objects indeed produces a locally bijective embedding (no triangle folds), for any set of asymmetric positive weights. The proof extends to non-triangular meshes as well.




Figure 20. (a) Global Conformal Surface Parameterization [Gu and Yau 2003]; (b) Periodic Global Parameterization [Ray et al. 2006].

Motivated by geometry processing applications, Ray et al. [2005; 2006] incorporate more geometric information into the problem setting (Figure 20 (b)). Given the curvature tensor of the surface, their method computes two piecewise linear periodic functions, aligned with the minimal and maximal curvatures. Together these two functions define a smooth parameterization almost everywhere, except in the vicinity of the singular points of the tensor field. The authors demonstrate how to extract a quadrilateral chart layout from this periodic parameterization. The method can construct quasi-isometric parameterizations at the expense of introducing additional singular points in non-developable regions where the curl of the input field is non-zero. They also propose a curvature-adapted parameterization method that minimizes the curl and removes those additional singular points by adaptively scaling the parameterization. The authors demonstrate that the method is well suited for quadrilateral remeshing and surface fitting.

Kharevych et al. [2006] extend their method of parameterization using circle patterns to global parameterization. They observe that in angle space formulation the only difference between parameterizing the mesh boundary and its interior mesh is in the constraint imposed on the sum of vertices in the interior, and that it is possible to define a global parameterization by specifying an unconnected subset of mesh vertices as boundary vertices. Thus they first compute a solution in angle space with a set of *cone singularity* vertices specified by the user as boundary. For planar parameterization, to perform the angle-to-uv conversion they later compute edge paths between these vertices. The obtained parameterization is globally continuous up to translation

and rotation, everywhere except at the cone singularities. The proposed approach can be directly applied to other angle-space methods such as ABF/ABF++.

Tong et al. [2006] compute a global parameterization by expressing the differential 1-form system back in terms of the associated potential fields (the 0-form, or uv variables per vertex) augmented with a set line singularities (Figure 17). This singularity graph is similar to a quadrilateral base domain used by the methods in Section 5.1, and can be either user-provided or generated semi-automatically from the mesh curvature.

5.4 Hyperbolic plane

Objects of genus n can be mapped to an infinite tiling of the hyperbolic disc with $4n$ -sided regular polygons. Straight lines from Euclidean geometry correspond to arcs that intersect the main hyperbolic disc at 90 degrees. The fundamental domain is a regular hyperbolic polygon centered at the origin, where each opposing pair of edges corresponds to one of the $2n$ loops that cut open the handles of the object. All the vertices of the polygon correspond to a single point on the object, where all the loops meet, which is also the only intersection point between any two loops. The sum of angles of the polygon is 2π (which is possible for a hyperbolic polygon with more than 3 sides). To produce the tiling, the polygon is repeatedly “flipped” over an edge, using linear fractional transforms: $(z) \rightarrow (az+b / cz + d)$, where z is a point inside the disc, represented as a complex number.

A parameterization based on hyperbolic geometry was proposed by Ferguson and Rockwood [1993], and Rockwood and Park [1999]. Grimm and Hughes [2003] use a hyperbolic parameterization as the starting point for creating a manifold: an atlas with several overlapping charts and transition functions between them. Each chart corresponds to one of the elements of the base domain. In the case of the original un-subdivided regular polygon: one face, one (shared) vertex, and $2n$ (pairs of) edges. Each chart overlaps charts of adjacent elements, resulting in a transition function, and does not overlap any charts of non-adjacent elements. Jin et al. [2006] provide a practical algorithm for computing hyperbolic parameterization for meshes of arbitrary genus (Figure 21).




Figure 21. Hyperbolic parameterization [Jin et al. 2006].

6 Constraints

Sometimes a parameterization needs to accommodate user input, in the form of correspondences between vertices of the mesh, or curves traced on the surface, and locations in the domain. The most important application of constrained parameterization is texture mapping 3D models from photographs [Lévy, 2001, Eckstein et al. 2001, Kraevoy et al. 2003]. Zhou et al. [2005] go further and allow the user to combine several images to produce a complete texture for a mesh. Constrained parameterization can also be used to hide cross-seam discontinuities [Kraevoy et al. 2003; Zhou et al. 2005].

Constraints are also used to establish maps with feature correspondence between different 3D objects [Schreiner et al., 2004; Kraevoy and Sheffer, 2004]. Such maps between meshes allow morphing, transferring of textures and other attributes between meshes, and principal component analysis of similar models in a database. Constrained parameterization methods can be grouped into two large categories, based on whether they treat this user input as “suggestions” or “help”, or rather, hard constraints to be satisfied exactly.

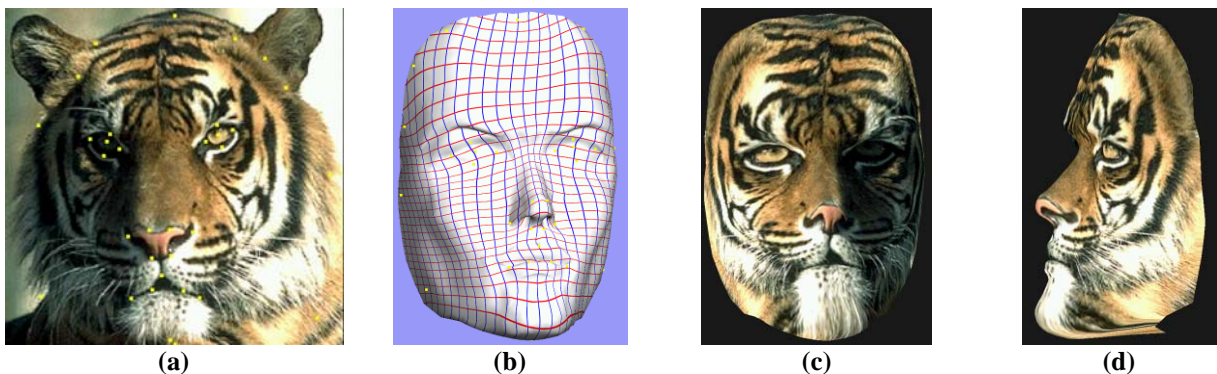


Figure 22. Soft constraints [Lévy 2001]. Mapping the tiger texture (a) onto the face in (b) yields (c) and (d).

