

Variational Methods in Surface Parameterization

Thesis by
Nathan J. Litke

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy



California Institute of Technology
Pasadena, California

2005
(Defended 17 May 2005)

© 2005

Nathan J. Litke

All Rights Reserved

For my parents, Richard and Doris Litke.

Acknowledgements

There are many people whose guidance and support made this work possible. First and foremost, I thank my advisor, Peter Schröder, for his wisdom and generosity. Through his kind guidance I have learned invaluable skills which have shaped me as a researcher and as a scientist. For these experiences I am forever grateful.

I am also deeply indebted to my thesis committee whose encouragement and enthusiasm not only elevated this work but have also profoundly affected me personally and philosophically. Alan Barr has been a voice of sage counsel, whose timely advice has gently steered me along an enlightened path. Mathieu Desbrun has graciously shared his brilliant intellect to help me overcome even my most trivial conundrums, and his kindness and optimism have been felt far beyond my academic pursuits. Martin Rumpf has been my colleague, my mentor and my friend through many exciting discoveries, and his kind hospitality during my time in Germany made these months the highlight of my graduate career.

I have been extremely fortunate to work with many exceptionally talented and gifted researchers over the years. Adi Levin, Andrei Khodakovsky, Ulrich Clarenz and Marc Droske have been a pleasure to work with and have enriched my experiences both in research and personally. My peers in the Multires Modeling Group and abroad made the journey fulfilling and entertaining, and the support that I received from my friends has truly provided me with the experience of a lifetime. Eitan Grinspun, Liliya Kharevych, Catherine Mastromarino, Gerhard Bendels, Mark Meyer, Tran Gieng, Zoë Wood, Sharif Elcott, Ilja Friedel, Steven Schkolne and Boris Springborn have left me with countless unforgettable memories.

Finally, and most importantly, I thank my family and those who have become part of my family throughout the years at Caltech. Your love and encouragement kept me sane and grounded in this wonderful chapter of my life.

Abstract

A surface parameterization is a function that maps coordinates in a 2-dimensional parameter space to points on a surface. This thesis investigates two kinds of parameterizations for surfaces that are disc-like in shape. The first is a map from a region of the plane to the surface. The second is a mapping from one surface to another, which defines a correspondence between them. The main challenge in both cases is the construction of a smooth map with low distortion. In this thesis we present a variational approach to surface parameterization that addresses these challenges.

The first contribution of this thesis is the development of a variational framework for parameterizations. This framework encompasses the mapping of a region of the plane to a surface that is isomorphic to a disc, and the mapping between such surfaces. It is based on the rich mathematical theory built up over decades in the study of rational mechanics. Because of its roots in mechanics, our parameterizations are guaranteed to be smooth and locally bijective, and optimal parameterizations which minimize a variational energy are known to exist. A proof of existence is given for the case of optimal parameterizations in the plane.

Our second contribution is a set of algorithms to construct parameterizations for surface triangulations. We describe in detail free-boundary methods that use standard numerical optimization algorithms for the computation of optimal parameterizations. A flexible set of parameters is offered to the user to formulate preferences for the trade-off between angle, area and length distortion in parameterizations in the plane. In the specification of a correspondence between surfaces, we provide user control through feature lines which are mapped as sets onto corresponding feature lines. Additionally we allow for a partial correspondence of the surfaces which is particularly important for correlating surfaces with boundaries.

Our third contribution is an analysis of the performance of the algorithms based on our implementations. Our testing focuses on parameterizations of physically-acquired triangle mesh data. The efficiency of our methods is measured by analyzing the rate of convergence of the energy minimization, and execution times are shown to be quite reasonable. Robustness is established in the

presence of large deformations in the parameterizations and the stability of our parameterizations is demonstrated under different discretizations of the surfaces.

Our fourth contribution is a set of concrete, compelling applications of surface parameterization. Non-trivial examples which draw from texture mapping, morphing and facial animation provide further evidence and insight into the versatility of our parameterization framework.

Contents

Acknowledgements	iv
Abstract	v
1 Introduction	1
1.1 Related Work	3
1.2 Contributions	6
1.3 Overview	7
2 Variational Framework	8
2.1 Surface Parameterization	9
2.1.1 Definition of a Parameterization	9
2.1.2 Measuring Distortion in a Parameterization	10
2.1.3 Measuring Distortion in a Deformation	10
2.1.4 A Variational Approach from First Principles	12
2.1.5 Minimal Energy Parameterizations	14
2.1.6 Existence of Optimal Parameterizations	15
2.2 Surface Matching	17
2.2.1 A Physical Interpretation of the Matching Deformation	17
2.2.2 Measuring Distortion in a Deformation	18
2.2.3 Measuring Bending in a Deformation	20
2.2.4 Matching Features	21
2.2.5 Partial Correspondence	22
2.2.6 Definition of the Matching Energy	23

3	Algorithms	25
3.1	The Parameterization Algorithm	26
3.1.1	Initial Parameterization	27
3.1.2	Finite Element Discretization	28
3.1.3	Numerical Optimization	29
3.2	The Matching Algorithm	32
3.2.1	Surface Patch Parameterization	32
3.2.2	Feature Set Construction	34
3.2.3	Evaluation of Surface Properties	35
3.2.4	Multiscale Finite Element Method	36
4	Analysis	38
4.1	Surface Parameterization	39
4.1.1	Convergence of the Optimization Algorithm	39
4.1.2	Convergence of the Continuation Method	41
4.1.3	Robustness to Surface Discretization	42
4.1.4	Distortion in Optimal Parameterizations	44
4.1.5	Comparison to Related Work	44
4.2	Surface Matching	45
4.2.1	Convergence of the Optimization Algorithm	46
4.2.2	Robustness to Image Discretization	47
4.2.3	Partial Correspondence	47
4.2.4	Comparison to Related Work	48
5	Applications	51
5.1	Texture Mapping	52
5.2	3D Morphing	54
5.3	Facial Animation	56
6	Conclusion	58
6.1	Summary	59
6.2	Future Work	60

A	Surface Parameterization	62
A.1	Representation of the Parameterization Energy	62
A.2	Existence of Optimal Parameterizations	65
A.3	Variations of the Parameterization Energy	66
B	Surface Matching	68
B.1	Shape Operators	68
B.2	Variations of the Matching Energy	69
	Bibliography	72

List of Figures

2.1	The optimal parameterization function $x[\phi]$	11
2.2	A physical interpretation of the matching deformation	18
2.3	The matching function ϕ_M	19
3.1	Overview of the surface parameterization algorithm	26
3.2	The parameterization optimization algorithm	29
3.3	The step size control algorithm	30
3.4	Overview of the surface matching algorithm	33
3.5	User-defined feature sets for matching	35
3.6	Geometry images of surface properties	36
3.7	The multiscale matching algorithm	37
3.8	The multigrid pyramid	37
4.1	Energy decay from the parameterization optimization	40
4.2	Sequence of parameterization deformations	41
4.3	Successive refinement of a triangulation	42
4.4	Trade-off between length, area and angle distortion	43
4.5	Flattening the earth	43
4.6	Comparison of parameterizations to related work	44
4.7	Energy decay from the matching optimization	45
4.8	Sequence of multiscale matching deformations	46
4.9	Large deformations in a match	48
4.10	Gross misalignment in the image domain	48
4.11	Partial correspondence	49
5.1	Photographs used for the texture image	52
5.2	Texture mapping 3D scan data	53
5.3	3D morphing between surface pairs	55
5.4	Morphing through keyframe poses	57

5.5 Facial animation using blend shapes 57

List of Tables

4.1	Parameterization performance figures	39
4.2	Performance figures for the continuation method	42
4.3	Matching performance figures	47

Chapter 1

Introduction

Parametric surfaces have a long and illustrious history in the fields of computer graphics, computer-aided geometric design and engineering. Their success is largely due to the ease in which they represent the physical, analytic and visual properties of surfaces in space, *i.e.*, the boundaries of three-dimensional shapes, which makes them amenable to simulation, manufacturing and visualization applications. Early boundary representations were built from piecewise-smooth parametric patches, which provide flexibility in design but have practical limitations in the level of detail that they can model. The advent of laser range scan technology in the last decade has provided the world with high-resolution models of physical objects with millions to billions of sample points. However, the triangulated surfaces generated from this data have no inherent parameterization, and thus parameterizing these surfaces has been an active area of research in recent years.

In this thesis we investigate parameterizations of surface triangulations in two contexts. The first context considers parameterizations which are defined as mappings from a disc-like region of the plane to the embedded surface. We present a rigorous variational framework for the construction of low-distortion parameterizations with free boundaries, based on the rich mathematical theory built up over decades in the study of rational mechanics. With a view toward the theory of elasticity, we treat the parameterization problem as one of finding an optimal deformation of an initial parameter domain in the plane. Starting with some basic axioms in the continuous setting, we derive a deformation energy which provides combined control over angle, area and length distortion in a unified framework. Because of its roots in mechanics, our method comes with analytic guarantees: our parameterizations are smooth and locally bijective, and we provide proof of the existence of optimal parameterizations that minimize our energy. Our continuous model of the parameterization problem lets us use straightforward finite element methods to discretize our energy. This in turn allows us to leverage the extensive and mature body of work in numerical optimization to construct optimal parameterizations using a free-boundary method that uses standard numerical optimization algorithms for the energy minimization, with well-known performance guarantees. The finite element approach allows us to retain analytic guarantees of smoothness and bijectivity for the discrete parameterizations. Furthermore we provide a proof of convergence to the continuous solution in the limit of triangle grid refinement. The performance of our parameterization algorithm is demonstrated on surfaces reconstructed from physical data measurements. This data is particularly difficult to process due to its geometric complexity and the high variability of the data samples. We use a range of energy parameters to highlight the flexibility and variety of parameterizations that one can achieve with our energy. The resulting parameterizations have very low angle, area and length distortion overall and the different distortion measurements reflect the trade-offs in the choice of parameters.

In the second context we describe a method to find a correspondence between two surfaces that are isomorphic to a disc. This correspondence can also be interpreted as a parameterization of a region of a surface over a manifold domain embedded in space. Our approach draws upon well-developed methods in global image matching to correlate two surfaces via a non-rigid deformation in their parameter domains. To do so we map geometric characteristics into the parameter plane and modify the standard image matching energies to correctly account for the geometry of the surfaces. This decouples the discretization of the surface from the discretization of the matching deformation, and thus the level of detail in the correspondence mapping can be chosen independently of the resolution of the surface triangulation. We present a multiresolution algorithm using regular image grids

that allows us to overcome problems with robustness and computational efficiency seen in previous methods that pursue mappings between surfaces directly in the embedding space. User control over the match is given through feature lines which are mapped as sets onto corresponding feature lines, which provides the user with a wide variety of techniques to fine-tune the correspondence if desired. We allow for a partial correspondence of the surfaces, which is particularly important for matching surfaces of disc topology, especially in boundary regions. The existence and global injectivity of the matching deformations is guaranteed such that the resulting deformations are smooth and bijective.

1.1 Related Work

The body of relevant literature is quite broad due to the wide span of fields in which parameterizations are studied. Most relevant to our treatment of parameterizations on triangulated surfaces is the recent body of research from the graphics literature. In the context of correspondences between surfaces, we draw upon work in image matching, in particular the non-linear approaches which deal directly with the large deformation setting. Relevant work from the graphics literature covers approaches which pursue direct mappings between surfaces in \mathbb{R}^3 .

Parameterization

The most desirable property to achieve in a parameterization is isometry. This implies intuitively that all of the properties of the surface are represented in the corresponding parameter domain. Strictly speaking, a map between two surfaces is an isometry if their first fundamental forms coincide. It is well known that isometric parameterizations exist only if the surface itself is locally flat, *i.e.* developable. A broad variety of algorithms have been proposed to construct such parameterizations for embedded triangle meshes (for a recent survey see the comprehensive overview by Floater and Hormann [19]). Generically these algorithms are distinguished by the way they measure distortion.

Distortion is a measurement of change in the first fundamental form of the mapping between the surface and the parameter domain, *i.e.*, how much the metric induced by the parameterization deviates from the identity. There are many possible ways to measure deviation from the identity. For example, Sander *et al.* consider the minimization of a variety of norms (ℓ^2 , ℓ^∞) of the *singular values* of the Jacobian [44], while Hormann and Greiner minimize the *condition number* of the Jacobian [25]. Degener *et al.* extend this work by adding a term that penalizes area distortion [14].

The *maximum eigenvalue* of the Jacobian and its inverse were used by Sorkine *et al.* [47]. These approaches all involve difficult algorithms and it is unclear what guarantees can be made of the resulting parameterizations.

Other free-boundary methods are primarily concerned with controlling angle distortion. On one hand we find methods which attempt to construct conformal parameterizations directly. These include the minimization of harmonic energy with natural boundary conditions [15], which turns out to be equivalent to a least-squares optimal discrete approximation of conformal parameterizations [29]. The case of closed, arbitrary genus surfaces (possibly after constructing a double cover) was treated by Gu and Yau [24] based on discrete Riemann surface theory [33]. Direct minimization of the change of angle was pursued by Sheffer and de Sturler [46], who used a highly non-linear minimization procedure. Kharevych *et al.* approached discrete conformal parameterizations by deriving a convex energy based on the intersection angles of circles on the surface triangulation [26]. In this class of methods the focus is entirely on controlling angle distortion, with the exception of Desbrun *et al.* who also define a separate authalic energy to control the change of area (although only in the presence of Dirichlet boundary conditions).

In all of these approaches, the interaction between area, angle and length distortion is not easily controlled or well understood from a mathematical perspective, *e.g.*, little is known about the existence of solutions. With the exception of methods based on harmonic or conformal maps [15, 24, 26, 29], most of the previous work is formulated in the discrete setting of surface triangulations where these results are unknown.

Image Matching

In image processing, registration is often approached as a variational problem. One asks for a deformation which maps structures in the reference image A onto corresponding structures in the template image B on some image domain ω . In the case of unimodal images with a direct correspondence of the image intensities I_A and I_B , the energy $\int_{\omega} (I_B(\xi) - I_A(\phi(\xi)))^2 d\xi$ measures the least-squares error of the match. We extend this idea to surface matching through a bending energy which measures the matching defect with respect to curvatures. It is well established that the associated minimization problem is ill-posed if one considers the infinite dimensional space of deformations [7, 48]. This is generally addressed by choosing a suitable regularization. Motivated by models from continuum mechanics, one may ask for a deformation that is additionally controlled by *elastic stresses* on images regarded as elastic sheets. For example see the early work of Bajcsy

and Broit [3] and more recent, significant extensions by Grenander and Miller [20]. In our surface matching problem, we consider surfaces as thin shells. Besides the bending which we mentioned, surface deformations also lead to tangential stretching and shearing, which gives a real physical interpretation to the elastic stresses that are treated as a regularization in the resulting model. In particular, if large displacements are necessary to ensure a proper match, a regularization based on non-linear elasticity with its built-in control of length, area and volume changes is indispensable. Cohen [11] considered polyconvex elastic functionals and Droske and Rumpf [18] used this type of regularization to guarantee global injectivity and well-posedness. We incorporate these ideas to avoid folding in our surface matches. In essence, non-rigid image matching is a well developed and powerful tool which we will exploit for surface matching.

3D Registration and Correspondence

Motivated by the ability to scan geometry with high fidelity, a number of approaches have been developed in the graphics literature to bring such scans into correspondence. Early work used parameterizations of the meshes over a common parameter domain to establish a direct correspondence between them [28]. Typically these methods are driven by user-supplied feature correspondences which are then used to drive a mutual parameterization. The main difficulty is the management of the proper alignment of selected features during the parameterization process [27, 43, 45] and the algorithmic issues associated with the management of irregular meshes and their effective overlay.

Special methods have been developed for situations in which a large number of scans of similar objects, albeit with great geometric variety, are to be brought into correspondence, for example for purposes of statistical analysis. These are typically geared towards establishing a correspondence against a template model. Blanz and Vetter [5] used cylindrical scans resulting in height “images” which were matched through a modified optical flow. Allen *et al.* [1] fit a high resolution template mesh to scans of the human body. They computed non-rigid deformations for the template by minimizing an error functional which performs well in the presence of holes and poorly sampled data, provided that the two surfaces are in similar poses. Such template-based approaches can also be very helpful during the acquisition itself. Zhang *et al.* [50] present a method for meshing dynamic range data using a surface fitting approach. In their method, a template mesh is fitted to a registered stereo pair of depth maps and the fitting is achieved by minimizing a depth matching energy and a regularization energy. Recently, Gu and Vemuri [23] considered matches of topological spheres through conformal maps with applications to brain matching. Their energy measures the defect of

the conformal factor and — similar to our approach — the defect of the mean curvature. However they do not measure the correspondence of feature sets or tangential distortion, and thus do not involve a regularization energy for the ill-posed energy minimization. Furthermore, they seek a one-to-one correspondence, whereas we must address the difficult problem of partial correspondences between surfaces with boundaries.

1.2 Contributions

The contributions documented in this thesis summarizes work performed by the author and collaborators in papers [10, 17, 30, 31]. The theoretical foundation for the variational methods was largely investigated by Dr. Martin Rumpf, with additional collaboration by Dr. Ulrich Clarenz, Marc Droske and the author. The theorems and their proofs contained herein are the work of Dr. Martin Rumpf. Some parts of the implementation of the surface matching algorithm are owned by Marc Droske.

The first contribution of this thesis is the development of a variational framework for parameterizations. This framework encompasses the mapping of a region of the plane to a surface that is isomorphic to a disc, and the mapping between such surfaces. It is based on the rich mathematical theory built up over decades in the study of rational mechanics. Because of its roots in mechanics, our parameterizations are guaranteed to be smooth and locally bijective, and optimal parameterizations which minimize a variational energy are known to exist. A proof of existence is given in this thesis for the case of optimal parameterizations in the plane.

Our second contribution is a set of algorithms to construct parameterizations for surface triangulations. We describe in detail free-boundary methods that use standard numerical optimization algorithms for the computation of optimal parameterizations. A flexible set of parameters is offered to the user to formulate preferences for the trade-off between angle, area and length distortion in parameterizations in the plane. In the specification of a correspondence between surfaces, we provide user control through feature lines which are mapped as sets onto corresponding feature lines. Additionally we allow for a partial correspondence of the surfaces which is particularly important for correlating surfaces with boundaries.