Soft Constraints. Methods based on energy minimization can accommodate soft constraints by adding a quadratic term to the energy function, measuring the distance between the constraint features in the current configuration and their desired location [Lévy 2001]. Such constraints work reasonably well in practice, and can be solved efficiently, since they only add linear terms to the energy, but sometimes break theoretical guarantees about the original parameterization method, such as bijectivity. Allen et al. [2003] deform a template mesh to match a series of laser scans of human bodies by softly matching a set of about 70 marker constraints and a dense set of automatically computed point correspondences between the surfaces. Sumner and Popovic [2004] use similar ideas to transfer deformation (and whole animation sequences) from one mesh to another, though they do not extract an explicit one-to-one mapping between surfaces. Note that in these last two papers the energy being minimized takes into account more than just the difference between the actual and desired locations of a constraint, and actually softly matches a whole affine transformation matrix defining the surfaces around each constraint. These methods typically work well when the constraints do not require significant changes in the output compared to the unconstrained result. In addition, the methods no longer guarantee that the resulting parameterization will remain bijective.

Hard constraints. Applications such as hiding texture discontinuities along seams in the parameterization [Kraevoy et al. 2003; Zhou et al. 2005] require hard constraints to achieve perfect alignment of the texture along the seams. In addition hard constraints can be enforced while preserving parameterization bijectivity.

It is possible to enforce constraints by adding them into a regular parameterization formulation using Lagrange multipliers [Desbrun et al. 2002]. This approach allows constraints to be defined on points inside triangles and on arbitrary line segments (the vertices of the triangulation can be constrained more easily by taking the corresponding variables out of the system). However, it is easy to show that for a given mesh connectivity not every set of constraints can be satisfied. Thus methods like this that preserve the mesh connectivity will fail to generate bijective parameterizations for many inputs.

It is known that an embedding can always be found if the mesh is enriched by adding a number of additional *Steiner* vertices [Pach and Wenger 1998], Regrettably, finding the minimal number of vertices that have to be added is NP-Hard [Pach and Wenger 1998].

Methods that enforce hard constraints therefore modify the provided meshes by inserting a few new Steiner vertices [Eckstein et al. 2001, Kraevoy et al. 2003] or by complete remeshing [Praun et al. 2001].

Eckstein et al. [2001] introduce a method that enforces “hard” (exact) positional constraints by deforming an existing embedding while adding a number of Steiner vertices (Figure 23). Theoretically, this method can handle large sets of constraints but is extremely complicated.

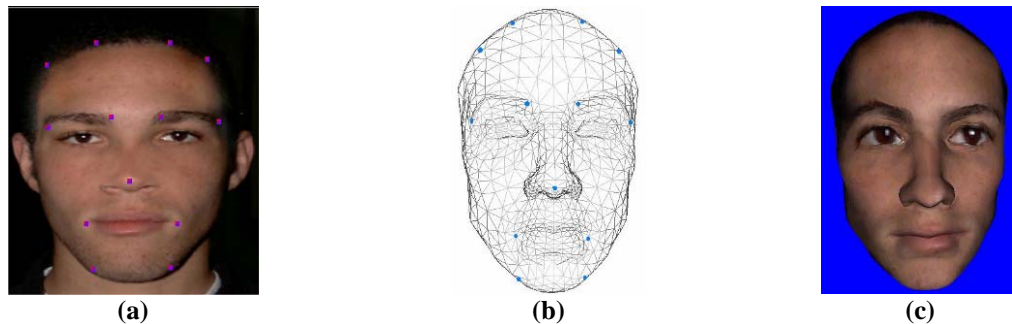


Figure 23. Eckstein et al. [2001] - Mapping a (a) texture of a face to a (b) mesh of a face. Result is in (c).

Instead, Praun et al. [2001] and Kraevoy et al. [2003] compute the parameterization by establishing coarse patch correspondences between the input and the parameter domain (Figure 24). The provided feature points on the input model and the parameter domain are connected using a network of curves that partition the surface into patches that are then parameterized while trying to maintain continuity and smoothness between them. The curve tracing process is guided by a set of topological rules that ensure that the resulting patches will be consistent between the object(s) being parameterized and the domain, and by heuristics that try to guide them on the correct side of constraints and prominent model features (such as limbs), avoiding awkward parameterization “swirls”. Praun et al. [2001] use a user provided triangulation of the features on the parameter domain, and enforce it on the model, while Kraevoy et al. [2003] compute the two triangulations simultaneously, solving a more challenging problem. Continuity and smoothness between patches can be obtained by parameterizing or relaxing the parameterization of several patches simultaneously.