Our third contribution is an analysis of the performance of the algorithms based on our implementations. Our testing focuses on parameterizations of physically-acquired triangle mesh data. The efficiency of our methods is measured by analyzing the rate of convergence of the energy minimization, and execution times are shown to be quite reasonable. Robustness is established in the

presence of large deformations in the parameterizations and our parameterizations are stable under different discretizations of the surfaces.

Our fourth contribution is a set of concrete, compelling applications of surface parameterization. Non-trivial examples which draw from texture mapping, morphing and facial animation provide further evidence and insight into the versatility of our parameterization framework.

1.3 Overview

This thesis is organized as follows. The next chapter presents a variational framework for parameterization in two parts. The first of these, entitled *surface parameterization*, describes a variational approach to parameterizations in the plane. The second, *surface matching*, is concerned with the correspondence mapping between two surfaces. The derivations therein define the basic components of the parameterization energies with particular attention given to the geometric interpretation of the energy integrands.

Chapter 3 describes the parameterization and matching algorithms in detail. Both procedures take user parameters that provide flexibility in the optimal parameterization, especially in the way it deals with distortion. The numerical optimization methods that perform the energy minimization are fully detailed to aid in implementation.

In Chapter 4 we analyze the performance and robustness of the algorithms through a series of tests using physically-acquired data. Chapter 5 describes a variety of compelling applications of surface parameterization. This thesis concludes with a summary of our contributions and a discussion of avenues for future research.

Chapter 2

Variational Framework

This chapter develops the variational framework for surface parameterizations. We begin with parameterizations that map disc-like surfaces to the plane in Section 2.1, and then turn our attention to mappings from one embedded surface to another in Section 2.2. In our approach we consider parameterizations as deformations of the surface. For parameterizations in the plane, this amounts to “flattening” the surface. For defining a correspondence between two surfaces, we think of deforming one surface to match the shape of the other. In fact both points of view are compatible, whereby a parameterization in the plane can be thought of as matching to a planar surface. From this perspective, our treatment of distortion falls within a general framework that separates tangential stress from stress due to bending of the surface normals. The latter is not considered for parameterizations in the plane because the curvature of the surface goes to zero everywhere under the parameterization deformation.

2.1 Surface Parameterization

The goal of surface parameterization is the construction of a smooth map with low distortion from a given surface patch to a region of the plane. Our approach starts with an initial parameterization from the embedded surface into \mathbb{R}^2 and then *composes* this map with another from \mathbb{R}^2 to itself. In doing so, we turn the problem of defining a parameterization into finding an optimal non-rigid deformation of the parameter domain. We develop a variational approach based on insights from rational mechanics which performs an energy relaxation over a set of non-rigid deformations in the plane. This gives us a unified framework for surface parameterization that incorporates the trade-offs between different distortion measurements and provides clear mathematical statements as to the properties of this energy. In what follows, we develop this variational approach in the continuous setting, which we will later discretize to the space of surface triangulations.

2.1.1 Definition of a Parameterization

A parameterization is a mapping from the plane onto a given surface, or in the case of its inverse, from the surface onto the plane. Consider a smooth surface $\mathcal{M} \subset \mathbb{R}^3$ and a parameterization

$$x : \omega \rightarrow \mathcal{M}$$

on a parameter domain $\omega \subset \mathbb{R}^2$. For a parameterization to be properly defined, its inverse x^{-1} cannot allow the surface to fold onto itself in the plane. In this case x is locally bijective, and we say that the parameterization is *admissible*. A metric g is defined on ω ,

$$g = Dx^T Dx$$

where $Dx \in \mathbb{R}^{3,2}$ is the Jacobian of the parameterization x . Roughly speaking, the metric g is a function which relates distances in the parameter domain ω to distances on the surface \mathcal{M} . It acts on tangent vectors v, w on ω with

$$(g v) \cdot w = Dx v \cdot Dx w$$

which is simply the inner product of tangent vectors $Dx v, Dx w$ on \mathcal{M} , where products are given in matrix notation. It follows that the metric describes the distortion of an infinitesimal area of the domain under the parameterization function.

2.1.2 Measuring Distortion in a Parameterization

Let us now focus on the distortion from the surface \mathcal{M} onto the parameter domain ω under the inverse parameterization x^{-1} . This distortion is measured by the inverse metric $g^{-1} \in \mathbb{R}^{2,2}$. Just as $\sqrt{\text{tr}(A^T A)}$ measures the average *change of length* under a linear mapping A , $\sqrt{\text{tr}(g^{-1})}$ measures the average change of length of tangent vectors under the mapping from the surface onto the parameter plane. Additionally, $\sqrt{\det(g^{-1})}$ measures the corresponding *change of area*. In terms of the eigenvalues $\Gamma \geq \gamma$ of g^{-1} we obtain

$$a = \text{tr}(g^{-1}) = \Gamma + \gamma, \quad d = \det(g^{-1}) = \Gamma \gamma.$$

The condition for a conformal parameterization is $\Gamma = \gamma$, and thus

$$(\Gamma - \gamma)^2/d = \frac{\Gamma}{\gamma} + \frac{\gamma}{\Gamma} - 2 = a^2/d - 4$$

is a normalized, scale-invariant measure of the lack of conformality. The normalization is done so that a conformal mapping yields a value of zero for better numerical conditioning.

In what follows, we will use the quantities a , d and $a^2/d - 4$ to account for length, area and angle distortion, respectively. For an admissible parameterization we have $\Gamma, \gamma > 0$ everywhere and therefore $a, d > 0$ and $a^2/d - 4 \geq 0$, *i.e.*, our measures of distortion are strictly non-negative.

2.1.3 Measuring Distortion in a Deformation

In our approach, we suppose that a parameter map x of the surface patch \mathcal{M} is defined in an initial step. We will assume that x and its parameter domain ω are fixed from now on. As mentioned above, we start with this *fixed* initial parameterization $x : \omega \rightarrow \mathcal{M}$ and use an invertible Euclidean deformation $\phi : \omega \rightarrow \mathbb{R}^2$ to derive a new parameterization

$$x[\phi] = x \circ \phi^{-1}$$

over the parameter domain $\omega[\phi] := \phi(\omega)$. For the inverse we obtain $x^{-1}[\phi] = \phi \circ x^{-1}$. Obviously, any smooth parameterization \tilde{x} of the surface patch \mathcal{M} can be written as $\tilde{x} = x[\phi]$ for $\phi := \tilde{x}^{-1} \circ x$. The complete mapping is illustrated in Figure 2.1.

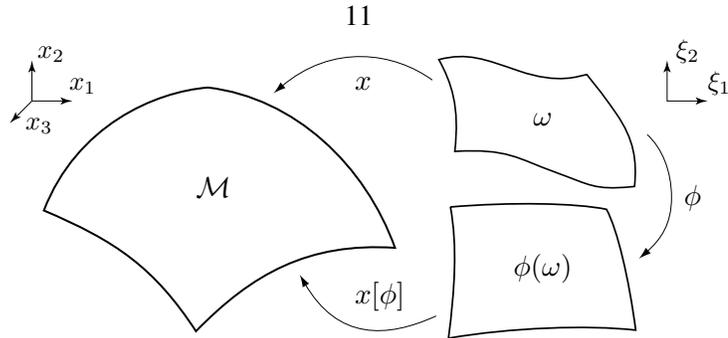


Figure 2.1: *The parameterization function $x[\phi] := x \circ \phi^{-1}$ is a mapping from the parameter domain $\omega[\phi] := \phi(\omega)$ to the surface patch \mathcal{M} .*

We will now study the distortion which arises from the deformation ϕ . To motivate our definition of distortion, for the moment let us replace the surface patch \mathcal{M} with a Euclidean domain $\Omega \subset \mathbb{R}^2$. A fundamental assumption from the theory of elasticity [9, 32] tells us that all distortion measures for a given deformation $\phi : \Omega \rightarrow \mathbb{R}^2$, such as length, area or angle distortion, can be expressed in terms of the deformation tensor $D\phi$. *Thus measuring the distortion of a mapping is directly linked to measuring the elastic energy of the associated deformation.* A fundamental result of rational mechanics states that the elastic stress induced in an isotropic elastic material by some deformation ϕ depends solely on the principal invariants of the Cauchy-Green strain tensor $\mathbf{C} = D\phi D\phi^T$. In the case of deformations in \mathbb{R}^2 , these invariants are the trace $\text{tr}(\mathbf{C})$ and determinant $\det(\mathbf{C})$. In our setting, we replace the elastic body $\Omega \subset \mathbb{R}^2$ by our surface patch \mathcal{M} and we consider parameterizations as elastic surface deformations to leverage these fundamental theorems from classical elasticity. The Cauchy-Green strain tensor we use here is defined by

$$\mathbf{C}[\phi] := D\phi g^{-1} D\phi^T$$

which generalizes the standard definition in that we obtain the classical Cauchy-Green tensor when \mathcal{M} is flat ($g = \mathbb{1}$, the Euclidean metric). Note that $(\mathbf{C}[\phi])^{-1}$ is the metric tensor $g[\phi]$ of the parameterization $x[\phi]$ induced by the deformation ϕ . With this observation we can give a geometric interpretation of the elastic surface deformation. As in the case of the initial parameterization, $a = \sqrt{\text{tr} \mathbf{C}[\phi]} = \sqrt{\text{tr}(g^{-1}[\phi])}$ is a measure of length distortion under the mapping $\phi \circ x^{-1}$ and $d = \sqrt{\det \mathbf{C}[\phi]} = \sqrt{\det(g^{-1}[\phi])}$ measures area distortion. Furthermore, $a^2/d - 4$ provides us with a measure of angle distortion. Thus $\text{tr}(\mathbf{C}[\phi])$ and $\det(\mathbf{C}[\phi])$ are natural variables for an energy that measures the distortion in a parameterization.

2.1.4 A Variational Approach from First Principles

A theory for the deformation of elastic bodies has been developed over many decades, and is covered in many classical texts [9, 32, 39]. By approaching parameterizations from the context of classical elasticity, we obtain mathematical guarantees in the continuous setting which we will later extend to triangulated surfaces via a standard finite element discretization. We will now derive our parameterization energy from the first principles of *elasticity* (2.1), *frame indifference* (2.2) and *isotropy* (2.3), taking into account our geometric setting.

- ◇ **Elasticity:** The energy depends solely on the distortion tensor $\mathbf{D}(\phi \circ x^{-1})$. Therefore the total energy of a parameterization deformation ϕ and thus of the corresponding parameterization $x[\phi]$ is given by

$$E[\phi] = \int_{\mathcal{M}} W_{\mathcal{M}}(\mathbf{D}(\phi \circ x^{-1})(x)) dA \quad (2.1)$$

where $W_{\mathcal{M}} : \mathbb{R}^{2,3} \rightarrow \mathbb{R}$ is an energy density which we will specify later.

- ◇ **Frame Indifference:** The energy is independent of Euclidean motions $r(\xi) = R\xi + b$ of the parameter domain $\phi(\omega)$, where $R \in SO(2)$ denotes a rotation and $b \in \mathbb{R}^2$ is a translation. Thus for any deformation $\phi : \omega \rightarrow \mathbb{R}^2$ we have

$$W_{\mathcal{M}}(R\mathbf{D}(\phi \circ x^{-1})) = W_{\mathcal{M}}(\mathbf{D}(\phi \circ x^{-1})) \quad (2.2)$$

where $\mathbf{D}(r \circ \phi \circ x^{-1}) = R\mathbf{D}(\phi \circ x^{-1})$ by the chain rule.

- ◇ **Isotropy:** The energy does not depend on directions on \mathcal{M} . If $u : \mathcal{M} \rightarrow \mathcal{M}$ is any smooth mapping from \mathcal{M} onto itself which is locally a rotation at some point $x \in \mathcal{M}$, then the energy density at x is not affected by this transformation. Being “locally a rotation” here means that up to first order, the deformation is a rotation Q around x in the tangent space $\mathcal{T}_x\mathcal{M}$. Using the chain rule again (*i.e.*, $\mathbf{D}(\phi \circ x^{-1} \circ u) = \mathbf{D}(\phi \circ x^{-1})Q$) we find

$$W_{\mathcal{M}}(\mathbf{D}(\phi \circ x^{-1})Q) = W_{\mathcal{M}}(\mathbf{D}(\phi \circ x^{-1})) \quad (2.3)$$

for any rotation Q on $\mathcal{T}_x\mathcal{M}$.

From these first principles we are able to deduce a representation formula for the parameterization energy. Here we make use of two additional facts. Because the initial parameterization is

fixed, we can integrate over the *parameter domain* in \mathbb{R}^2 instead of the embedded surface. It is then through the metric tensor $g \in \mathbb{R}^{2,2}$ that the initial parameterization from the region in \mathbb{R}^2 to the surface enters into our energy. Furthermore, since the elastic deformation is given as the minimizer of an elastic energy integrand, the energy density too must depend *only* on the principal invariants of $\mathbf{C}[\phi]$ and not the full deformation tensor $D\phi \in \mathbb{R}^{2,2}$. This leads us to the following general form for our parameterization energy.

Theorem 2.1 (Representation of the energy) *Let \mathcal{M} be a smooth surface patch, which is parameterized by an initial parameterization x over a set $\omega \subset \mathbb{R}^2$. Under the assumptions of elasticity (2.1), frame indifference (2.2) and isotropy (2.3), there exists a function $W : \mathbb{R}^2 \rightarrow \mathbb{R}$ such that $W_{\mathcal{M}}(\mathbf{D}(\phi \circ x^{-1})(x)) = W(\iota_{\mathbf{C}[\phi]})\sqrt{\det g}$ where $\iota_A = (\text{tr}(A), \det(A))$ are the two principal invariants of $A \in \mathbb{R}^{2,2}$ and the Cauchy-Green strain tensor is defined by*

$$\mathbf{C}[\phi] = D\phi g^{-1} D\phi^T .$$

Furthermore, the expression $g \in \mathbb{R}^{2,2}$ is the metric tensor w.r.t. the parameterization x , i.e., $g = Dx^T Dx$. Thus, the parameterization energy can be written as

$$E[\phi] = \int_{\omega} W(\text{tr}(\mathbf{C}[\phi]), \det(\mathbf{C}[\phi])) \sqrt{\det g} d\xi . \quad (2.4)$$

A proof of this theorem appears in Appendix A.1. The representation (2.4) has two significant practical advantages. First, we simplify the energy density W to a function of two variables, namely the principal invariants $\text{tr}(\mathbf{C}[\phi]), \det(\mathbf{C}[\phi])$, which reduces the computational overhead of treating the full deformation tensor $D\phi$ as a 4-dimensional variable. Second, by describing the new parameterization in terms of a deformation ϕ of the initial parameter domain ω , we avoid a free-boundary problem which would involve complicated integrals on the boundary.

It comes as no surprise that the principal invariants of the distortion tensor $\mathbf{C}[\phi]$ appear in our energy (2.4). Indeed, a number of earlier works have arrived at energies based on the principal invariants by using similar considerations in the discrete geometric setting. However as we shall see below, by deriving from these fundamental axioms of classical elasticity we are able to make concrete statements about its theoretical properties.

2.1.5 Minimal Energy Parameterizations

The previous theorem describes an energy which is supposed to capture the distortion of a parameterization as a function of the principal invariants of the Cauchy-Green strain tensor. As suggested in Section 2.1.3, our energy density should therefore be a function of $a = \text{tr}(\mathbf{C}[\phi])$ (length control), $d = \det(\mathbf{C}[\phi])$ (area control) and $a^2/d - 4$ (conformality control). Following Theorem 2.1, we propose to use the energy density

$$\boxed{W(a, d) = \alpha_l a + \alpha_a \left(d + \left(\frac{\alpha_l}{\alpha_a} + 1 \right) d^{-1} \right) + \alpha_c (a^2/d - 4)} \quad (2.5)$$

The three terms control the preservation of length, area and conformality of the chart $x \circ \phi^{-1}$, and thus simultaneously capture all of our measures of distortion. Note that the second term controls both area expansion with d and area compression through d^{-1} . The parameters α_l , α_a and α_c are chosen by the user according to the relative importance of length, area and angle distortion in the parameterization. These parameters must satisfy certain conditions to guarantee that the parameterization problem is well-posed. We start with the requirement that an isometry is an α -optimal parameterization, *i.e.*, a parameterization which minimizes the energy defined by a given set of parameters $\alpha = (\alpha_l, \alpha_a, \alpha_c)$. We will refer to these simply as “optimal parameterizations” in what follows. Based on a Taylor expansion of the energy density (2.5) close to an isometry, we show in Appendix A.1 that an isometry is at least a local minimizer of the energy if

$$\boxed{\alpha_l > 0, \quad \alpha_a > 0, \quad \alpha_c \geq 0} \quad (2.6)$$

A further analysis provides additional insight into these restrictions:

- ◊ We require a non-zero length parameter $\alpha_l > 0$ to guarantee the existence of a minimal energy deformation. Without this restriction one will not have a bounded minimizing sequence in the space of deformations, *e.g.*, without a penalty on length the energy is invariant under shearing deformations.
- ◊ A non-zero area parameter $\alpha_a > 0$ is critical to ensure local injectivity, and is directly responsible for preventing folds in the discrete setting of triangulations. Control over area compression through d^{-1} also prevents the parameterization from collapsing to a point in the domain, *i.e.*, as occurs in linear elasticity without the presence of boundary constraints.
- ◊ The balancing factor $\left(\frac{\alpha_l}{\alpha_a} + 1 \right)$ between area expansion and compression appears when one asks

for a positive definite Hessian at an energy minimum, *e.g.*, for the sake of applying standard algorithms in numerical optimization to find a minimizing deformation.

Consequently, the direct control of length and area through $\alpha_l, \alpha_a > 0$ plays a crucial role in the proof of existence.

2.1.6 Existence of Optimal Parameterizations

Guaranteeing the existence of a solution is essential for a numerical algorithm to robustly generate a parameterization. Therefore we will now give a rigorous and fairly general statement on the existence of an optimal parameterization.

Theorem 2.2 (Existence of an Optimal Parameterization) *Let \mathcal{M} be a smooth surface patch, which is parameterized by an initial parameterization x over a set $\omega \subset \mathbb{R}^2$. Furthermore, consider the energy in (2.5) over a set of admissible deformations ϕ , which are continuous and have $\det \phi > 0$. Then for parameters $\alpha_l, \alpha_a, \alpha_c$ which satisfy (2.6), there exists an admissible, minimizing deformation $\tilde{\phi}$.*

We defer the proof to Appendix A.2. The theorem states that for any smooth surface patch \mathcal{M} and initial parameterization x , there exists an optimal parameterization for any *valid* choice of parameters α satisfying (2.6). This parameterization is formulated by $x[\tilde{\phi}] = x \circ \tilde{\phi}^{-1}$, where $\tilde{\phi}$ is an *admissible* deformation that minimizes the energy (2.5). Furthermore, $\tilde{\phi}$ is in the set of *admissible* deformations, meaning that $\tilde{\phi}$ is continuous and locally bijective (*i.e.*, $\det \tilde{\phi} > 0$ everywhere). Since the initial parameterization x is required to be locally bijective, it follows that the optimal parameterization $x[\tilde{\phi}] := x \circ \tilde{\phi}^{-1}$ is guaranteed to be free of folds in its parameter domain $\omega[\tilde{\phi}] := \tilde{\phi}(\omega)$.

In certain applications, one is interested in obtaining an exactly area-preserving parameterization (*i.e.*, $\Gamma \gamma = 1$). This is summarized in the following statement.

Corollary 2.3 (Existence of an Optimal, Area-Preserving Parameterization) *Under the assumptions of Theorem 2.2, and for the modified set of admissible deformations ϕ with $\det g^{-1}[\phi] = 1$, there exists a minimizing deformation $\tilde{\phi}$.*

The proof is given in Appendix A.2. The existence of such parameterizations is well known [13, 24, 37], and in this case there exist entire families of parameterizations. This additional degree

of freedom can be exploited in the general energy (2.5) to optimize a secondary criterion while maintaining an exactly area-preserving parameterization.

2.2 Surface Matching

A correspondence between two surface patches, \mathcal{M}_A and \mathcal{M}_B , can be regarded as a non-rigid spatial deformation

$$\phi_{\mathcal{M}} : \mathcal{M}_A \rightarrow \mathbb{R}^3$$

such that corresponding regions of \mathcal{M}_A are mapped onto regions of \mathcal{M}_B . Instead of formulating these maps directly in \mathbb{R}^3 , we match geometric characteristics and user-defined features in the parameter domain. The main benefit of this approach is that it simplifies the problem of finding correspondences for surfaces embedded in \mathbb{R}^3 to a matching problem in two dimensions.

Our motivation comes from a variational approach for matching images through an energy relaxation over a set of non-rigid deformations in the plane [18, 20], where the optimal match is achieved by the mapping that minimizes a suitable energy. To ensure that the actual geometry of the surface patches is treated properly in our approach, the energy on the deformations from one parameter space to the other will measure:

- ◇ **(regularization energy)** smoothness of the deformation in terms of tangential distortion,
- ◇ **(bending energy)** bending of normals through the proper correspondence of curvature, and
- ◇ **(feature energy)** the proper correspondence of important surface and texture features.

Furthermore, it will consistently take into account the proper metrics on the parameter domains, which ensures that we are actually treating a deformation from one surface onto the other even though all computations are performed in 2D.

2.2.1 A Physical Interpretation of the Matching Deformation

Consider the first surface to be a thin shell which we press into a mould of the second surface (Figure 2.2). One can distinguish between stresses induced by stretching and compression, and stresses induced by bending that occurs in the surface as it is being pressed. Thus $\phi_{\mathcal{M}}$ can be regarded as the deformation of such a thin shell. We assume this deformation to be elastic. The regularization energy in (2.7) will measure the induced in-plane stresses, and the concrete energy density in (2.8) allows control over length and area-distortion in this surface-to-surface deformation. Since we are aiming for a proper correspondence of shape, we will incorporate the bending of normals in our energy with (2.9). Finally, the matching of feature sets in (2.10) will provide user-specified landmarks to guide the surface deformation. In what follows, we will develop the variational approach step-by-step.

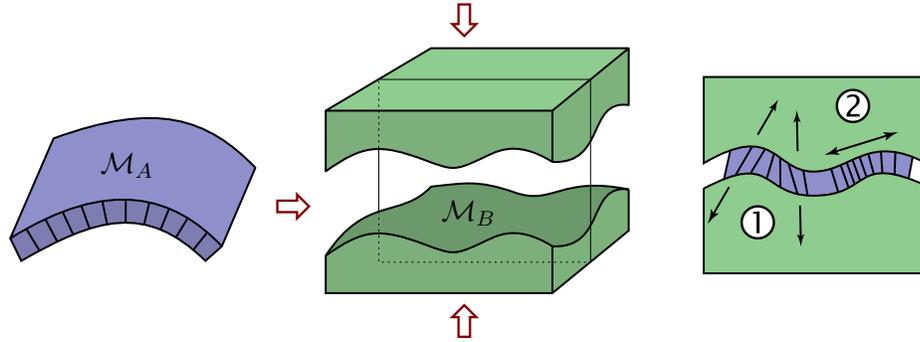


Figure 2.2: A physical interpretation of ϕ_M as pressing a thin shell \mathcal{M}_A into a mould of the surface \mathcal{M}_B being matched. The bending (1) and stretching (2) of the thin shell is measured in our matching energy, and minimized by the optimal match ϕ_M .

2.2.2 Measuring Distortion in a Deformation

The distortion in a parameterization was the subject of Section 2.1.2 in the previous chapter, where it was shown that the distortion from the surface patch \mathcal{M} onto the parameter domain ω under the inverse parameterization x^{-1} is measured by the inverse metric $g^{-1} \in \mathbb{R}^{2,2}$. Specifically, $\sqrt{\text{tr}(g^{-1})}$ measures the average *change of length* of tangent vectors under the mapping from the surface onto the parameter plane. Additionally, $\sqrt{\det(g^{-1})}$ measures the corresponding *change of area*. We will use these quantities in the following sections to account for the distortion of length and area on the surface as we formulate our matching energy in the parameter domain.

The above discussion now applies to the parameter maps x_A and x_B of the surfaces \mathcal{M}_A and \mathcal{M}_B . We suppose that these parameterizations are defined in an initial step and we assume that x_A and x_B as well as the corresponding parameter domains ω_A and ω_B are fixed from now on. Their metrics are denoted by g_A and g_B , respectively. We will now study the distortion which arises from a deformation of the first parameter domain onto the second parameter domain. First, let us consider deformations $\phi : \omega_A \rightarrow \omega_B$ which are one-to-one. This deformation between parameter domains induces a deformation between the surface patches $\phi_M : \mathcal{M}_A \rightarrow \mathcal{M}_B$ defined by

$$\phi_M := x_B \circ \phi \circ x_A^{-1}.$$

Let us emphasize that we do not actually expect a one-to-one correspondence between surface patches. Later we will relax this assumption and in particular allow for deformations ϕ with $\phi(\omega_A) \not\subset \omega_B$. The complete mapping is illustrated in Figure 2.3.

Now let us focus on the distortion from the surface \mathcal{M}_A onto the surface \mathcal{M}_B . In elasticity, the

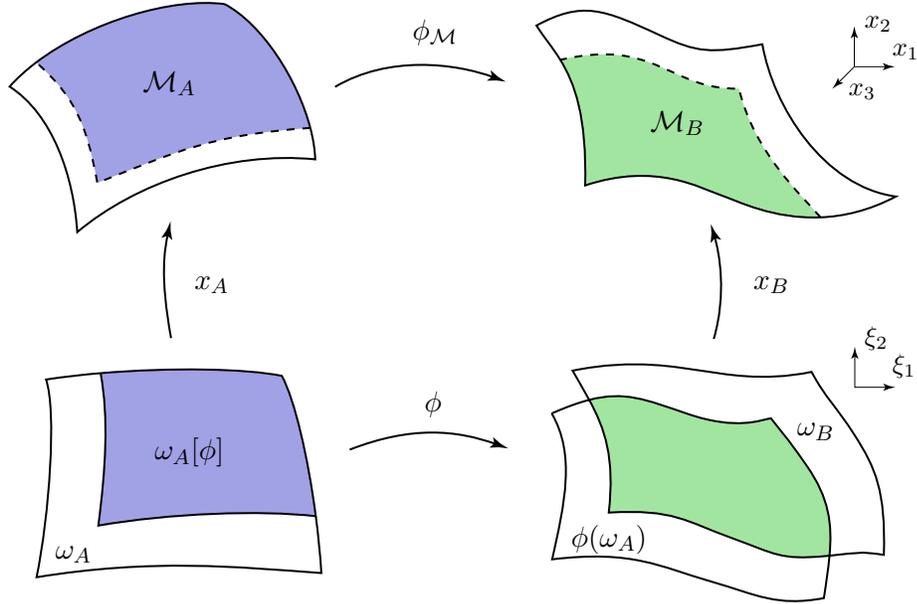


Figure 2.3: The matching function $\phi_{\mathcal{M}} := x_B \circ \phi \circ x_A^{-1}$ is a mapping between the corresponding shaded regions of the two surfaces. The partial correspondence is defined on $\omega_A[\phi] := \omega_A \cap \phi^{-1}(\omega_B)$.

distortion under an elastic deformation ϕ is measured by the Cauchy-Green strain tensor $D\phi^T D\phi$. We wish to adapt this definition to measure distortion between tangent vectors on the two surfaces, as we did with the metric g in Section 2.1.3. Therefore, we properly incorporate the metrics g_A and g_B at the deformed position and obtain the distortion tensor $\mathcal{G}[\phi] \in \mathbb{R}^{2,2}$ given by

$$\mathcal{G}[\phi] = g_A^{-1} D\phi^T (g_B \circ \phi) D\phi,$$

which acts on tangent vectors on the parameter domain ω_A , where products are denoted in matrix notation. Mathematically, this tensor is defined implicitly via the identity $(g_A \mathcal{G}[\phi] v) \cdot w = (g_B \circ \phi) D\phi v \cdot D\phi w$ for tangent vectors v, w on the surface \mathcal{M}_A and their images as tangent vectors $D\phi v, D\phi w$ on \mathcal{M}_B , where here we have identified tangent vectors on the surfaces with vectors in the parameter domains.

As in the parameterization case, one observes that $\sqrt{\text{tr}(\mathcal{G}[\phi])}$ measures the average *change of length* of tangent vectors from \mathcal{M}_A when being mapped to tangent vectors on \mathcal{M}_B and $\sqrt{\det(\mathcal{G}[\phi])}$ measures the *change of area* under the deformation $\phi_{\mathcal{M}}$. Thus $\text{tr}(\mathcal{G}[\phi])$ and $\det(\mathcal{G}[\phi])$ are natural variables for an energy density in a variational approach measuring the *regularity* of a surface de-

formation,

$$E_{reg}[\phi] = \int_{\omega_A} W(\text{tr}(\mathcal{G}[\phi]), \det(\mathcal{G}[\phi])) \sqrt{\det g_A} d\xi. \quad (2.7)$$

This simple class of energy functionals was previously derived in Chapter 2.1 from a set of natural axioms for measuring the distortion of a single parameterization. In particular, the following energy density

$$W(a, d) = \alpha_l a + \alpha_a \left(d + \left(\frac{\alpha_l}{\alpha_a} + 1 \right) d^{-1} \right) \quad (2.8)$$

accounts for length distortion with $a = a(A) = \text{tr}(\mathcal{G}[\phi])$, area expansion with $d = d(A) = \det(\mathcal{G}[\phi])$ and area compression with d^{-1} in a similar manner to (2.5). The parameters $\alpha_l, \alpha_a > 0$ are typically chosen by the user according to the relative importance of length and area distortion.

2.2.3 Measuring Bending in a Deformation

When we press a given surface \mathcal{M}_A into the thin mould of the surface \mathcal{M}_B , a major source of stress results from the bending of normals. We assume these stresses to be elastic as well and to depend on changes in normal variations under the deformation. Variations of normals are represented in the metric by the shape operator. We defer the derivation of the shape operators S_A and S_B of the surface patches \mathcal{M}_A and \mathcal{M}_B to Appendix B.1, where we end up with $\text{tr}(S_B \circ \phi) - \text{tr}(S_A)$ as a *measure for the bending of normals*. Since the trace of the shape operator is the mean curvature, we can instead aim to compare the mean curvature $h_B = \text{tr}(S_B)$ of the surface \mathcal{M}_B at the deformed position $\phi_{\mathcal{M}}(x)$ and the mean curvature $h_A = \text{tr}(S_A)$ of the surface \mathcal{M}_A . A similar observation was used in [21] to define a bending energy for discrete thin shells. This proposed simplification neglects any rotation of directions due to the deformation, *e.g.*, if the deformation aligns a curve with positive curvature on the first surface to a curve with negative curvature on the second surface and vice versa, an energy depending solely on $h_B \circ \phi - h_A$ does not recognize this mismatch. Nevertheless, in practice the bending energy

$$E_{bend}[\phi] = \int_{\omega_A} (h_B \circ \phi - h_A)^2 \sqrt{\det g_A} d\xi \quad (2.9)$$

turns out to be effective and sufficient. By minimizing this energy, we ensure that the deformation properly matches mean curvature on the surfaces.

2.2.4 Matching Features

Frequently, surfaces are characterized by similar geometric or texture features, which should be matched properly as well. Therefore we will incorporate a correspondence between one-dimensional feature sets in our variational approach to match characteristic lines drawn on the surface. In particular, we prefer feature lines to points for the flexibility afforded to the user, as well as to avoid the theoretical problems introduced by point constraints [9]. We will denote the feature sets by $\mathcal{F}_{\mathcal{M}_A} \subset \mathcal{M}_A$ and $\mathcal{F}_{\mathcal{M}_B} \subset \mathcal{M}_B$ on the respective surfaces. Furthermore, let $\mathcal{F}_A \subset \omega_A$ and $\mathcal{F}_B \subset \omega_B$ be the corresponding sets on the parameter domains. We are aiming for a proper match of these sets via the deformation, *i.e.*,

$$\phi_{\mathcal{M}}(\mathcal{F}_{\mathcal{M}_A}) = \mathcal{F}_{\mathcal{M}_B}$$

or in terms of differences, $\mathcal{F}_{\mathcal{M}_A} \setminus \phi_{\mathcal{M}}^{-1}(\mathcal{F}_{\mathcal{M}_B}) = \emptyset$ and $\mathcal{F}_{\mathcal{M}_B} \setminus \phi_{\mathcal{M}}(\mathcal{F}_{\mathcal{M}_A}) = \emptyset$. A rigorous way to reflect this in our variational approach is with a third energy contribution,

$$\begin{aligned} E_{\mathcal{F}}[\phi] &= \mathcal{H}^1(\mathcal{F}_{\mathcal{M}_A} \setminus \phi_{\mathcal{M}}^{-1}(\mathcal{F}_{\mathcal{M}_B})) + \\ &\quad \mathcal{H}^1(\mathcal{F}_{\mathcal{M}_B} \setminus \phi_{\mathcal{M}}(\mathcal{F}_{\mathcal{M}_A})) \end{aligned} \quad (2.10)$$

where $\mathcal{H}^1(\mathcal{A})$ is the one-dimensional Hausdorff measure of a set \mathcal{A} on the corresponding surface. Roughly speaking, this gives a symmetric measurement of the size of the mismatch of the features. This type of energy does not lend itself to a robust numerical minimization. Therefore, we will instead consider a suitable approximation of (2.10) that involves the *distance on the surface to the feature sets*, and define

$$\begin{aligned} \tilde{E}_{\mathcal{F}}^{\epsilon}[\phi] &= \int_{\omega_A} (\eta^{\epsilon} \circ d_A(\xi)) (\theta^{\epsilon} \circ d_B(\phi(\xi))) \sqrt{\det g_A} d\xi + \\ &\quad \int_{\omega_B} (\eta^{\epsilon} \circ d_B(\xi)) (\theta^{\epsilon} \circ d_A(\phi^{-1}(\xi))) \sqrt{\det g_B} d\xi \end{aligned} \quad (2.11)$$

where $d_A(\cdot) = \text{dist}_A(\cdot, \mathcal{A})$ and $d_B(\cdot) = \text{dist}_B(\cdot, \mathcal{A})$ are distance functions on the parameter domains ω_A and ω_B with respect to some set \mathcal{A} on the corresponding surface. Note that we measure distance either in the metric g_A on ω_A or in the metric g_B on ω_B . Additionally, we define the

localization functions

$$\eta^\epsilon(s) = \frac{1}{\epsilon} \max\left(1 - \frac{s}{\epsilon}, 0\right), \quad \theta^\epsilon(s) = \min\left(\frac{s^2}{\epsilon}, 1\right)$$

which act as cut-off functions. For Lipschitz continuous feature sets and bi-Lipschitz continuous deformations, we observe that $\tilde{E}_{\mathcal{F}}^\epsilon[\phi] \rightarrow E_{\mathcal{F}}[\phi]$ as $\epsilon \rightarrow 0$, which motivates our approximation. In view of the later discretization, we can reformulate the second term in (2.11) as

$$\int_{\omega_A} (\eta^\epsilon \circ d_B(\phi(\xi))) (\theta^\epsilon \circ d_A(\xi)) \sqrt{\det g_B(\phi(\xi))} \det D\phi \, d\xi.$$

2.2.5 Partial Correspondence

Usually, we cannot expect that $\phi_{\mathcal{M}}(\mathcal{M}_A) = \mathcal{M}_B$, particularly near the boundary where certain subregions of \mathcal{M}_A will have no corresponding counterpart on \mathcal{M}_B and vice versa. Therefore, we must allow for points on \mathcal{M}_B with no pre-image in \mathcal{M}_A under a matching deformation $\phi_{\mathcal{M}}$, and points on \mathcal{M}_A which are not correlated to points on \mathcal{M}_B via $\phi_{\mathcal{M}}$ (cf. Figure 2.3). Thus we must adapt the variational formulation accordingly. If $\phi(\omega_A) \neq \omega_B$, then $\phi_{\mathcal{M}}$ is now defined on $x_A(\omega_A[\phi])$ only, where

$$\omega_A[\phi] := \phi^{-1}(\phi(\omega_A) \cap \omega_B)$$

is the corresponding subset of the parameter domain ω_A . Furthermore, we define new energies (with modifications marked in red):

$$E_{bend}[\phi] = \int_{\omega_A[\phi]} (h_B \circ \phi - h_A)^2 \sqrt{\det g_A} \, d\xi, \quad (2.12)$$

$$E_{\mathcal{F}}[\phi] = \mathcal{H}^1(\omega_A[\phi] \cap \mathcal{F}_{\mathcal{M}_A} \setminus \phi_{\mathcal{M}}^{-1}(\mathcal{F}_{\mathcal{M}_B})) + \mathcal{H}^1(\mathcal{F}_{\mathcal{M}_B} \setminus \phi_{\mathcal{M}}(\omega_A[\phi] \cap \mathcal{F}_{\mathcal{M}_A})). \quad (2.13)$$

For an energy that controls tangential distortion, it is still helpful to control the regularity of the deformation outside the actual matching domain $\omega_A[\phi]$, where we would like to allow significantly larger deformations by using a “softer” elastic material. Hence we will suppose that g_B , which is initially only defined on ω_B , is extended to \mathbb{R}^2 and takes on values that are relatively small to allow for greater stretching.