Zhou et al. [2005] allow the user to combine several images to produce texture for a mesh, by assigning some surface patches to different images, as well as using in-painting techniques to

create texture for any unassigned transition patches between them (Figure 25). In addition to geometric smoothness of the map, Zhou et al. [2005] take into account the continuity of the texture signal being applied since it may come from different sources for two neighboring patches.




Figure 24. Mapping a photograph onto a model using Matchmaker [Kraevoy et al. 2003]. Pink dots in (a) and (b) represent constraints. (c) and (d) show various views of the texture mapped model.




Figure 25. Texturing Buddha from textures of multiple objects [Zhou et al. 2005]. From left to right: input images, generated texture atlas and mapping results.

7 Conclusions

In this paper we reviewed recent mesh parameterization techniques, as well as related mesh processing tools such as mesh segmentation and seam cutting. We focused on the more practical aspects of using parameterization for computer graphics applications, discussing method properties such as speed, robustness, and ease of implementation.

We would like to note, that judging by the amount of active research on parameterization, no survey on the topic can, realistically, be truly complete, and ours is no exception. We are sure to have missed at least a few of the related techniques introduced in the last couple of years.

Lastly we observe that in recent years many links have emerged between parameterization and other mesh processing techniques, such as mesh editing and surface approximation which are not addressed by this survey. We also did not address parameterization of other types of surface

representations, such as free-form surfaces and point-based surfaces, which provide challenges of their own, and for which there is a separate, vast body of work.

We believe that while there are numerous methods for surface parameterization available, many challenges remain. There is room for further research on topics such as global parameterization techniques, techniques that address complex constraints, and segmentation. There are also many open questions in terms of finding suitable numerical tools and better understanding the theoretical aspects of parameterization techniques as well as convergence properties for the numerical tools used.

Acknowledgements

The authors would like to thank the reviewers for their comments and Ciarán Llachlan Leavit for help with editing the paper. Sheffer and Rose gratefully acknowledge the research support of NSERC and IBM.

References

- [Alexa, 2000] ALEXA, M. 2000. Merging polyhedral shapes with scattered features. *The Visual Computer*, 16 (1), pp. 26-37.
- [Alexa, 2002] ALEXA, M. 2002. [Recent advances in mesh morphing](#). *Computer Graphics Forum* 21(2), pp. 173-196.
- [Allen et al., 2003] ALLEN, B., CURLESS, B., AND POPOVIĆ, Z. 2003. [The space of human body shapes: reconstruction and parameterization from range scans](#). *ACM SIGGRAPH 2003*, pp. 587-594.
- [Alliez et al., 2005] ALLIEZ, P., UCELLI, G., GOTSMAN, C., AND ATTENE, M. 2005. [Recent advances in remeshing of surfaces](#). Springer.
- [Alliez and Gotsman, 2005] ALLIEZ, P., AND GOTSMAN, C. 2005. [Recent advances in compression of 3D meshes](#). *Advances in Multiresolution for Geometric Modelling*, pp.3 -26.
- [Anguelov et al., 2005] ANGUELOV, D., SRINIVASAN, P., KOLLER, D., THRUN, S., RODGERS, J. DAVIS, J. 2005. [SCAPE: Shape Completion and Animation of People](#). *ACM SIGGRAPH 2005*.
- [Bennis et al., 1991] BENNIS, C., VÉZIEN, J.-M., AND IGLÉSIAS, G. 1991. Piecewise surface flattening for nondistorted texture mapping. *ACM SIGGRAPH '91*, pp. 237-246.
- [Blinn, 1978] BLINN, J. F. 1978. Simulation of Wrinkled Surfaces, *ACM SIGGRAPH '78*, pp. 286-292.
- [Biermann et al., 2002] BIERMANN, H., MARTIN, I., BERNARDINI, F., AND ZORIN, D., 2002. [Cut-and-paste editing of multiresolution surfaces](#). *ACM SIGGRAPH 2002*, pp. 312-321.
- [Birkholz, 2004] BIRKHOLOZ, H. 2004. Shape-preserving parametrization of genus 0 surfaces. *Proc. Winter Conference on Computer Graphics (WSCG)*.
- [Blanz and Vetter, 1999] BLANZ, V., AND VETTER, T., 1999. [A morphable model for the synthesis of 3D faces](#). *ACM SIGGRAPH 1999*.
- [Blanz et al., 2003] BLANZ, V., BASSO, C., POGGIO, T. AND VETTER, T., 2003 [Reanimating Faces in Images and Video](#), *Computer Graphics Forum* 22(3), *EUROGRAPHICS 2003*, pp. 641 – 650.
- [Blanz et al., 2004] BLANZ, V., SCHERBAUM, K., VETTER, T., AND SEIDEL, H.-P., 2004. [Exchanging faces in images](#), *EUROGRAPHICS 2004*.
- [Boier-Martin, 2004] BOIER-MARTIN, I., RUSHMEIER, H., JIN, J. 2004. [Parameterization of Triangle Meshes over Quadrilateral Domains](#), *Eurographics Symposium on Geometry Processing 2004*, pp. 197-208.
- [Bowers, and Hurdal, 2003] BOWERS, P. AND HURDAL, M. 2003. [Planar Conformal Mappings of Piecewise Flat Surfaces](#). *Visualization and Mathematics III*, Springer Verlag, pp. 3-34.
- [do Carmo, 1976] DO CARMO, M. 1976 *Differential geometry of curves and surfaces*, Prentice Hall.
- [Carr and Hart, 2002] CARR, N., AND HART, J. 2002. [Meshed Atlases for Real-Time Procedural Solid Texturing](#). *ACM Transactions on Graphics* 21(2), Apr. 2002
- [Carr and Hart, 2004] CARR, N., AND HART, J. 2004. [Painting Detail](#), *ACM SIGGRAPH 2004*.
- [Carr et al., 2006] CARR, N., HOBEROCK, J., CRANE, K., AND HART, J. 2006. [Rectangular Multi-Chart Geometry Images](#), *Proc. Symposium on Geometry Processing (SGP)*.
- [Chung, 1997] CHUNG, F. R. K. 1997. [Spectral Graph Theory](#). Providence, RI: Amer. Math. Soc..
- [Clarenz et al., 2004] CLARENZ, U., LITKE, N., AND RUMPF, M. 2004. Axioms and variational problems in surface parameterization. *Comput. Aided Geom. Des.* 21, 8,
- [Cohen-Steiner et al., 2004] COHEN-STEINER, D., ALLIEZ, P., AND DESBRUN, M. 2004. [Variational shape approximation](#). *ACM SIGGRAPH 2004*, pp. 905-914.