In the minimization algorithm, we need descent directions which will involve derivatives of these energies with respect to the deformation ϕ . In taking these derivatives, integrals over the

variable boundary $\partial\omega_A[\phi]$ will appear. Since these are tedious to treat numerically, we will rely on another approximation for the sake of simplicity. Our strategy here is to change the domain of integration $\omega_A[\phi]$ to a superset ω which extends beyond the boundary $\partial\omega_A[\phi]$. Doing so means that special treatment of boundary integrals is no longer necessary, although we are now required to evaluate the integrands of the energies outside of ω_A , and similarly for deformed positions outside of ω_B . To achieve this, we will extend our surface quantities onto $\omega \setminus \omega_A$ and $\omega \setminus \omega_B$, respectively, by applying a harmonic extension with natural boundary conditions on $\partial\omega$ to g_A, g_B and h_A, h_B (e.g., we define h_A as the solution of Laplace's equation on $\omega \setminus \omega_A$ with vanishing flux on $\partial\omega$). Additionally, we introduce a regularized characteristic function

$$\chi_{\mathcal{A}}^\epsilon(\xi) = \max(1 - \epsilon^{-1} \text{dist}(\xi, \mathcal{A}), 0) \quad (2.14)$$

to cause the energy contributions to be ignored at some distance ϵ away from $\omega_A[\phi]$. Thus, instead of dealing with a deformation dependent-domain $\omega[\phi]$ in the definition of our different energy contributions, we always integrate over the *whole* image domain ω and insert the product of the two regularized characteristic functions

$$\chi^\epsilon(\xi) = \chi_{\omega_A}^\epsilon(\xi) \chi_{\omega_B}^\epsilon(\phi(\xi))$$

as an additional factor in the energy integrand. We apply this modification to the energy E_{bend} (2.9) and the already regularized energy $\tilde{E}_{\mathcal{F}}^\epsilon$ (2.11) and denote the resulting energies by

$$E_{bend}^\epsilon \quad \text{and} \quad E_{\mathcal{F}}^\epsilon, \quad (2.15)$$

respectively.

2.2.6 Definition of the Matching Energy

We are now ready to collect the different cost functionals and define the global matching energy. Depending on the user's preference, we introduce weights $\beta_{bend}, \beta_{reg}, \beta_{\mathcal{F}}$ for the approximate energies in (2.15) — we have found that $\beta_{bend} = 1, \beta_{reg} = 0.01, \beta_{\mathcal{F}} = 5$ work well — and define the global energy

$$\boxed{E^\epsilon[\phi] = \beta_{bend} E_{bend}^\epsilon[\phi] + \beta_{reg} E_{reg}^\epsilon[\phi] + \beta_{\mathcal{F}} E_{\mathcal{F}}^\epsilon[\phi]} \quad (2.16)$$

which measures the quality of a matching deformation ϕ on the domain ω . Finally, in the limit $\epsilon \rightarrow 0$ we obtain a weighted sum of (2.7), (2.12) and (2.13):

$$E[\phi] = \beta_{bend} E_{bend}[\phi] + \beta_{reg} E_{reg}[\phi] + \beta_{\mathcal{F}} E_{\mathcal{F}}[\phi]. \quad (2.17)$$

We call a matching deformation β -*optimal* if it minimizes the energy (2.17) for a specific set of parameters $\beta = (\beta_{bend}, \beta_{reg}, \beta_{\mathcal{F}})$. One can prove the *existence, global injectivity, and regularity of β -optimal matching deformations* [17]. Because of this proof, we can expect to obtain smooth deformations that are free of folds and singularities.

Chapter 3

Algorithms

The variational framework described in the previous chapter provides us with a solid foundation on which to design numerical methods for generating parameterizations. In the discrete setting, we assume that the surface patches are given as triangle meshes. In this chapter, we provide the algorithms for generating optimal, discrete surface parameterizations and surface matches, respectively. The general procedure is the same for both of these methods. We begin with an initial mapping of the surface to a triangulated parameter domain. A straightforward finite element discretization follows, which describes the function space of elastic deformations over this domain. We then employ standard methods from numerical optimization to generate an optimal deformation that minimizes a variational energy. The energy, which is defined at the discrete level, properly accounts for the geometry of the surface through the discrete metric of the initial parameterization. The algorithms are guaranteed to generate an optimal mapping that retains the analytic properties which we derived earlier, including smoothness and local bijectivity. We will now turn to the description and analysis of the surface parameterization algorithm, followed by the surface matching algorithm in the second half of this chapter.

3.1 The Parameterization Algorithm

We shall now describe our algorithm to construct an optimal parameterization for a discrete surface patch. We assume that the surface patch \mathcal{M} is provided as a manifold triangle mesh with boundary. A discrete parameterization $x : \omega \rightarrow \mathcal{M}$ is a piecewise-affine mapping on the triangles in the parameter domain $\omega \subset \mathbb{R}^2$ to the corresponding triangles in \mathbb{R}^3 . The algorithm generates an optimal, discrete deformation $\tilde{\phi} : \omega \rightarrow \mathbb{R}^2$ which we apply directly to the planar triangulation ω to obtain the optimal parameter domain $\omega[\tilde{\phi}] := \tilde{\phi}(\omega)$ (cf. Figure 2.1). The final parameterization is fully defined by the coordinates assigned to the vertices of the triangles in $\omega[\tilde{\phi}]$ (e.g., as “texture coordinates” for texture mapping, Figure 3.1).

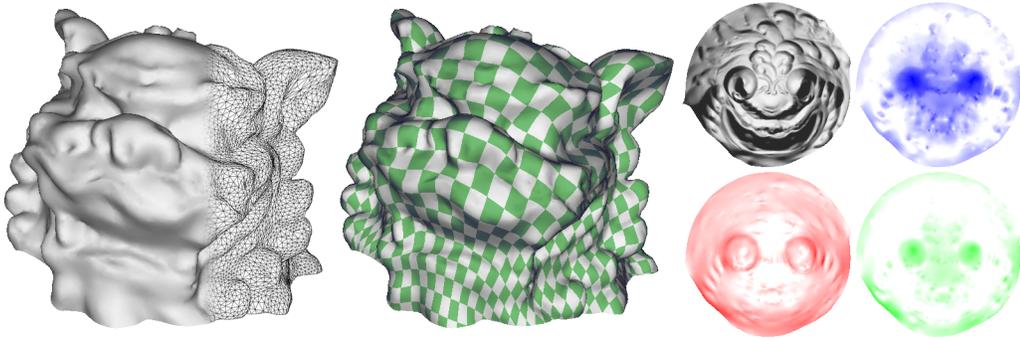


Figure 3.1: For a surface mesh (left, 36k triangles), a parameterization is computed, depicted through a texture map (middle). The parameter domain (right) is optimized through the minimization of energy functions which control length, area and angle distortion (shown clockwise from top right; dark regions indicate higher energy).

The user controls the parameterization through the choice of parameters $\alpha = (\alpha_l, \alpha_a, \alpha_c)$ that balance the trade-off between length, area and angle distortion, respectively. For example, specifying a larger value for α_a will result in a parameterization that is biased toward low area distortion at the expense of greater length and angle distortion, *etc.* The parameters must satisfy the conditions stated in (2.6), namely

$$\alpha_l > 0, \quad \alpha_a > 0, \quad \alpha_c \geq 0.$$

In practice, we normalize the parameters using $\alpha_l + \alpha_a + \alpha_c = 1$ for the sake of better numerics.

We will now turn to the main steps in the algorithm:

1. Construct an initial parameterization of the surface.
2. Optimize the parameterization through a discrete deformation of the parameter domain that minimizes the energy $E[\cdot]$ defined in (2.5).

3.1.1 Initial Parameterization

Our algorithm requires that an initial parameterization be provided for the subsequent optimization of the parameter domain. With this optimization step in mind, we desire an initial parameterization that is close to optimal to make the energy minimization more efficient. It follows that this parameterization should already have low distortion. Moreover we advocate choosing an initial parameterization with natural boundary conditions over a constrained boundary, *e.g.*, a convex polygon in the plane. In the latter case one often encounters extreme shearing and compression in the boundary elements which is difficult to correct through numerical optimization.

For this purpose we use the natural conformal map [15] in our examples (also known as the least-squares conformal map [29]). This parameterization has low angle distortion and natural boundary conditions, and is cheap to compute since it only requires the solution of a sparse linear system. Other free-boundary parameterization algorithms can be used (*e.g.*, [26, 46]) provided that the initial parameterization is admissible (cf. Section 2.1.2).

Recall from Section 2.1.5 that the energy (2.5) is a function of the principal invariants of the inverse metric g^{-1} . On inspection only the conformal (third) term in the energy is unaffected by uniformly scaling the surface patch \mathcal{M} or the parameter domain ω . In practice we would like the parameterization energy to be scale-invariant. In many applications the scale of the ambient space in \mathbb{R}^3 and the parameter space in \mathbb{R}^2 are chosen independently and therefore a parameterization energy that identifies these scales does not fit within the application framework. Furthermore, it is unreasonable to expect the numerical optimization method to make up the difference in scale during the energy minimization, as this requires extra computation and can become a source of instability due to the large magnitude of the gradients for the length and area terms. One solution is to reformulate the energy so that these terms are scale-invariant. However a simpler solution is to uniformly scale the parameter domain by a factor σ prior to the optimization step. Following [44], we let σ be the ratio of the surface area versus the area of the domain triangulation, so that $\sigma = 1$

for any isometric parameterization of a developable surface. When the optimal parameterization is constructed, the parameter domain is returned to its original scale using the inverse scale factor σ^{-1} .

3.1.2 Finite Element Discretization

Suppose that ω is an admissible triangle mesh in the parameter space \mathbb{R}^2 , and let T denote the triangles of ω . A discrete deformation $\phi : \omega \rightarrow \mathbb{R}^2$ is a vector-valued function, whose components are piecewise-affine, continuous functions. For these deformations we have to evaluate the discrete counterpart of energy functional (2.5),

$$E[\phi] = \sum_{T \in \omega} W(\text{tr}(\mathbf{C}[\phi]), \det(\mathbf{C}[\phi])) \sqrt{\det g} |T|$$

with the area of a triangle in the plane denoted by $|T|$. On each triangle $T \in \omega$, a constant Cauchy-Green strain tensor is defined by $\mathbf{C}[\phi] := D\phi g^{-1} D\phi^T$ where the derivative $D\phi$ and the discrete metric g are constant matrices in $\mathbb{R}^{2,2}$.

Besides evaluating the discrete energy for a given deformation ϕ , we have to compute the gradient $\nabla E[\phi]$ and (optionally) the Hessian $\nabla^2 E[\phi]$ in the actual minimization algorithm in Section 3.1.3, which requires the differentiation of the discrete energy with respect to the discrete deformation. All the necessary expressions to assemble these matrices are provided in Appendix B.2.

A necessary condition for a discrete deformation ϕ to minimize the energy is $\nabla E[\phi] = 0$. Due to our assumption of frame indifference (2.2) the energy is invariant under Euclidean motion. To remove these degrees of freedom we add two constraints on the zero moment $M_0(\phi)$ to eliminate translations, and one on the angular momentum $M_1(\phi)$ to cancel rotations:

$$\begin{aligned} M_0(\phi) &:= \int_{\omega} \phi(\xi) d\xi = 0 \\ M_1(\phi) &:= \int_{\omega} [\phi(\xi)_2 \xi_1 - \phi(\xi)_1 \xi_2] d\xi = 0 \end{aligned}$$

To provide the necessary normalization with respect to rigid body motions, we employ Lagrange multipliers for the moment constraints. Effectively, the discrete energy $E[\phi]$ is replaced by a modified energy

$$\bar{E}[\phi, \lambda] = E[\phi] + \begin{pmatrix} \lambda_1 \\ \lambda_2 \end{pmatrix} \cdot M_0(\phi) + \lambda_3 M_1(\phi).$$

Instead of solving $\nabla E[\phi] = 0$ for ϕ one now solves the equation $\nabla \bar{E}[\phi] = 0$ for the deformation ϕ and the Lagrange multipliers $\lambda_1, \lambda_2, \lambda_3$. To find the minimum of this energy with a Newton method one also needs the Hessian. Based on $\nabla^2 E[\phi]$ we obtain

$$\nabla^2 \bar{E}[\phi, \lambda] = \begin{pmatrix} \nabla^2 \bar{E}[\phi] & M \\ M^T & 0 \end{pmatrix}$$

using

$$M = \begin{pmatrix} \mathbf{m} & 0 & -\mathbf{m} \\ 0 & \mathbf{m} & \mathbf{m} \end{pmatrix},$$

where \mathbf{m} is the vector of masses of the nodal basis functions, *i.e.*, $\mathbf{m}^k := \sum_{\mathcal{N}^k} \frac{1}{3} |T|$ with \mathcal{N}^k denoting the set of triangles incident to vertex $k = 1, \dots, n$. Note that while this new Hessian $\nabla^2 \bar{E}[\phi, \lambda]$ is not positive definite, it does not pose a problem for our algorithm as we explain below.

3.1.3 Numerical Optimization

We use an iterative approach to find a sequence of discrete deformations $\{\phi_i\}_{i=0 \dots N}$ such that ϕ_N minimizes the discrete energy $E[\cdot]$, *i.e.*, $E[\phi_0] > E[\phi_1] > \dots > E[\phi_N]$ and $\nabla E[\phi_N] = 0$. The optimal deformation $\tilde{\phi} = \phi_N$ thus generates the final, optimal parameterization $x[\tilde{\phi}] := x \circ \tilde{\phi}^{-1}$.

Because our energy is polyconvex, it has multiple local minima and therefore it is non-trivial to guarantee that one will find a global minimum. Our strategy is to find a local energy minimum starting from a reasonable initial guess, namely, the initial parameterization in Section 3.1.1. It follows that the minimizing sequence of deformations begins with ϕ_0 being the identity function.

```

// Find an optimal deformation  $\tilde{\phi}$  that minimizes  $E[\cdot]$ 
 $\phi_0 \leftarrow \mathbf{0}$ ,  $i \leftarrow 0$ 
repeat
  // Choose a search direction and step size
  compute the search direction  $\psi_i$ 
  calculate the step size  $\tau_i$ 
  if  $\|\tau_i \psi_i\|_2 < \varepsilon$  then  $\tilde{\phi} \leftarrow \phi_i$  and stop

  // Update the deformation
   $\phi_{i+1} \leftarrow \phi_i + \tau_i \psi_i$ 
   $i \leftarrow i + 1$ 

```

Figure 3.2: Pseudo-code for the optimization algorithm.

We use a simple line search method to find $\tilde{\phi}$ (Figure 3.2). Each iteration of the line search computes a search direction ψ_i and then decides how far to step in this direction. The iteration is given by

$$\phi_{i+1} \leftarrow \phi_i + \tau_i \psi_i$$

Many classical numerical methods can be used to compute ψ_i (Figure 3.2, **search direction**). There is a typical trade-off between performance, in terms of the number of iterations to converge to the minimum, and the complexity of algorithm, which is related to what local model of the energy landscape is used in selecting the search direction. For example, the simplest gradient descent method relies only on the local gradient, *i.e.*, $\psi_i = \nabla E[\phi_i]$, but it has poor convergence properties. We chose to use a Newton method in our implementation because it guarantees a quadratic rate of convergence close to an optimal solution (Chapter 3, [38]). It computes ψ_i by solving a sparse linear system,

$$\nabla^2 \bar{E}[\phi_i, \lambda] \begin{pmatrix} \psi_i \\ \mu \end{pmatrix} = \begin{pmatrix} \nabla E[\phi_i, \lambda] \\ 0 \end{pmatrix} \quad (3.1)$$

where $\mu \in \mathbb{R}^3$ appears due to the Lagrange multipliers in $\bar{E}[\cdot]$. For the solution of the linear saddle-point problem (3.1), one could apply Uzawa's algorithm [8]. We simply use a preconditioned conjugate gradient algorithm. We have never observed any difficulty in practice even though theoretically the matrix $\nabla^2 \bar{E}[\phi_i, \lambda]$ is not positive definite. As a preconditioner, we apply a diagonal preconditioning to $\nabla^2 \bar{E}$.

The next step is to determine the step size τ_i for the iteration (Figure 3.2, **step size**). We require τ_i to be chosen such that the energy decreases at each iteration. We accomplish this with a simple step size control algorithm that uses the Armijo condition [38] to guarantee sufficient decrease of the energy functional (Figure 3.3). The parameters control how aggressively the algorithm is in taking large steps — we find that $\tau_i = 1$, $\rho = 0.8$, $c = 0.5$ work well in practice.

```

// Choose  $\tau_i > 0$ ,  $\rho, c \in (0, 1)$ 
repeat until  $E[\phi_i + \tau_i \psi_i] \leq E[\phi_i] + c\tau_i \nabla E[\phi_i]^T \psi_i$ 
   $\tau_i \leftarrow \rho \tau_i$ 
```

Figure 3.3: Pseudo-code for the step size control algorithm.

Beside asking for sufficient decrease of the energy, one must be concerned with taking *legal* steps. Our concern lies in the fact that an illegal step $\phi_i + \tau_i \psi_i$ that generates an *inadmissible* deformation could actually lower the energy. This problem can be seen directly from our energy functional (2.5). If a triangle element T reverses its orientation due to a fold in the deformation, one can experience a negative energy value. One can address this issue by defining $E[\phi]$ to be unbounded for $\det \phi < 0$, which is a natural extension of $E[\cdot]$ since $E[\phi] \rightarrow \infty$ as $\det \phi \rightarrow 0$. This is equivalent to simply disallowing a step for which $\det(\phi_i + \tau_i \psi_i) \leq 0$ on any triangle T .

We stop the line search method when the update is sufficiently small in the L^2 norm, controlled by the threshold parameter $\varepsilon > 0$ (Figure 3.2, **stop**). If we have reached a local minimum, *i.e.*, $\nabla E[\phi_i] \approx 0$, then we have achieved the optimal deformation $\tilde{\phi} = \phi_i$.

The line search will also terminate if no further progress can be made to reduce the energy. This “locking” can occur when the line search gets lost in the energy landscape because of a poor choice for the initial parameterization. In this case, a simple continuation method on the energy parameters α can be used to set up a series of easier optimization problems,

$$H(\phi, \theta) = E[\phi]|_{\alpha=(1-\theta)\alpha_0+\theta\alpha_1}$$

where the blend parameter $\theta \in [0, 1]$ is chosen from a monotonically increasing sequence $\{\theta_j\}_{j=0\dots J}$ that transforms the energy landscape from $E[\cdot]|_{\alpha_0}$ to the desired energy $E[\cdot]|_{\alpha_1}$. The motivation comes from choosing α_0 so that the initial parameterization is close to an optimal solution of $E[\cdot]|_{\alpha_0} = H(\cdot, \theta_0)$. The minimizer of $H(\cdot, \theta_j)$ is used as the initial parameterization for the next problem in the sequence, $H(\cdot, \theta_{j+1})$. Thus by completing this sequence one obtains the optimal deformation $\tilde{\phi}$ as the minimizer of $E[\cdot]|_{\alpha_1} = H(\cdot, \theta_J)$. Generally speaking, naïve continuation methods of this form can fail in practice [38]. However we have found this approach to be entirely sufficient for transforming solutions of one energy to another.

3.2 The Matching Algorithm

In Chapter 2.2 we developed a variational framework for matching surface patches without regard to a particular discretization. Now we will describe our method for constructing a match based on a straightforward discretization using finite elements. We assume that the surface patches \mathcal{M}_A and \mathcal{M}_B are given as triangle meshes. In the initial step, we generate parameterizations which define triangulated parameter domains ω_A and ω_B using the approach described in Section 3.1. Because of the difficult algorithmic details, we do not wish to deal with effectively overlaying two triangulations during the numerical solver stage. Consequently we discretize the domain onto a regular grid (“image”) and evaluate the associated surface quantities needed in the energies at each pixel (in effect we may think of this as a geometry image [22]).

This setup has two principal advantages: (1) the resolution of the original meshes is decoupled from the resolution used in the image domain and (2) multiscale algorithms are far simpler to implement in the regularly sampled image grid than over arbitrary triangle meshes (even if flat). In particular, we can use higher sampling rates in the image domain to alleviate aliasing problems. Additionally, the image pyramids used by a hierarchical solver have far more efficient memory access patterns on modern processors than one achieves with arbitrary meshes.

We now turn to the basic components of the implementation:

1. Construct parameterizations for the surface patches.
2. (optional) Select matching features on the surfaces with separate texture maps.
3. Evaluate the metric and mean curvature by scan converting the surface triangulation in the parameter domain.
4. Apply a finite element discretization and optimize the matching deformation using a multiscale approach to minimize the energy $E^\epsilon[\cdot]$ defined in (2.16).

3.2.1 Surface Patch Parameterization

We are interested in low distortion parameterizations to ensure an adequate sampling and to keep the energy landscape as nice as possible. Recall that the metric of the parameterization enters into our energy formulation. Therefore, parameterizations with unnecessarily large gradients only serve to make the energy minimization problem harder. We address these issues by applying the algorithm

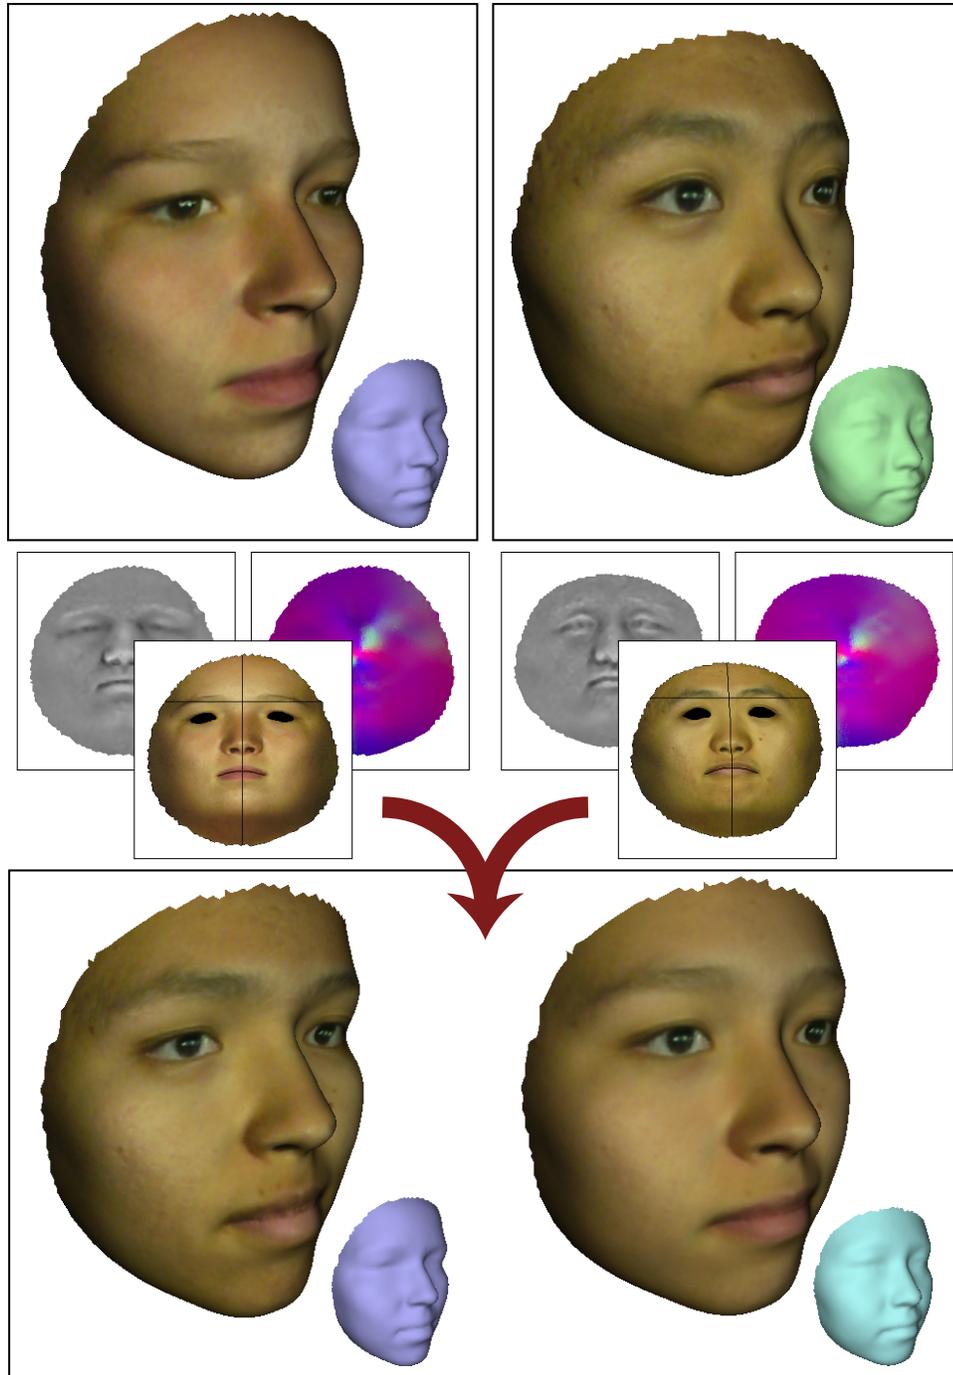


Figure 3.4: A brief overview of our surface matching process: The two surfaces to be matched are parameterized (top); we generate images of mean curvature, the metric and user-defined feature sets on a 256×256 grid (middle); an optimal matching deformation is then found through a global energy relaxation in the parameter domain using fast multiscale algorithms. The match is properly defined in 3D by incorporating the metric in the matching energy. A finely detailed texture is transferred to the first surface, and a 50% morph is created (bottom). The match took under 3 min.

in Section 3.1 to parameterize the surfaces. In practice, we find that $\alpha_l = \alpha_a = \alpha_c = \frac{1}{3}$ works nicely for this purpose.

Once both surfaces have been parameterized, we perform a rough alignment by applying a normalizing transformation such that the parameter domains are subsets of the domain $\omega := [0, 1]^2$. Due to the feature energy contribution from (2.11) and the hierarchical nature of the minimization algorithm this rough alignment is entirely sufficient.

3.2.2 Feature Set Construction

The user can control the match by identifying sets of similar features on the surface patches. The cost functional $E_{\mathcal{F}}^c[\cdot]$ from (2.11) helps to guide the energy minimization to match the corresponding features. Marking the desired feature set \mathcal{F}_A is most easily accomplished in the image domains. Figure 3.5 gives examples of these, showing the texture images mapped onto the surfaces. The actual feature sets are the boundaries of the (pixel) regions drawn by the user on the texture image. Note that when we match features, we do *not* constrain points since this would break the regularity of the deformation. Instead, we match *feature curves* which permits sliding of the deformation along the curve.

There are many ways in which features can be used to control a match:

- ◇ Coarse control of the match is achieved by roughly selecting corresponding geometric features and gradually decreasing $\beta_{\mathcal{F}}$ to zero as the multiscale method goes to finer resolutions (Figure 3.5a).
- ◇ Precise control over matching texture features (*e.g.*, on the face, *etc.*) is achieved by selecting the corresponding pixels in the feature set image (Figure 3.5b).
- ◇ Lines of symmetry drawn as features allow deformations tangential to the feature boundary, but prevent deformations that are transverse to it (Figure 3.5c).

In general, features tend to localize the matching deformation, *i.e.*, the feature boundaries partition the domain into deformable regions and minimize the deformation between these regions. This is useful when the surface is composed of regions with different elastic properties and prevents excessive “sliding” of the surface (*e.g.*, the highly deformable face versus the forehead, Figure 3.5d). Sometimes it is necessary to relax the regularization energy to allow very large deformations in a certain region (*e.g.*, opening the mouth, Figure 3.5e). We achieve this by decreasing β_{reg} in the selected region.

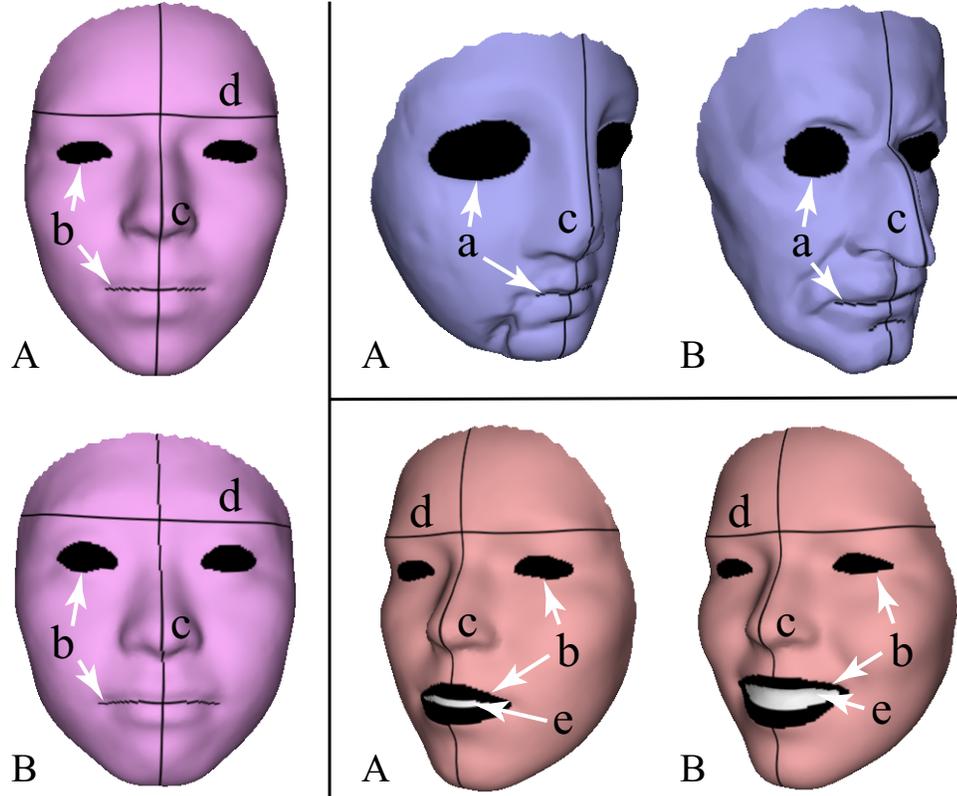


Figure 3.5: Examples of user-defined feature sets: (a) coarse registration of geometric features; (b) aligning texture features; (c) lines of symmetry; (d) preventing smooth, rigid regions from sliding; and (e) increasing the elasticity of highly deformable regions.

The distance maps used in the definition of the feature energy are discretized by an upwind scheme for the corresponding Eikonal equations [40] (Figure 3.6c). We used the particular upwind finite element algorithm of Bornemann [6] since it fits well with our overall finite element framework. Multiple sets of overlapping features are accounted for by taking the distance to the nearest feature to create a single distance map.

3.2.3 Evaluation of Surface Properties

The matching energy needs to evaluate surface quantities such as the mean curvatures h_A, h_B , metric tensors g_A, g_B , distance to the feature sets $\mathcal{F}_A, \mathcal{F}_B$, and signed distance to the domain boundaries $\partial\omega_A, \partial\omega_B$ in the case of (2.14). Since all of these are constants in the energy they only need to be discretized into appropriate textures once in the beginning. We achieve this through rendering the flattened triangulations with appropriate values at the vertices interpolated across triangles (Figure 3.6).

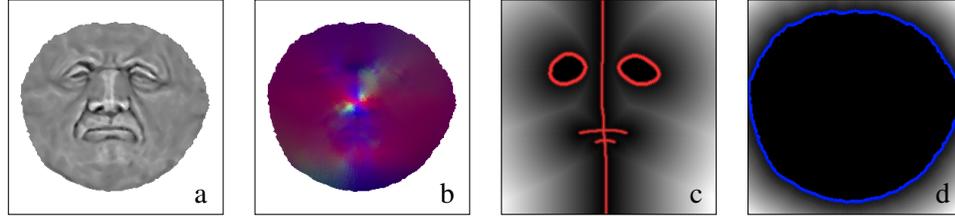


Figure 3.6: Surface properties are evaluated once and rasterized into images for (a) mean curvature h_A ; (b) the metric tensor g_A , with components shown as rgb values; (c) the distance map for the feature set \mathcal{F}_A , shown in red; and (d) the distance map for the characteristic function χ_A^ϵ , with the domain boundary $\partial\omega_A$ in blue.

For h_A we use the magnitude of the mean curvature normal [16, 34]. The sign is chosen according to the surface normal, which we take to be the area-weighted sum of triangle normals around a vertex. Other measurements for the mean curvature would work equally well (see [34] for a survey). Since g_A is symmetric and constant over each triangle element, we can evaluate its three unique components as a function of the triangle vertices. The calculation of the Jacobian of the parameterization over a triangle is well documented in the parameterization literature (e.g., see [44]). The distance map for \mathcal{F}_A is described above and illustrated in Figure 3.6c. To compute the distance map for the characteristic function, we rasterize the domain of \mathcal{M}_A and then generate its signed distance field in a similar manner to \mathcal{F}_A (Figure 3.6d).

3.2.4 Multiscale Finite Element Method

The total energy $E[\cdot]$ is highly non-linear. In particular the bending energy causes many local minima in the energy landscape over the space of deformations. We take a multiscale approach, solving a sequence of matching problems from coarse to fine scales. This type of method is frequently applied and well understood in image processing [2], allowing for a robust and efficient global minimization on complicated energy landscapes (Figure 3.7).

To begin, let us define a sequence of energies $(E^{\sigma_k})_{k=0,\dots,m}$ corresponding to scale parameters $\sigma_0 > \sigma_1 > \dots > \sigma_m = 0$, which act as filter parameters and range from coarse to fine scale. In essence the energy landscape is smoothed, enabling “basin catching” at coarse levels to provide good starting guesses for subsequently finer levels. Note that it is not necessary to compute the exact minimizer on every coarse scale. Instead we apply a non-linear conjugate gradient method with respect to the standard L^2 metric on the space of deformations, and use the Armijo condition for step size control [38] (Figure 3.7, **minimize**). We stop iterating as soon as the update is sufficiently small in the L^2 norm. In practice this proves to be a good heuristic to ensure that at the time we stop

```

// Build the image pyramid
for  $k = m$  downto 0 do
    filter images at level  $k$  using  $\sigma_k$ 

// Optimize  $\phi^k$  from coarse to fine scales
 $\phi^0 \leftarrow \mathbf{0}$ 
for  $k = 0$  to  $m$  do
    minimize  $E^{\sigma_k}$  starting with  $\phi^k$ 
    if  $k < m$  then prolongate  $\phi^k$  to  $\phi^{k+1}$ 

```

Figure 3.7: Pseudo-code for the multiscale algorithm.

on level k with a deformation ϕ^k , this deformation is already in the contraction region of the global minimum on the next finer scale $k + 1$ (Figure 3.7, **prolongate**). The prolongation of ϕ^k to ϕ^{k+1} is performed using bilinear interpolation.

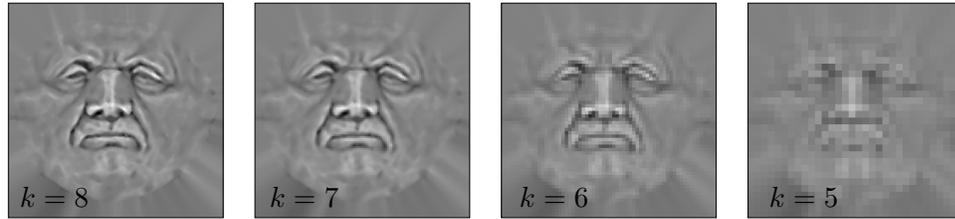


Figure 3.8: The mean curvature function h_A is extended to the full image domain ω and successively restricted to coarser grids $k = 8, \dots, 0$ from the multigrid pyramid.