- [Colin de Verdière, 1990] COLIN DE VERDIÈRE, Y. 1990. On a new graph invariant and a criterion for planarity. *Graph structure Theory/Contemporary Mathematics*, pp. 137-147.
- [Collins et al., 2003] COLLINS, C., AND STEPHENSON, K. 2003. A Circle Packing Algorithm. *Computational Geometry: Theory and Applications* 25, pp. 233-256.
- [Degener et al., 2003] DEGENER, P., MESETH, J., AND KLEIN, R. 2003. [An adaptable surface parameterization method](#), *Proceedings, 12th International Meshing Roundtable*, pp. 201-213.
- [Desbrun et al., 2002] DESBRUN, M., MEYER, M., AND ALLIEZ, P. 2002. [Intrinsic parameterizations of surface meshes](#). *Computer Graphics Forum, 21 (2002)*, pp. 209-218.
- [Dong et al., 2006] DONG, S., BREMER, P.-T., GARLAND, M., PASCUCCI, V., AND HART, J. C. 2006. [Spectral Surface Quadrangulation](#). *ACM SIGGRAPH 2006*
- [Eck et al., 1995] ECK, M., DE ROSE, T., DUCHAMP, T., HOPPE, H., LOUNSBERY, M., AND STUETZLE, W. 1995. [Multiresolution analysis of arbitrary meshes](#). *ACM SIGGRAPH 95*, pp. 173-182.
- [Eckstein et al., 2001] ECKSTEIN, I., SURAZHISKY, V. AND GOTSMAN, C. 2001. [Texture mapping with hard constraints](#). *Eurographics 2001*, pp. 95-104.
- [Erickson, Jand Har-Peled, 2004] ERICKSON, J. AND HAR-PELED, S. 2004. [Optimally Cutting a Surface into a Disk](#). *Discrete Computational Geometry* 31(1), pp. 37-59.
- [Ferguson and Rockwood, 1993] FERGUSON, H., AND ROCKWOOD, A. 1993. [Multiperiodic functions for surface design](#). *Computer Aided Geometric Design, 10(3)*, pp. 315-328.
- [Floater, 1997] FLOATER, M., 1997. [Parameterization and smooth approximation of surface triangulations](#). *CAGD 1997, 14(3)*, pp. 231-250.
- [Floater and Hormann 2002] FLOATER, M. AND HORMANN K. 2002. Parameterization of Triangulations and Unorganized Points. *Tutorials on Multiresolution in Geometric Modelling*.
- [Floater, 2003] FLOATER, M., 2003. [Mean value coordinates](#). *CAGD, 20(1)*, 19-27.
- [Floater and Hormann, 2005] FLOATER, M. AND HORMANN, K. 2005. [Surface Parameterization: a Tutorial and Survey](#). *Advances in Multiresolution for Geometric Modelling*, pp. 157-186.
- [Garland et al., 2001] GARLAND, M., WILLMOTT, A., AND HECKBERT, P. 2001. [Hierarchical Face Clustering on Polygonal Surfaces](#). *Proceedings of ACM Symposium on Interactive 3D Graphics 2001*.
- [Gortler et al., 2005] GORTLER, S.J., GOTSMAN, C., AND THURSTON, D. 2005. [Discrete One-Forms on Meshes and Applications to 3D Mesh Parameterization](#). *Computer Aided Geometric Design*.
- [Gotsman et al., 2003] GOTSMAN, C., GU, X. AND SHEFFER, A. 2003. [Fundamentals of spherical parameterization for 3D meshes](#). *ACM SIGGRAPH 2003*, pp. 358-364.
- [Graphite, 2003] GRAPHITE, 2003, <http://www.loria.fr/~levy/Graphite/index.html>.
- [Grimm and Hughes, 2003] GRIMM, C. AND HUGHES, J. 2003. [Parameterizing N-holed Tori](#). *The Mathematics of Surfaces IX*.
- [Gu et al., 2002] GU, X., GORTLER, S., AND HOPPE, H. 2002. [Geometry images](#). *ACM SIGGRAPH 2002*, pp. 356-361.
- [Gu and Yau, 2002] GU, X. AND YAU, S.-T. 2002. [Computing conformal structures of surfaces](#). *Communications in Information and Systems, 2(2)*, pp. 121-146.
- [Gu, and Yau, 2002] GU, X., AND YAU, S.-T. 2003. [Global conformal surface parameterization](#). *Symposium on Geometry Processing*, 127-137.
- [Gu et al., 2003] GU, X., WANG, Y., AND YAU, S.-T. 2003. [Geometric Compression Using Riemann Surface Structure](#). *Communication in Information and Systems*.
- [Gu et al., 2004a] GU, X., WANG, Y., AND YAU, S.-T. 2004 [Multiresolution Computation of Conformal Structures of Surfaces](#), *Journal of Systemics, Cybernetics and Informatics* 1(6).
- [Gu et al., 2004b] GU, X., WANG, Y., CHAN, T. F., THOMPSON, P. M., AND YAU, S.-T. 2004. [Genus Zero Surface Conformal Mapping and Its Application to Brain Surface Mapping](#), *IEEE Transaction on Medical Imaging* 23(7).
- [Guskov, et al., 2000] GUSKOV, I., VIDIMČE, K., SWELDENS, W., AND SCHRÖDER, P. 2000. [Normal meshes](#). *ACM SIGGRAPH 2000*, pp. 95-102.
- [Guskov, 2002] GUSKOV, I. 2002. [An anisotropic mesh parameterization scheme](#). *Proceedings of the 11th International Meshing Roundtable*, pp. 325-332.
- [Guskov et al., 2002] GUSKOV, I., KHODAKOVSKY, A., SCHRÖDER, P., AND SWELDENS, W. 2002. [Hybrid meshes: multiresolution using regular and irregular refinement](#). *ACM Symposium on Computational Geometry 2002*, pp. 264-272.
- [Haker et al., 2000] HAKER, S., ANGENENT, S., TANNENBAUM, S., KIKINIS, R., SAPIRO, G., AND HALLE, M. 2000. [Conformal surface parametrization for texture mapping](#). *IEEE TVCG, 6(2)*, pp. 181-189.
- [He et al., 2005] HE, Y., GU, X., AND QIN, H., 2005, Rational spherical splines for genus zero shape modeling, Proceedings of IEEE International Conference on Shape Modeling and Applications 2005
- [Hoppe and Praun, 2005] HOPPE, H., AND PRAUN, E., 2005. [Shape compression using spherical geometry images](#), *Advances in multiresolution for geometric modeling*, pp. 3 -26.
- [Hormann and Greiner, 2000] HORMANN, K., AND GREINER, G. 2000. [MIPS: An efficient global parametrization method](#). *Curve and Surface Design*, pp. 153-162.
- [Hormann, 2000] HORMANN, K. 2000. Fitting free form surfaces. *Principles of 3d image analysis and synthesis*.
- [Hormann, et al., 1999] HORMANN, K., GREINER, G., AND CAMPAGNA, S. 1999. [Hierarchical parametrization of triangulated surfaces](#). *Vision, Modeling, and Visualization*, pp. 219-226.

- [Hurdal et al., 1999] HURDAL, M. K., BOWERS, P. L., STEPHENSON, K., SUMNERS, D. W. L., REHM, K., SCHAPER, K., ROTTENBERG, D. A. 1999. Quasi-conformally flat mapping the human cerebellum, *Medical Image Computing and Computer-Assisted Intervention 1679*, pp. 279-286.
- [Igarashi et al., 2001] IGARASHI, T., AND COSGROVE, D. 2001. [Adaptive Unwrapping for Interactive Texture Painting](#). *ACM Symposium on Interactive 3D Graphics*, pp. 209 – 216.
- [Isenburg et al., 2001] ISENBURG, M. GUMHOLD, S. AND GOTSMAN, C. 2001. [Connectivity Shapes](#), *Proceedings of IEEE Visualization*, pp. 135-142.
- [Jin et al., 2006] JIN, M., LUO, F., AND GU, X., 2006, Computing Surface Hyperbolic Structure and Real Projective Structure, *Solid and Physics Modeling 2006*
- [Julius et al., 2006] JULIUS D., KRAEVOY, AND SHEFFER, A. 2005. D_n-charts: Quasi-developable mesh segmentation., *Computer Graphics Forum (Proceedings of Eurographics 2005)*, 24(3), pp. 981-990.:581–590.,
- [Karni et al., 2005] KARNI, Z., GOTSMAN, C., GORTLER, S., 2005. [Free-boundary linear parameterization of 3D meshes in the presence of constraints](#), *Proceedings of Shape Modeling and Applications 2005*.
- [Kharevych et al., 2006] KHAREVYCH, L., SPRINGBORN, B., AND SCHRÖDER, P. 2006. [Discrete conformal mappings via circle patterns](#). *ACM Transactions on Graphics* 25(2).
- [Khodakovsky et al., 2003] KHODAKOVSKY, A., LITKE, N., AND SCHRÖDER, P. 2003. [Globally smooth parameterizations with low distortion](#). *ACM SIGGRAPH 2003*, pp. 350-357.
- [Kobbelt et al., 1999] KOBBELT, L.P., VORSATZ, J., LABISK, U. AND SEIDEL, H.-P. 1999. A shrink-wrapping approach to remeshing polygonal surfaces. *Proceedings of Eurographics*.
- [Kos and Varady, 2003] KOS, G. AND VARADY, T. 2003. Parameterizing complex triangular meshes. *Curve and surface design*, pp. 265-274.
- [Kraevoy et al., 2003] KRAEVOY, V., SHEFFER, A. AND GOTSMAN, C. 2003. [Matchmaker: constructing constrained texture maps](#). *ACM SIGGRAPH 2003*, pp. 326-333.
- [Kraevoy and Sheffer, 2004] KRAEVOY, V., AND SHEFFER, A. 2004. [Cross-parameterization and compatible remeshing of 3D models](#). *ACM SIGGRAPH 2004*.
- [Kraevoy and Sheffer, 2005] KRAEVOY, V., AND SHEFFER, A. 2005., Template Based Mesh Completion, *Proc. Symposium on Geometry Processing (SGP)*.
- [Lazarus et al., 2001] LAZARUS, F., POCCHIOLA, M., VEGTER, G., AND VERRONST, A. 2001. Computing a canonical polygonal schema of an orientable triangulated surface. In *Proceedings of the Seventeenth Annual Symposium on Computational Geometry (Medford, Massachusetts, United States)*.
- [Labsik et al., 2000] LABSIK, U., HORMANN, K., GREINER, G. 2000. Using most isometric parametrizations for remeshing polygonal surfaces. *Geometric Modeling and Processing*.