It remains for us to define the scale of energies. First, we replace the functions on the surfaces as they appear in the different energy functionals by pre-filtered, smoothed representations (Figure 3.7, **filter**). A Gaussian filter of width σ_k is used to define the smoothing on the surfaces \mathcal{M}_A and \mathcal{M}_B . Exploiting the connection between Gaussian filtering and the fundamental solution of the heat equation, we replace the mean curvatures h_A and h_B (appearing in the bending energy) by pre-filtered mean curvature functions $h_A^{\sigma_k}$ and $h_B^{\sigma_k}$. This amounts to applying the appropriate filter kernels to the corresponding h_A and h_B images. Figure 3.8 shows images representing a scale of filtered mean curvature functions $h_A^{\sigma_k}$ on the parameter domain of a surface \mathcal{M}_A . Similarly, we filter the metric tensors g_A and g_B component-wise.

The regularization parameter ϵ in the definition of the energies also depends on the sequence of scale parameters and is set to $\epsilon(\sigma_k) = 2\sigma_k$. For the matching problems considered in this paper, we start with $\sigma_0 = 1$, and define $\sigma_k = \frac{1}{2}\sigma_{k-1}$ for $k = 1, \dots, m$. In our examples the parameter domains are discretized with a 256×256 grid and have a maximal number of $m = 8$ scales.

Chapter 4

Analysis

This chapter investigates the performance and robustness of the algorithms. We will demonstrate through a series of experiments the characteristic behaviour of these algorithms in practical test scenarios. The tests focus on parameterizations of physically-acquired triangle mesh data. Physical surfaces undeniably provide the largest variety of detail across many scales, so much so that modern acquisition methods often have difficulty producing data that is suitable for use without further processing. Establishing parameterizations for this data is one important form of pre-processing. Our challenge lies in the wide variety of shapes and the varying quality of the triangle meshes, which tax the performance and robustness of our numerical methods in different ways. We show evidence of efficiency and stability through our implementation of parameterization.

4.1 Surface Parameterization

The following provides demonstrations of the parameterization algorithm in Section 3.1. These feature typical but nonetheless challenging examples of parameterizations aimed at measuring the performance of our method. Fundamentally, we expect that the parameterizations generated for discrete surfaces uphold the principles that we established in our variational model for continuous surfaces (cf. Section 2.1). The central tenet is the existence of a continuous and locally bijective parameterization that is optimal for a given set of parameters $\alpha = (\alpha_l, \alpha_a, \alpha_c)$, which is summarized in Theorem 2.2. Moreover, the parameterization should be robust subject to different triangulations of the surface geometry, due to the finite element discretization of the parameterization energy in Section 3.1.2. Finally, we expect the distortion in the parameterizations to be distributed in relative accordance with the user’s choice of parameters α . The following examples show that all of these requirements are satisfied in our implementation.

4.1.1 Convergence of the Optimization Algorithm

For global energy minimization methods one is typically concerned with the rate of convergence to an optimal solution (if convergence occurs at all) and the number of iterations spent in the minimization algorithm. Figure 4.1 plots the energy decay for a series of examples, in which a quasi-conformal initial parameterization due to [15, 29] is optimized based on our own nearly conformal energy ($\alpha_c \gg \alpha_l, \alpha_a$). The vertical axis measures the L^2 norm of the energy gradient on a logarithmic scale, *i.e.*, $\log_{10}(\|\nabla E[\cdot]\|_2)$. The trend of these graphs shows that when the deformation ϕ_i is sufficiently close to the optimal solution, the Newton method achieves its optimal quadratic rate of convergence. In all of these examples the maximum step size ($\beta = 1$) was chosen at each iteration which further supports the choice of a Newton method for our problem.

\mathcal{M}	V	F	$\alpha = (\alpha_l, \alpha_a, \alpha_c)$	Iterations	Time
igea	7427	14611	(0.5, 0.5, 99)	4	0m59s
maxplanck	8991	17755	(1, 1, 98)	7	2m27s
mannequin	2732	5420	(0.5, 0.5, 99)	12	0m16s
sphere	4159	8184	(0.5, 0.5, 99)	5	0m17s
stampedpart	476	864	(1, 1, 98)	4	515ms
	1815	3456	(1, 1, 98)	4	0m03s
	7085	13824	(1, 1, 98)	5	1m03s

Table 4.1: Performance figures for examples of nearly conformal parameterizations (V, F denote # vertices and # faces, respectively, for triangle mesh \mathcal{M}).

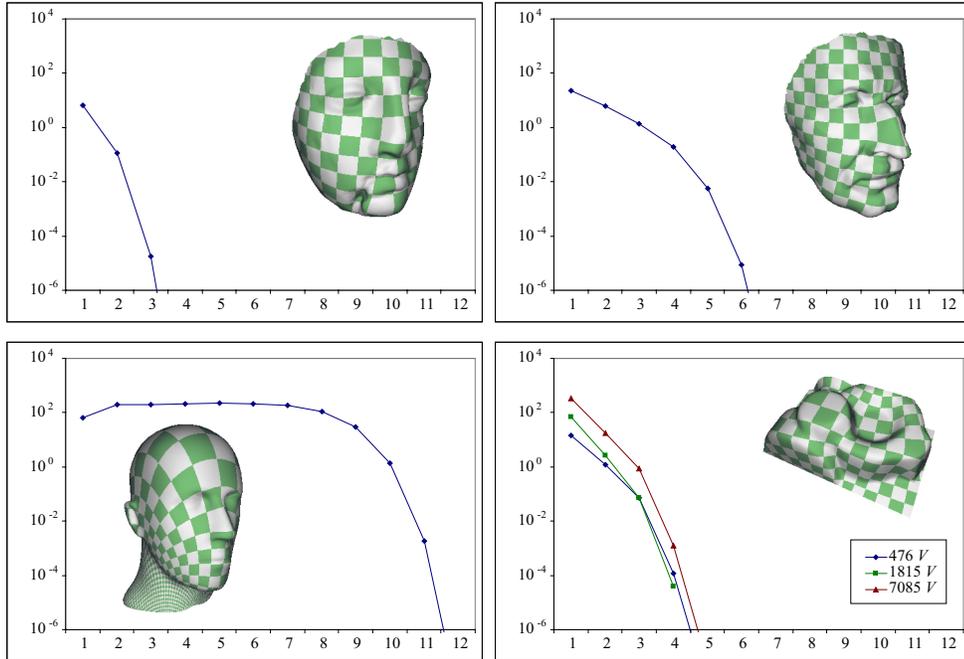


Figure 4.1: Plots of the energy decay for the sequence of deformations computed in the parameterization energy minimization algorithm. A quadratic rate of convergence is experienced near the optimal solution.

We can take a closer look at this process by inspecting the intermediate parameterizations that are generated during the minimization. Due to the invariants of our algorithm, these parameterizations are guaranteed to be admissible. Figure 4.2 depicts the intermediate deformations acting on a rectilinear grid in the initial parameter domain ω . Notice that much of the optimization is involved in improving the area distortion of the initial parameterizations. Since solutions of the natural conformal parameterization do not have direct control over the area factor, this result is not surprising. The final, optimal deformation is a *Möbius transformation* since the conformality of the parameterization is preserved under this mapping.

The overall performance of our implementation is summarized in Table 4.1. Performance figures were recorded on a 2 GHz Pentium IV processor. For each example, an initial quasi-conformal parameterization is optimized according to the energy defined by α (shown as percentile values of 1). In all of our tests, the global energy minimization algorithm converges in relatively few iterations. More iterations are typically required when the initial parameterization is significantly more distorted than the optimal parameterization.

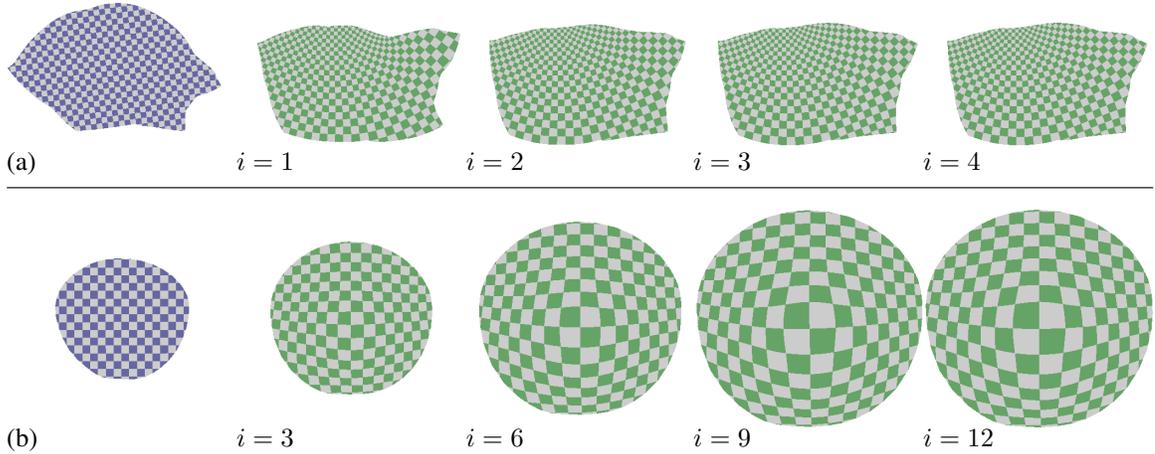


Figure 4.2: A sequence of deformations is generated by the global energy minimization algorithm, shown here acting on a grid in the parameter domain for (a) the stamped mechanical part and (b) the mannequin.

4.1.2 Convergence of the Continuation Method

The energy landscape that expresses the relationship between the initial parameterization and the optimal parameterization can be very complex: since the parameterization energy is polyconvex, it contains many local minima. Global optimization methods perform quite well when the initial parameterization is sufficiently close to the optimal solution, but they can quickly break down if a global minimum is nowhere in sight.

Continuation methods provide a strategy to gradually change the energy landscape allowing one to transform the optimal parameterization for one set of parameters α_0 to the solution for a second parameter set α_1 (cf. Section 3.1.3). Our strategy is quite naïve: after finding the optimal parameterization for the energy corresponding to parameters α_0 , we use this solution as the initial parameterization for the optimization of the energy defined by α_1 . Table 4.2 lists the number of iterations and the execution time for the optimization of $E[\phi]_{\alpha_1}$, starting from a parameterization $\tilde{\phi}$ that minimizes $E[\cdot]_{\alpha_0}$ (cf. Table 4.1). Although our continuation method is simple, it proves to be entirely effective. One can often take large steps between α_0 and α_1 in the continuation method when the initial parameterization minimizes $E[\cdot]_{\alpha_0}$. In practice, smaller steps must be taken when a large deformation is needed to obtain an optimal parameterization. This is most noticeable when one of the three energy terms $(\alpha_l, \alpha_a, \alpha_c)$ dominates the others (cf. the mannequin, Table 4.2). This issue is exaggerated when one desires an area-preserving parameterization, as these solutions often come at the price of large shear deformations that hinder numerical performance (cf. Figure 4.6-5).

\mathcal{M}	α_0	α_1	Iterations	Time
igea	(0.5, 0.5, 99)	(33.3, 33.3, 33.3)	6	0m54s
maxplanck	(1, 1, 98)	(33.3, 33.3, 33.3)	16	3m56s
sphere	(0.5, 0.5, 99)	(99, 0.5, 0.5)	7	0m25s
mannequin	(0.5, 0.5, 99)	(3, 3, 94)	10	0m10s
	(90, 5, 5)	(20, 20, 60)	12	0m12s
	(60, 20, 20)	(33.3, 33.3, 33.3)	5	0m05s

Table 4.2: *Performance figures for optimal parameterizations obtained using the continuation method. The iterations refer to the steps taken in the energy minimization of $E[\cdot]_{\alpha_1}$, starting from an initial parameterization that minimizes $E[\cdot]_{\alpha_0}$.*

4.1.3 Robustness to Surface Discretization

The continuous energy model that we use has several distinct advantages. Since our energy is defined in the continuous setting, the solution to the minimization problem is insensitive to the particular discretization of the surface. This fact can be observed because different discretizations minimize consistent approximations of the same continuous energy functional. In Figure 4.3, we consider the successive refinement of a triangular grid that represents a stamped mechanical part. Notice how the parameterization converges rapidly, and its main features are apparent even at the coarsest resolution. Not surprisingly, the number of Newton steps taken for each of these examples is nearly identical (cf. Table 4.1).

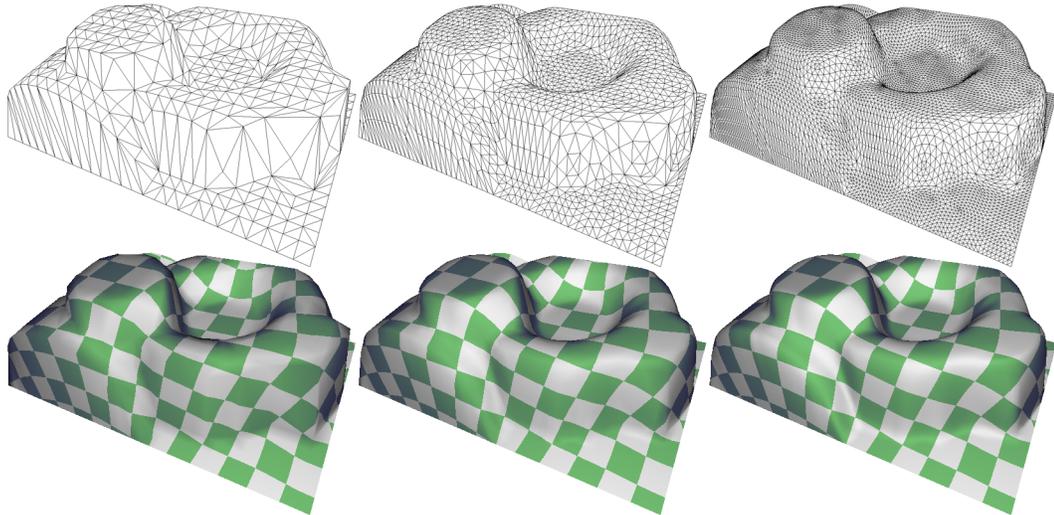


Figure 4.3: *Successive refinement of a triangulation demonstrates the stability of the parameterization, due to the fact that the different discretizations minimize consistent approximations of the same continuous energy functional.*

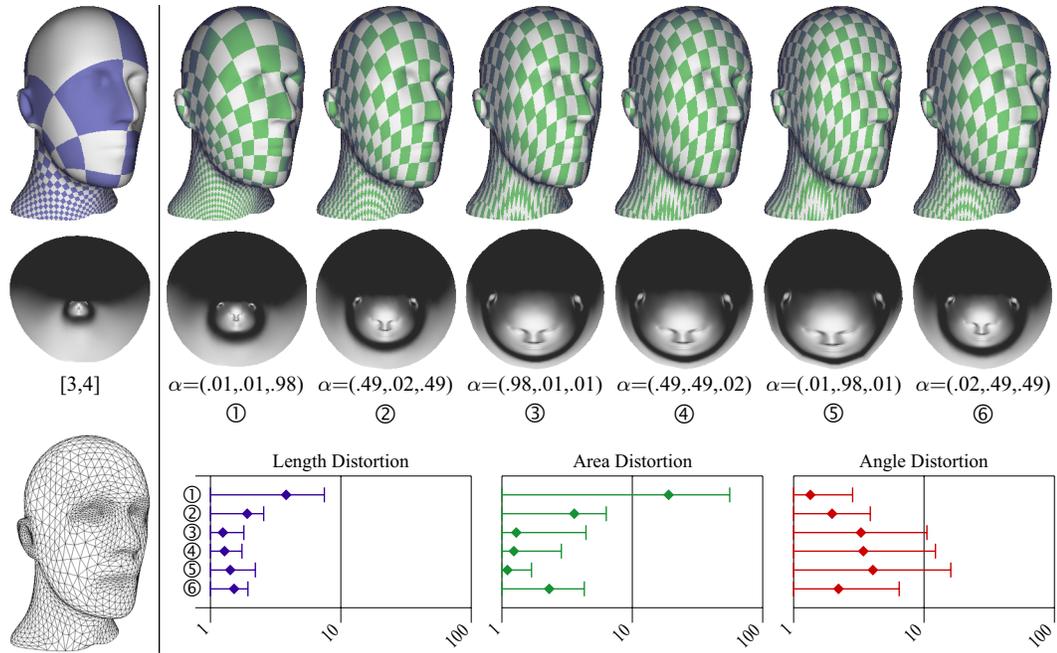


Figure 4.4: The parameterization is controlled by parameters $\alpha = (\alpha_l, \alpha_a, \alpha_c)$ which capture the trade-off between length, area and angle preservation, respectively. Texture maps of the initial parameterization on the left and solutions for different parameter settings are shown on top, with the parameter domain beneath. A plot of the maximum, minimum and mean distortion is shown below. Notice that the measured distortion tends to reflect the trade-offs in the choice of parameters.

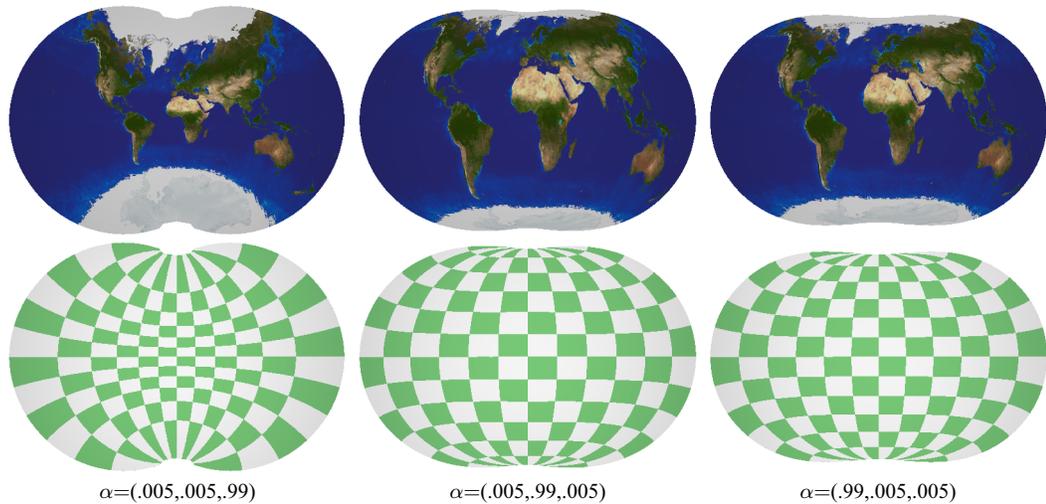


Figure 4.5: Flattening the earth. The sphere is opened with a semi-circular cut and its parameterization is optimized for angle, area and length preservation, respectively. Here texture maps show the deformation of the surface in the parameter domain.

4.1.4 Distortion in Optimal Parameterizations

Figure 4.4 demonstrates our algorithm for a variety of parameter values. Given the initial parameterization on the left, each parameterization on the right is obtained by minimizing the corresponding energy. The relative importance of the three competing parameters $\alpha_l, \alpha_a, \alpha_c$ is reflected in the distortion of the final parameterization. Recall that distortion can be measured as a function of the maximum and minimum eigenvalues Γ, γ of the first fundamental form of the parameterization function. We calculate the mean distortion for a parameterization as the arithmetic sum of the contributions from each triangle in our discretization. Popular measures of distortion for angles ($\sqrt{\Gamma/\gamma}$), area ($\sqrt{\Gamma\gamma}$) and length ($\sqrt{\Gamma + \gamma}$) for Figure 4.4 are reported in the bottom graphs, where the extreme deformations in the parameterization functions clearly show the trade-offs between the different distortion measures. This correspondence is further explored in Figure 4.5, which demonstrates the effect of minimizing the different energy terms on parameterizations of a sphere that has been cut.

4.1.5 Comparison to Related Work

We have presented a parameterization method that gives the user flexible control over the trade-off between area, length and angle distortion. Figure 4.4 demonstrates this control with examples that range from angle preserving to area preserving, and a comprehensive balance with length distortion

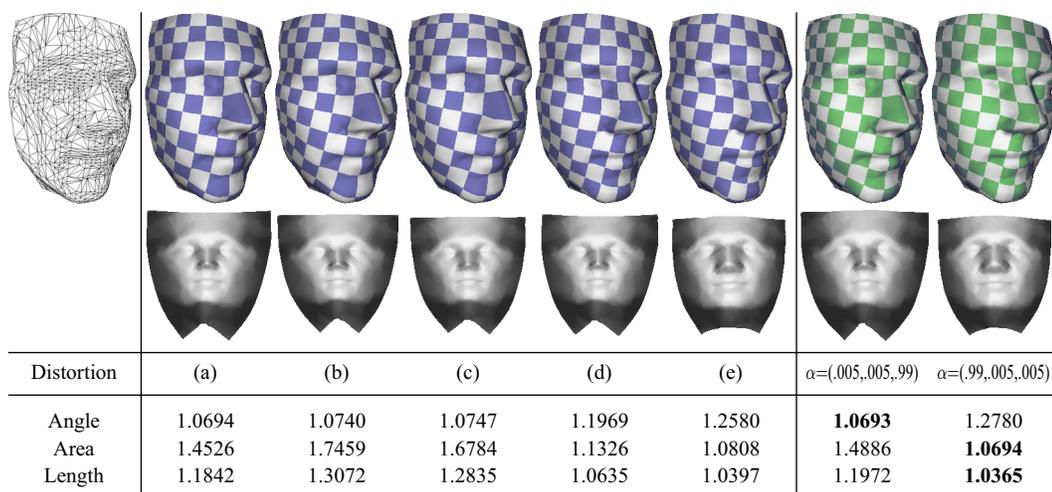


Figure 4.6: A comparison of related work to parameterizations generated by our method. Unlike other methods, we are not restricted to optimizing for only one notion of distortion, as demonstrated in the table of mean distortion measurements above. The related work considered here includes (a) [15, 29], (b) [46], (c) [25], (d) [14] with $\theta = 1$ and (e) [44].

in the rest. This contrasts our method from previous work, which largely focused on optimizing for only one measure of distortion, or possibly using linear combinations of energies. Figure 4.6 demonstrates that this trade-off is highly effective. For example, when emphasizing conformality, we achieve a smaller mean distortion than methods which were geared towards the construction of conformal mappings. The same is true for the stretch minimization (note however that we only minimize length distortion with respect to the L^2 norm, not the L^∞ norm as in [44]). Because of the inherent non-linearity of our method, we find that the discrete conformal mappings which result from linear systems provide us with an excellent starting point for our optimization.

4.2 Surface Matching

We continue our investigation of the parameterization algorithms with an analysis of the surface matching method. At the heart of this approach lie the parameterizations generated in the previous section which are rasterized as images of the geometric properties to be matched. Our experiments feature models of faces because of the high variability in detail and the size of features. As well, the degree of scrutiny that we subject them to allows one to be more critical of the correspondences in our examples.

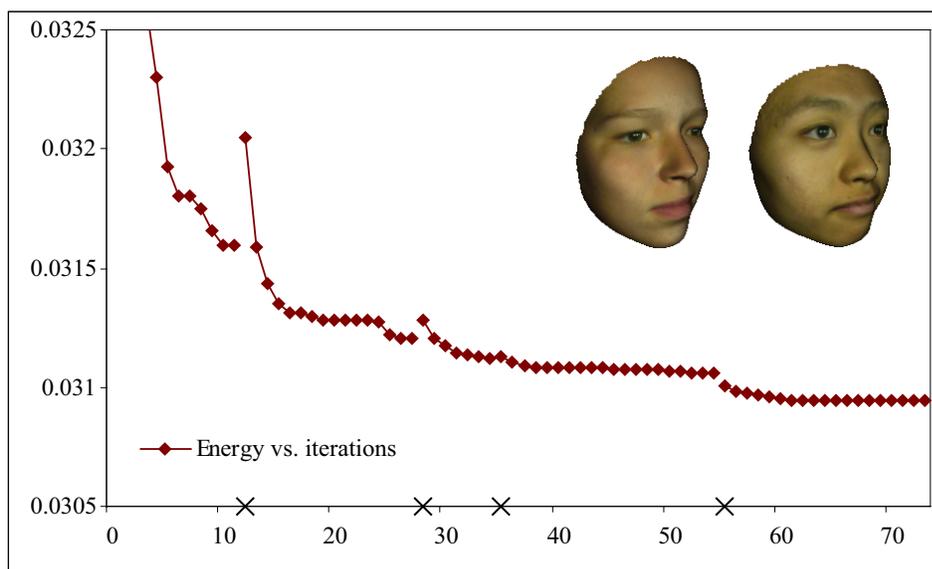


Figure 4.7: *Plot of the energy decay for the sequence of deformations computed in the matching energy minimization algorithm on multiple scales. The staircase-like behaviour results from the jumps onto the next finer scale (marked “×”), which may increase the energy temporarily.*

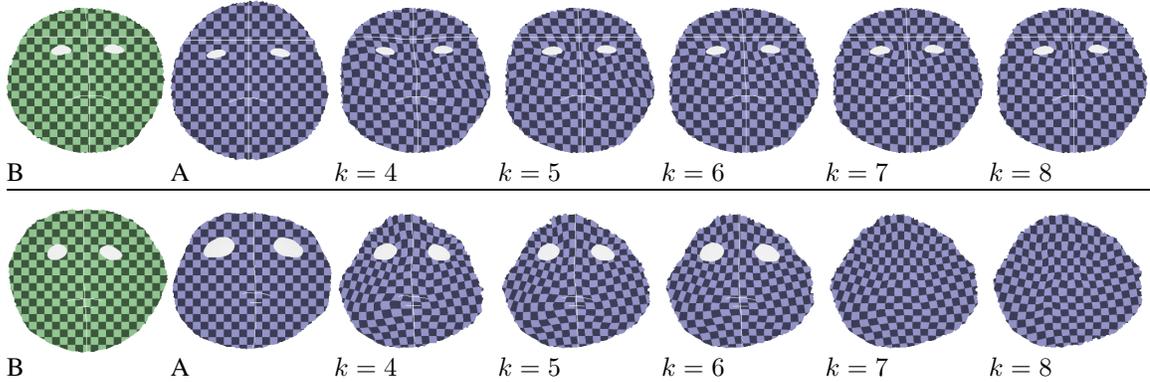


Figure 4.8: For a matching pair A, B , the matching algorithm generates deformations from coarse to fine scales. The feature sets effectively guide the optimization algorithm at coarser scales. At finer scales, features can be removed to match purely by shape (bottom).

4.2.1 Convergence of the Optimization Algorithm

To analyze the rate of convergence of the matching algorithm we consider plotting the energy of the system at each iteration of the global minimization algorithm. Figure 4.7 shows a representative energy decay which exhibits a characteristic staircase-like behaviour at the change in levels in the multiscale algorithm. One can experience a temporary increase in the energy at these transitional points as more detail is introduced to the refined image.

It is also useful to track the progress of the algorithm by analyzing the intermediate deformations in the plane. Figure 4.8 illustrates the sequence of deformations that are generated by the multiscale optimization algorithm. The optimal deformations $\tilde{\phi}^k$ that minimize $E^{\sigma_k}[\cdot]$ are guaranteed to be smooth and bijective. Notice how the features serve to guide the match in the coarser stages of the optimization. In the lower example, the feature energy is eliminated at finer scales to allow the match to be optimized based purely on the geometry of the surfaces. This technique is very effective at providing guidance at coarser levels for matches that involve large deformations.

Execution times for our examples are given in Table 4.3, which lists the matching pairs, the sizes of the triangle meshes, and the number of iterations and computation times for the energy minimization. Performance figures were recorded on a 3.6 GHz Intel Pentium IV processor, on which the computation time to generate the geometry images in the examples is negligible. The number of iterations at each level of the multiscale energy minimization algorithm is determined by a threshold on the energy decrease (we use a cut-off value of 10^{-6} in our experiments). In general we have observed that a larger cut-off gives very similar results in a fraction of the time, since this greatly reduces the number of iterations at the finest levels.

\mathcal{M}_A	\mathcal{M}_B	V_A	V_B	Iterations	Time
lily	lilygrin	12614	13512	44	0m59s
lilygrin	lilysmile	13512	14032	59	2m31s
lily	weiwei	12614	14265	65	2m57s
igea	maxplanck	14611	17755	101	5m45s
armadillo	gargoyle	84935	75914	55	0m17s

Table 4.3: *Performance figures for the matching examples (V_A, V_B denote # vertices for the triangle meshes $\mathcal{M}_A, \mathcal{M}_B$ respectively).*

4.2.2 Robustness to Image Discretization

In taking an image processing approach to defining a correspondence between two surfaces, one must be concerned with the discretization of the geometric properties into images. This discretization is controlled by the parameterizations of the surfaces to the plane, and therefore the approximation of the surface properties in the images is subject to the distortion in the parameterizations. Figure 4.9 features a match between two surfaces which have a high degree of distortion in their parameterizations. The distortion is balanced according to the algorithm in Section 3.1, however the shapes of the surfaces nonetheless involve a large area factor in flattening them to the plane. As we illustrate here, our correspondence behaves well despite the large deformation between the surfaces and the distortion in the image domain.

In the presence of a gross misalignment of the images, such as a 180° rotation of one of the image domains in Figure 4.10, one cannot expect the optimization algorithm to correct the image registration. However this example shows that our algorithm robustly computes an optimal deformation that is smooth and free of folds. For illustration purposes, the second model is rotated in the blend to make the marked features more visible.

4.2.3 Partial Correspondence

A difficult challenge for matching disc-like surfaces is the presence of a partial correspondence, where certain regions have no matching counterpart under the correspondence. This is particularly prevalent in boundary regions, where it is actually quite undesirable to define the mapping because no practical correspondence exists.

Figure 4.11 demonstrates the robust treatment of a partial correspondence in the presence of large deformations in the matching. This example is made even more challenging due to the gross differences in the surfaces and mismatched features. We find that using coarse features as hints and relying on the bending energy for accuracy at finer resolutions is a good strategy in this case.

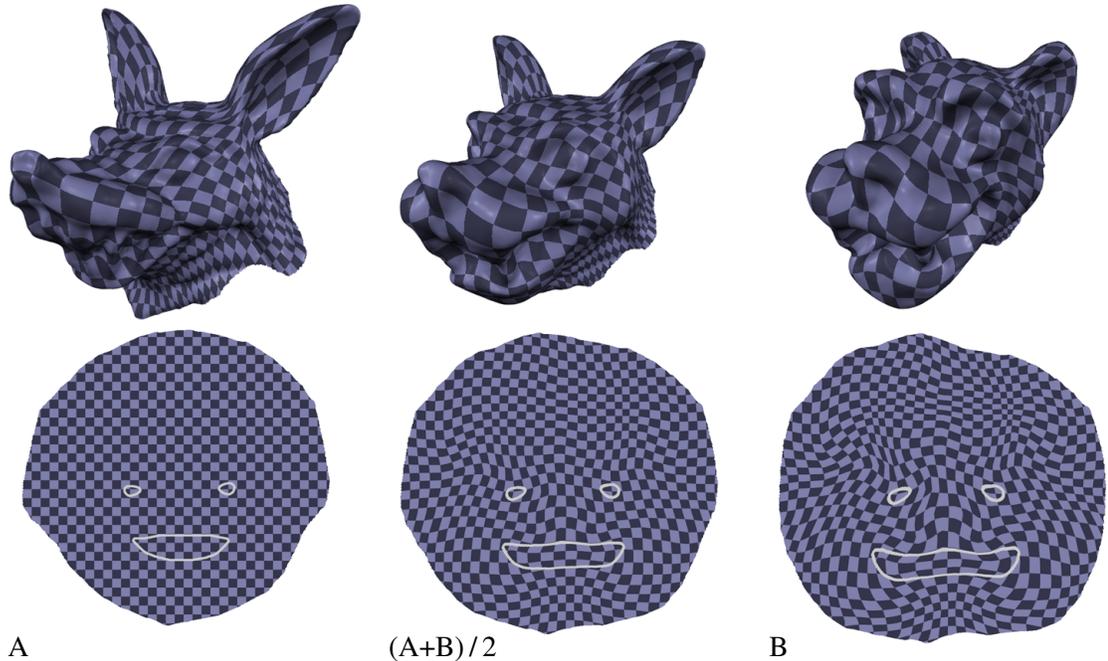


Figure 4.9: Large deformations are often needed to match surfaces that have very different shapes. A checkerboard is texture mapped onto the first surface as it morphs to the second surface (top). The matching deformation shown in the parameter domain (bottom) is smooth and regular, even where the distortion is high (e.g., around the outlines of the mouth and eyes).

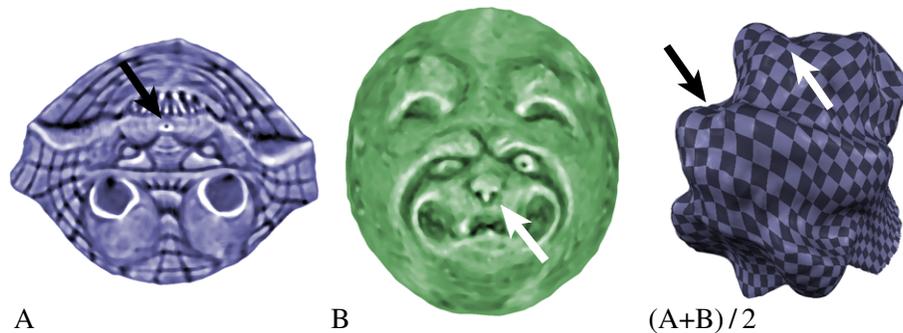


Figure 4.10: Our method guarantees that the resulting match is well-behaved, even in the presence of gross misalignments such as the 180° rotation of the left image shown here. The 50% morph on the right is smooth and regular (features in the domain are marked with arrows on the surface).

4.2.4 Comparison to Related Work

We have presented an approach to surface matching that addresses the problems with robustness and computational efficiency seen in previous methods. One class of algorithms is based on partitioning the surfaces into a common network of parametric domains. This either requires a great deal of user interaction to do the partitioning by hand [28, 43] or involves a difficult algorithm with some user guidance to get a good partitioning [27, 45]. The combinatoric complexity of the meshes

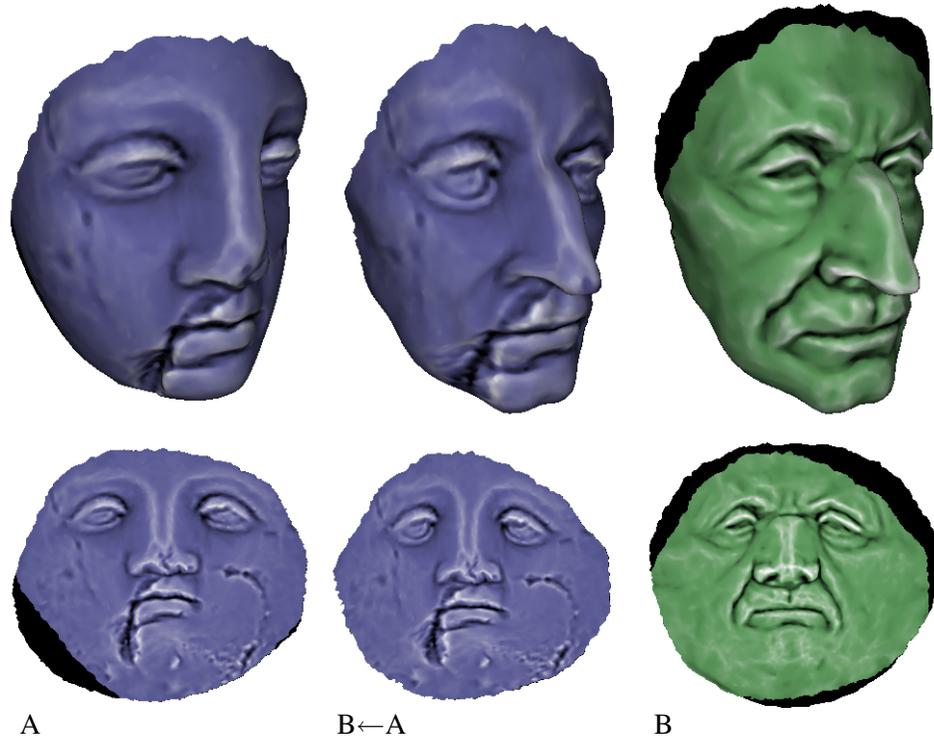


Figure 4.11: *For surfaces with boundaries, a partial correspondence is often desired. The correspondence is defined where their parameter domains intersect under the matching deformation (bottom). In this domain, quantities such as texture maps can be mapped between the surfaces (center). The unmatched regions are in black.*

determines the performance of these methods, which can be quite costly (on the order of hours). The final matching can only be controlled by changing the parametric domain layout, and it is not clear whether very fine features, for example from texture maps, can be matched with high precision. Our method neither experiences these performance drawbacks nor is the user responsible for manually partitioning the surfaces.