- [Lee et al., 1998] LEE, A., SWELDENS, W., SCHRÖDER, P., COWSAR, L., AND DOBKIN, D. 1998. [MAPS: Multiresolution adaptive parameterization of surfaces](#). *ACM SIGGRAPH '98*, pp. 85-94.
- [Lee et al., 1999] LEE, A., DOBKIN, D., SWELDENS, W., AND SCHRÖDER, P. 1999. [Multiresolution mesh morphing](#). *ACM SIGGRAPH '99*, 343-350.
- [Lee et al., 2000] LEE, A., HOPPE, H., MORETON, H. 2000. Displaced Subdivision Surfaces. *ACM SIGGRAPH 2000*.
- [Lee et al., 2002] LEE, Y., KIM, H. S., AND LEE, S. 2002. Mesh parameterization with a virtual boundary. *Computers And Graphics* 26(5), pp. 677-686.
- [Lengyel et al., 2001] LENGYEL, J., PRAUN, E., FINKELSTEIN, A., AND HOPPE, H. 2001 [Real-time fur on arbitrary surfaces](#), *Symposium on Interactive 3D Graphics*, pp.227-232.
- [Lévy, 2001] LÉVY, B. 2001. [Constrained texture mapping for polygonal meshes](#). *ACM SIGGRAPH 2001*, pp. 417-424.
- [Lévy and Mallet, 1998] LÉVY, B. AND MALLET, J.-L. 1998. Non-distorted texture mapping for sheared triangulated meshes. *ACM SIGGRAPH '98*, pp. 343-352.
- [Lévy et al., 2002] LÉVY, B., PETITJEAN, S., RAY, N., AND MAILLOT, J. 2002. [Least squares conformal maps for automatic texture atlas generation](#). *ACM SIGGRAPH 2002*, pp. 362-371.
- [Lévy, 2003] LÉVY, B. 2003. [Dual domain extrapolation](#). *ACM SIGGRAPH 2003*.
- [Li et al., 2006] LI, W. C., RAY, N., AND LEVY, B., Automatic and Interactive Mesh to T-Spline Conversion , 2006, *EG/ACM Symposium on Geometry Processing*
- [Litke et al., 2005] LITKE, N., DROSKE, M., RUMPF, M., SCHRÖDER, P., 2005, An Image Processing Approach to Surface Matching., *Proceedings of the Symposium on Geometry Processing 2005*
- [Losasso et al., 2003] LOSASSO, F., HOPPE, H., SCHAEFER, S., AND WARREN, J. 2003 [Smooth geometry images](#). *Eurographics Symposium on Geometry Processing*, pp. 138-145.
- [Lovasz and Schrijver, 1999] LOVASZ, L. AND SCHRIJVER, A. 1999. [On the null space of a Colin de Verdière matrix](#). *Annales de l'Institut Fourier* 40(3), pp. 1017-1026
- [Maillot et al., 1993] MAILLOT, J., YAHIA, H., AND VERRONST, A. 1993. [Interactive texture mapping](#). *ACM SIGGRAPH '93*, pp. 27-34.
- [Marschner et al., 2000] MARSCHNER, S., GUENTER, B., AND RAGUPATHY, S. 2000. [Modeling and rendering for realistic facial animation](#). *Proceedings of the Eurographics Workshop on Rendering Techniques*, pp. 231-242.
- [McCartney et al., 1999] MCCARTNEY J., HINDS, B. K., SEOW B. L. 1999 The flattening of triangulated surfaces incorporating darts and gussets. *Computer-Aided Design* 31,(4) pp. 249–260.
- [Milenkovic, 1998] MILENKOVIC, V. J. 1998. [Rotational polygon containment and minimum enclosure](#). *ACM Symposium on Computational Geometry*.

- [Mitani and Suzuki, 2004] MITANI J., AND SUZUKI H., 2004.: Making papercraft toys from meshes using strip-based approximate unfolding. *ACM SIGGRAPH 2004*, pp. 259–263.
- [Morgan, 1998] MORGAN, F., 1998. Riemann Geometry – A beginner’s guide. *A. K. Peters, Wellesley, Massachusetts, Second Edition*.
- [Ni et al., 2004] NI, X., GARLAND, M., AND HART, J. C. 2004. [Fair Morse functions for extracting the topological structure of a surface mesh](#). *ACM SIGGRAPH 2004*, pp. 613–622.
- [Pach and Wenger, 1998] PACH, J. AND WENGER, R. 1998. [Embedding planar graphs at fixed vertex locations](#). *Graph Drawing*, pp. 263-274.
- [Pederson, 1995] PEDERSON, H.-K. 1995. [Decorating Implicit Surfaces](#). *ACM SIGGRAPH '95*.
- [Peng et al., 2004] PENG, J., KRISTJANSSON, D., AND ZORIN, D. 2004. [Interactive modeling of Topologically Complex Geometric Detail](#). *ACM SIGGRAPH 2004*.
- [Pinkall and Polthier, 1993] PINKALL, U. AND POLTHIER, K. 1993. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics 2(1)*, pp. 15-36.
- [Piponi and Borshukov, 2000] PIPONI, D. AND BORSHUKOV, G. 2000. [Seamless texture mapping of subdivision surfaces by model pelting and texture blending](#). *ACM SIGGRAPH 2000*, pp. 471-478.
- [Porumbescu, et al., 2005] PORUMBESCU, S, BUDGE, B., FENG, L., AND JOY, K. I. 2005. [Shell Maps](#). *ACM SIGGRAPH 2005*.
- [Praun et al., 2000] PRAUN, E., FINKELSTEIN, A., AND HOPPE, H. 2000. [Lapped Textures](#). *ACM SIGGRAPH 2000*, pp. 465 – 470.
- [Praun and Hoppe, 2003] PRAUN, E. AND HOPPE, H. 2003. [Spherical parameterization and remeshing](#). *ACM SIGGRAPH 2003*, pp. 340-350.
- [Praun et al., 2001] PRAUN, E., SWELDENS, W. AND SCHRÖDER, P. 2001. [Consistent mesh parameterizations](#). *ACM SIGGRAPH 2001*, pp. 179-184.
- [Purnomo et al., 2004] PURNOMO, B., COHEN, J. AND KUMAR, S. 2004. [Seamless Texture Atlases](#). *Eurographics/ACM SIGGRAPH Symposium on Geometry Processing 2004*, pp. 67-76.
- [Ray and Lévy, 2003] RAY, N. AND LÉVY, B. 2003. Hierarchical least squares conformal maps. *Proceedings of the 11th Pacific Conference on Computer Graphics and Applications*, pp. 263-270.
- [Ray et al., 2006] RAY, N., LI, W-C., LÉVY, B., SHEFFER, A., AND ALLIEZ, P. 2006, [Periodic global parameterization](#). *ACM Transactions on Graphic*, 25(3).
- [Rockwood and Park, 1999] ROCKWOOD, A., AND PARK, H. 1999. [Interactive design of smooth genus N objects using multiperiodic functions and applications](#). *International Journal of Shape Modeling*, 5, pp. 135–157.
- [Rodin and Sullivan, 1987] RODIN, B., AND SULLIVAN, D. 1987. The Convergence of Circle Packings to the Riemann Mapping. *Journal of Differential Geom.* 26(2), pp. 349–360.
- [Saba et al., 2005] SABA, S., YAVNEH, I., GOTSMAN, C. AND SHEFFER, A. 2005. [Practical Spherical Embedding of Manifold Triangle Meshes](#). *Proceedings of Shape Modeling International*.
- [Sander et al., 2001] SANDER, P., SNYDER, J., GORTLER, S., AND HOPPE, H. 2001. [Texture mapping progressive meshes](#). *ACM SIGGRAPH 2001*, pp. 409-416.
- [Sander et al., 2002] SANDER, P., GORTLER, S., SNYDER, J., AND HOPPE, H. 2002. [Signal-specialized parametrization](#). *Eurographics Workshop on Rendering*, 87-100.
- [Sander et al., 2003] SANDER, P., WOOD, Z., GORTLER, S., SNYDER, J. AND HOPPE, H. 2003. [Multi-chart geometry images](#). *ACM Symposium on Geometry Processing 2003*.
- [Schreiner et al., 2004] SCHREINER, J., ASIRVATHAM, A., PRAUN, E. AND HOPPE, H. 2004. [Inter-Surface Mapping](#). *ACM SIGGRAPH 2004*.
- [Shamir, 2004] SHAMIR A. 2004. [A formulation of boundary mesh segmentation](#). *International Symposium on 3D Data Processing, Visualization, and Transmission*, pp. 82-89.
- [Shapiro and Tal,1998] SHAPIRO, A. AND TAL,A. 1998. [Polyhedron realization for shape transformation](#). *The Visual Computer*14(8), pp. 429 - 444.
- [Sheffer and de Sturler, 2000] SHEFFER, A., AND DE STURLER, E. 2000. Surface Parameterization for Meshing by Triangulation Flattening. *Proc. 9th International Meshing Roundtable*, pp. 161-172.
- [Sheffer and de Sturler, 2001] SHEFFER, A. AND DE STURLER, E. 2001, Parameterization of Faceted Surfaces for Meshing Using Angle Based Flattening, *Engineering with Computers* 17(3), pp. 326-337.
- [Sheffer and de Sturler, 2002] SHEFFER, A. AND DE STURLER, E. 2002. [Smoothing an Overlay Grid to Minimize Linear Distortion in Texture Mapping](#). *ACM Transaction on Graphics* 21(4), pp. 874-890.
- [Sheffer, 2002] SHEFFER, A., 2002, Spanning Tree Seams for Reducing Parameterization Distortion of Triangulated Surfaces, *Shape Modelling International*, pp. 61-66.
- [Sheffer and Hart, 2002] SHEFFER, A., AND HART, J. 2002. [Seamster: Inconspicuous low-distortion texture seam layout](#). *IEEE Visualization*, pp. 291-298.
- [Sheffer et al., 2003] SHEFFER, A., GOTSMAN, C., AND DYN, N. 2003. [Robust spherical parameterization of triangular meshes](#). *4th Israel-Korea Bi-National Conf. on Geometric Modeling and Computer Graphics*, pp. 94-99.
- [Sheffer et al., 2005] SHEFFER, A., LÉVY, B., MOGILNITSKY, M., AND BOGOMYAKOV, A. 2005. [ABF++: fast and robust angle based flattening](#). *ACM Transactions on Graphics* 24(2), pp. 311-330.
- [Surazhky and Gotsman, 2003] SURAZHKY, V. AND GOTSMAN, C. 2003. [Explicit surface remeshing](#). *ACM/Eurographics Symposium on Geometry Processing*.
- [Soler et al., 2002] SOLER, C., CANI, M.-P., ANGELIDIS, A. 2002. [Hierarchical Pattern Mapping](#). *ACM SIGGRAPH 2002*, pp 673 – 680.

- [Sorkine et al., 2002] SORKINE, O., COHEN-OR, D., GOLDENTHAL, R., AND LISCHINSKI, D. 2002. [Bounded-distortion piecewise mesh parametrization](#). *IEEE Visualization*, 355-362.
- [Sorkine et al., 2004] SORKINE, O., COHEN-OR, D., LIPMAN, Y., ALEXA, M., RÖSSL, C., AND SEIDEL, H.-P. 2004. [Laplacian Surface Editing](#). *Proceedings of Eurographics/ACM Symposium on Geometry Processing*, pp. 179-188.
- [Sumner and Popovic, 2004] SUMNER, R., AND POPOVIC, J. 2004. [Deformation transfer for triangle meshes](#). *ACM SIGGRAPH 2004*, pp. 399-405.
- [Tarini et al., 2004] TARINI, M., HORMANN, K., CIGNONI, P., AND MONTANI, C. 2004. Poly-cube maps. *ACM SIGGRAPH 2004*, pp.853-860.
- [Tewari et al., 2004] TEWARI, G., SNYDER, J., SANDER, P., GORTLER, S., AND HOPPE, H. 2004. [Signal-Specialized Parameterization for Piecewise Linear Reconstruction](#). *Eurographics Symposium on Geometry Processing*, pp.57-66.
- [Tong et al., 2006] TONG, Y., ALLIEZ, P., COHEN-STEINER, D., AND DESBRUN, M. 2006. [Designing Quadrangulations with Discrete Harmonic Forms](#). *Symposium on Geometry Processing*
- [Turk, 2001] TURK, G. 2001 [Texture synthesis on surfaces](#). *ACM SIGGRAPH 2001*, pp. 347-354.
- [Tutte, 1963] TUTTE, W., T. 1963. [How to draw a graph](#). *London Mathematical Society 1963*, 13, pp. 743-768.
- [Wei and Levoy, 2001] WEI, L.-Y., AND LEVOY, M. 2001. [Texture Synthesis over Arbitrary Manifold Surfaces](#), *ACM SIGGRAPH 2001*.
- [Wikipedia] WIKIPEDIA, NORMAL MAPPING. http://en.wikipedia.org/wiki/Normal_mapping
- [Ying et al., 2001] YING, L., HERTZMANN, A., BIERMANN, H., ZORIN, D. 2001. [Texture and Shape Synthesis on Surfaces](#). *Eurographics Workshop on Rendering*.
- [Yoshizawa et al., 2004] YOSHIZAWA, S., BELYAEV, A., AND SEIDEL, H.-P. 2004. [A fast and simple stretch-minimizing mesh parameterization](#). *Shape Modeling International*.
- [Zayer et al., 2003] ZAYER, R., RÖSSL, C., AND SEIDEL, H.-P. 2003. Variations on angle based flattening. *Proceedings of Multiresolution in Geometric Modelling*, pp. 285-296.
- [Zayer et al., 2005a] ZAYER, R., RÖSSL, C., AND SEIDEL, H.-P. 2005. Discrete Tensorial Quasi-Harmonic Maps *Proceedings of Shape Modeling and Applications*, pp.276-285.
- [Zayer et al., 2005b] ZAYER, R., RÖSSL, C., AND SEIDEL, H.-P. 2005. Setting the Boundary Free: A Composite Approach to Surface Parameterization. *Symposium on Geometry Processing* , pp. 91-100.
- [Zayer et al., 2006] ZZAYER, R., RÖSSL, C., AND SEIDEL, H.-P. 2006. Curvilinear Spherical Parameterization. *Proceedings of Shape Modeling and Applications*, pp. 57-64.
- [Zigelman et al., 2002] ZIGELMAN, G., KIMMEL, R., AND KIRYATI, N. 2002. Texture mapping using surface flattening via multi-dimensional scaling. *IEEE Transactions on Visualization and Computer Graphics 8(2)*, pp. 198-207.
- [Zhang et al., 2005] ZHANG, E., MISCHAIKOW, AND TURK, G. 2005. [Feature-Based Surface Parameterization and Texture Mapping](#), *ACM Transaction on Graphics 24(1)*, pp 1-27.
- [Zhou et al., 2004] ZHOU, K., SNYDER, J., GUO, B., AND SHUM, H.-Y. 2004. [Iso-charts: Stretch-driven Mesh Parameterization using Spectral Analysis](#). *Eurographics Symposium on Geometry Processing*, pp. 47-56.
- [Zhou et al., 2005] ZHOU, K., WANG, X., TONG, Y., DESBRUN, M., GUO, B., SHUM, H.-Y. 2005. [TextureMontage: Seamless Texturing of Arbitrary Surfaces From Multiple Images](#). *ACM SIGGRAPH 2005*.