Another class of methods matches a template surface to range scan data directly in 3D [1, 50]. These methods are particularly well suited to matching surfaces when a parameterization cannot easily be computed, such as in the presence of holes. However these algorithms require a manual 3D alignment and the surfaces must be in similar poses to get good results from their matching of normals. In contrast, we can handle large deformations (cf. Figure 4.9) and different poses (cf. Figure 5.4), and only require a coarse alignment of the parameter domains to get a good match. Unlike previous methods, the smoothness, bijectivity and existence of our solutions is guaranteed. In practice, our matching algorithm robustly computes an optimal deformation that is smooth and free of folds (Figure 4.10). It is well known in image matching [20] that uniqueness of solutions

cannot be guaranteed, nor is this a normal practical requirement: typically an application does not seek a unique match, but instead desires a smooth correspondence function with a good correlation between salient features. We provide this through the *analytic guarantees* of our approach and the *flexibility* afforded to the user to construct good matches.

Chapter 5

Applications

With our new tools in hand, we will now consider some classical applications of surface parameterizations. Common tasks in computer graphics like texture mapping and 3D morphing benefit from the versatility and robustness of our approach. This chapter features non-trivial examples of these applications and also looks at a novel approach to facial animation. These examples are particularly challenging because they use physically-acquired data that have unique problems to overcome.

5.1 Texture Mapping

Texture mapping is one of the most venerable applications of parameterizations for triangle meshes. This simple form of surface illustration is accomplished by assigning 2D texture coordinates to the vertices in the mesh. Each triangle is rendered with the pixels of a texture image that lie within its 2D pre-image defined by the coordinates at its vertices. This rudimentary technique nonetheless offers an incredible improvement in the level of detail one can achieve from rendering triangles, and it is an indispensable tool in both low- and high-end computer-generated visual effects.



Figure 5.1: *Photographs of the subject are used to generate the texture image.*

The central goal and biggest challenge for texture mapping is to achieve high quality (possibly photo-realistic) effects, often on a limited memory budget. This affects the choice of texture coordinates (*i.e.*, the surface parameterization) in two ways. Primarily, the sampling of the image space must provide enough detail for the 3D rendered triangles. One can always guarantee a sufficient sampling rate by increasing the resolution of the image itself, but this is not a practical solution under memory constraints. Therefore to achieve a suitable sampling rate, one desires a parameterization with low distortion, particularly through the control of length [44]. However since an artist often uses image editing software to design the texture image, it is desirable to have low shear distortion by introducing some control over angles. A secondary concern is the appearance of visual artifacts, especially when the surface is animated (*e.g.*, through a time-varying free-form deformation). Kinks and folds in the texture map are unacceptable if one aims to promote the suspension of disbelief. Thus one also requires that the parameterization is smooth and bijective.

Our method of surface parameterization is well suited to these challenges, since the texture coordinates that we generate have all of these qualities by construction. The following example is a demonstration of texture mapping that uses 3D facial data captured with structured-light for the geometry acquisition and digital photography for the image acquisition. The data consists of a triangle mesh (Figure 5.2a) and two colour photographs of the subject taken from different angles (Figure 5.1). Texture coordinates are assigned to each triangle in the mesh, corresponding to the appropriate region of the photograph. As is, the data is unsuitable for further processing for the following reasons:

1. The mesh surface is non-manifold and has extraneous data around the boundaries.
2. The triangle mesh is noisy due to low-confidence sample points from the acquisition, particularly at grazing angles and on reflective surfaces such as the eyes.
3. The parameterization is not continuous since the texture map is built from two photographs.

First the mesh is repaired by removing non-manifold artifacts, and then the face is extracted as a manifold subset of the original surface (Figure 5.2b). To address the second problem, we smooth the surface by displacing the mesh vertices in the normal direction [16]. The resulting surface is fit for texture mapping. We proceed to construct a parameterization of this surface that balances angle, area and length distortion appropriately. Then through this parameterization, a new texture image is generated by rendering the triangles of the *original* triangle mesh in the new parameter domain (Figure 5.2d). This image is now ready for final editing and is then applied to the new surface as a texture map (Figure 5.2c).

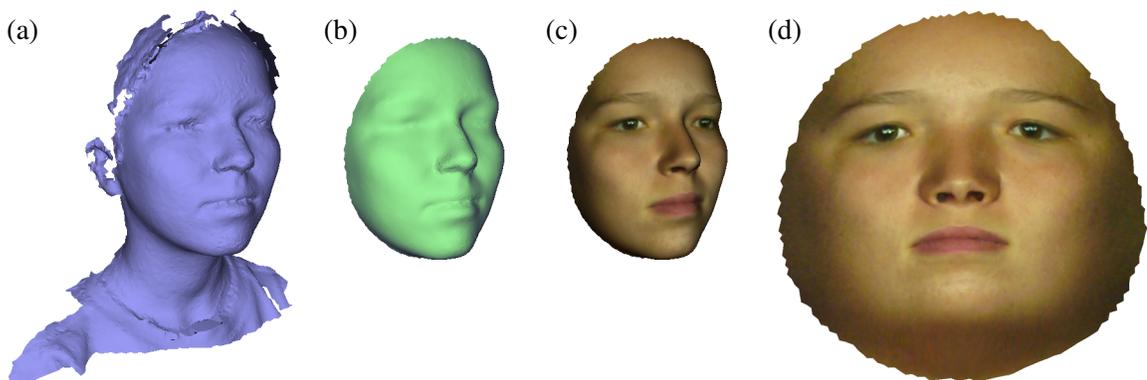


Figure 5.2: *Texture mapping physically-acquired data: (a) the original triangle mesh; (b) the mesh is smoothed and the face extracted; (c,d) a new texture image is generated in the parameter domain and texture mapped on the face, using a parameterization constructed with our method.*

5.2 3D Morphing

Morphing is a celebrated technique for special effects, famous for “wolf-man” transformations on the silver screen [49]. Simply put, a morph is an animation that transforms one object into another, with the most common subject being human (or human-like) faces. The central element in a morph is the correspondence between the before and after images. In 2D morphing, this correspondence is a map from one image to another, while in 3D one needs a point-wise correspondence between two surfaces. The simplest morphing technique is a linear blend between the objects which is evaluated over time to generate an animation. The quality of the correspondence is critical because any mismatch will be glaringly obvious in the cross-fade between the objects.

Our surface matching method provides an approach to defining a correspondence for 3D morphing. One begins by constructing a correspondence $\phi_{\mathcal{M}}$ for surfaces \mathcal{M}_A and \mathcal{M}_B . Then for each point x on \mathcal{M}_A , the corresponding point on \mathcal{M}_B is given by

$$\phi_{\mathcal{M}}(x) := (x_B \circ \phi \circ x_A^{-1})(x)$$

which is defined on $x_A(\omega_A[\phi])$. The morph is then animated over time $t \in [0, 1]$ with

$$\mathcal{M}(t) := (1 - t)\mathcal{M}_A + t\phi_{\mathcal{M}}(\mathcal{M}_B)$$

Examples of 3D morphing are given as sequences of animation stills in Figure 5.3. Matching faces is particularly challenging because of the incredible degree of scrutiny that they receive. Notice that the salient features are perfectly matched, and the deformation in time is smooth. Both of these qualities are a direct result of the control of shape and feature matching that is inherent in our approach. The high quality of these examples is prohibitively difficult to achieve with other methods, where the correspondences must often be painstakingly generated by hand. In particular, semi-automatic methods typically require the user to specify point constraints that can break the regularity of the deformation, and manually producing consistent mesh representations of the surfaces is both tedious and limited by the resolution of meshes that a human can process. In contrast, the user interaction in our approach is limited to drawing corresponding features on the surfaces. The time to generate a polished morph is about an hour for a user who is familiar with the system, and is typically spent by adding features and rendering an animation sequence, which only takes a few minutes to generate.



Figure 5.3: *Examples of 3D morphing. The correspondence between the surface pairs is generated by matching the shape and user-selected features of one surface to the other.*

5.3 Facial Animation

“Blend shapes” is a popular technique for facial animation [41, 42]. In it, multiple key expressions are modeled for the subject. Intermediate expressions are then generated by a linear combination of the key models. Like 3D morphing, this blending requires a consistent parameterization between the keys, and therefore this approach to facial animation has the same challenges. In light of the shortcomings of other 3D morphing techniques, the artist typically resorts to deforming a template mesh into the key expressions by hand, so that the correspondence between the poses is defined implicitly by the mesh structure.

Our surface matching technique provides an alternative approach that lifts the user from these restrictions. Since our correspondence is defined independently from the actual triangulation, we can use high resolution meshes for the key expressions. The only consistency that we require is an identification of the salient features of the face at each key pose, for example the eyes, mouth and lines of symmetry in the face (cf. Figure 3.5).

Figure 5.4 shows an example of a proof-of-concept implementation of this technique. Here the goal is to produce a continuous animation sequence by realistically blending between three keyframe poses of the same subject. This is particularly challenging, due to the presence of features that are either rigid or highly deformable (*i.e.* the eyes and mouth, respectively). We mask out these “holes” with features in order to match their boundaries precisely. The animation has two segments, one which morphs between $\mathcal{M}_A, \mathcal{M}_B$ and the other between $\mathcal{M}_B, \mathcal{M}_C$. We begin by matching the surfaces pairwise to get ϕ_{AB} and ϕ_{BC} . Then we use ϕ_{AB} to morph between \mathcal{M}_A and \mathcal{M}_B in the first segment, and in the second segment we apply the composition $\phi_{BC} \circ \phi_{AB}$ to blend $\phi_{AB}(\mathcal{M}_A)$ with \mathcal{M}_C , thus guaranteeing a seamless transition. A sequence of animation stills is shown in Figure 5.5. The eyes and mouth are masked out because unlike the skin, they do not undergo elastic deformations. The same mesh can be used throughout the animation in order to save on memory resources in the application. More key expressions can be easily added by matching to a similar key model and remeshing it with the template structure (typically the neutral expression) using an extension of our 3D morphing technique. In this scenario, selective regions of the key poses are typically given different blending ratios to compose more complicated expressions.

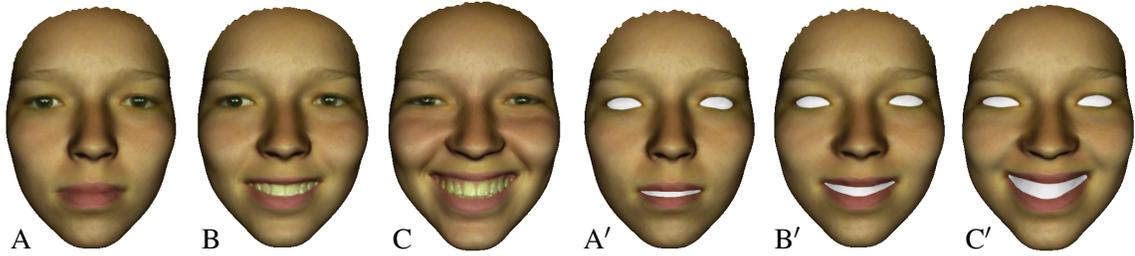


Figure 5.4: Morphing through keyframe poses A, B, C is accomplished through pair-wise matches $A \rightarrow B$ and $B \rightarrow C$. Starting with A we blend its shape into B using $A \rightarrow B$, and then morph to C by applying $A \rightarrow B$ followed by $B \rightarrow C$. The texture from A is used throughout. Because of the close similarity in the poses, one can expect the intermediate blends A', B', C' to correspond very well with the original keyframes A, B, C , respectively.



Figure 5.5: Facial animation using blend shapes. The initial surface and texture is animated through three key expressions that were generated from 3D scans of the same subject.

Chapter 6

Conclusion

We have presented two methods of surface parameterization. By deriving our approach from a rigorous variational formulation in the continuous setting, we have guaranteed the quality of our parameterizations and the performance of the algorithms which construct them. This chapter summarizes our contributions and motivates further exploration in directions such as parameterizations of higher genus surfaces, geodesic curves in deformation space and free-form deformations of discrete parametric surfaces.

6.1 Summary

In chapter 2 of this thesis, we developed a rigorous variational framework for surface parameterization. We addressed two important classes of parametric mappings, namely, parameterizations over a region of the plane, and correspondences that map between two surfaces. A critical distinguishing trait of our approach is its roots in the classical theory of rational mechanics. From this foundation we are able to bring theoretical guarantees along with our treatment of parameterizations, including proof of the existence of variationally optimal parameterizations.

Chapter 3 described a set of algorithms for constructing parameterizations of triangle meshes. Due to the theoretical underpinnings of the previous chapter, we apply a straightforward finite element method toward the discretization of our parameterization energies over the space of surface triangulations. This leads to flexibility in the design of the energy minimization method, which is chosen from the mature and extensive body of work in numerical optimization. Our implementation provides user control over the parameterization, particularly in the way that distortion is handled.

An analysis of our algorithms was provided in Chapter 4. We focus our experimentation on range scan data of physical surfaces. The triangle meshes obtained from this process are particularly challenging to parameterize due to the high degree of variability in the shapes of the surfaces and the quality of the discretization. We demonstrated the efficiency of our algorithms through an analysis of the energy minimization and by the fast execution times for our examples. Finally, we showed our methods are robust in a challenging set of examples. In all cases, we generate parameterizations that are smooth, bijective and have low distortion.

Chapter 5 presented a variety of applications of surface parameterization. First high-quality texture maps were created from range scan data and photographs of the subject's face. These parameterizations were then featured in novel applications of 3D morphing to highlight the versatility of our surface correspondences. Our approach to surface matching is a major contribution to surface parameterizations in its own right, as well as a clever application of our parameterizations in the plane.

6.2 Future Work

The work presented here invites many possible directions for continued research.

Parameterizations of Higher Genus Surfaces Continuing the theme of parameterization to surfaces that are not homeomorphic to a disc requires one to address challenging theoretical and practical problems. Cylindrical, spherical, toroidal and higher genus domains do not embed into the plane and therefore new variational approaches to optimal parameterization must be derived to extend the properties of our planar parameterizations to the arbitrary topology setting. A new approach based on overlapping parametric patches on the surface addresses the instability of high dynamic range parameterizations and has a direct application toward texture mapping in character animation. As the resolution of triangulated surfaces escalates due to high-precision range scan data and next-generation modeling tools, the desire to process these surfaces must be met by non-linear multigrid methods on unstructured triangle grids that provide stable and efficient algorithms for parameterization and other variational methods.

Geodesic Curves in Deformation Space Given a correspondence between surfaces in a discrete sequence sampled over time, it is natural to ask for an optimal surface deformation which interpolates each surface in turn to achieve a realistic, smooth animation (*e.g.*, facial animation of a live actor or a synthetic character). One may also ask to apply this deformation to an entirely new surface, whereby it must adapt to the shape of the second subject. Both scenarios are related to finding shortest paths in a space of deformations that is described by a time-dependent variational energy. Our previous work on surface matching provides both a tool for the correspondences and a framework for a deformation energy that measures the variation of shape over time subject to user control over the animation.

Free-Form Deformation of Discrete Parametric Surfaces The practical and theoretical guarantees of our previous work on surface parameterization introduces discrete parametric surfaces as a new geometric modeling primitive. Leveraging the mature body of modeling and animation techniques built up through the decades of study of parametric surface patches, one can develop new methods for the free-form deformation of discrete parametric surfaces. The aim of this research is to overcome the limitations of geometric modeling with subdivision surfaces, the de facto standard of modern-day commercial animation, while providing all of its benefits. The challenges of this en-

deavour will require new methods in surface remeshing, parametric correspondence, and free-form deformation techniques.

Appendix A

Surface Parameterization

A.1 Representation of the Parameterization Energy

Theorem A.1 (Representation of the energy) *Let \mathcal{M} be a smooth surface patch, which is parameterized by an initial parameterization x over a set $\omega \subset \mathbb{R}^2$. Under the assumptions of elasticity (2.1), frame indifference (2.2) and isotropy (2.3), there exists a function $W : \mathbb{R}^2 \rightarrow \mathbb{R}$ such that $W_{\mathcal{M}}(\mathbf{D}(\phi \circ x^{-1})(x)) = W(\iota_{\mathbf{C}[\phi]})\sqrt{\det g}$ where $\iota_A = (\text{tr}(A), \det(A))$ are the two principal invariants of $A \in \mathbb{R}^{2,2}$ and the Cauchy-Green strain tensor is defined by*

$$\mathbf{C}[\phi] = D\phi g^{-1} D\phi^T .$$

Furthermore, the expression $g \in \mathbb{R}^{2,2}$ is the metric tensor w.r.t. the parameterization x , i.e., $g = Dx^T Dx$. Thus, the parameterization energy can be written as

$$E[\phi] = \int_{\omega} W(\text{tr}(\mathbf{C}[\phi]), \det(\mathbf{C}[\phi]))\sqrt{\det g} d\xi .$$

Proof Due to our assumptions, g^{-1} is symmetric, uniformly positive definite and uniformly bounded. Suppose A is an arbitrary invertible matrix in $\mathbb{R}^{2,2}$. By the polar decomposition theorem there exists an orthogonal matrix Q such that $A = QU$ where $U = (A^T A)^{\frac{1}{2}}$ is symmetric and positive definite. Hence, by the assumption of frame indifference (2.2) we obtain

$$\hat{W}(A) = \hat{W}(QU) = \hat{W}(U) = \hat{W}((A^T A)^{\frac{1}{2}})$$

Next, we define $\tilde{W}(C) := \hat{W}(C^{\frac{1}{2}})$ for all symmetric, positive definite matrices C and obtain

$$\hat{W}(A) = \tilde{W}(A^T A)$$

From this and the assumption of isotropy (2.3) one deduces

$$\tilde{W}(A^T A) = \hat{W}(A) = \hat{W}(AQ) = \tilde{W}(Q^T A^T AQ)$$

for all g -orthogonal matrices Q . In particular, for symmetric, positive definite matrices C we get

$$\tilde{W}(C) = \tilde{W}(Q^T C Q) \tag{A.1}$$

Next, we show that whenever $\iota_{g^{-1}A} = \iota_{g^{-1}B}$ for symmetric, positive definite matrices $A, B \in \mathbb{R}^{2,2}$, then $\tilde{W}(A) = \tilde{W}(B)$. Thus, we immediately deduce that there is a function W such that $\tilde{W}(A) = W(\iota_{g^{-1}A})$ and hence for general matrices $A \in \mathbb{R}^{2,2}$,

$$\hat{W}(A) = \tilde{W}(A^T A) = W(\iota_{g^{-1}A^T A})$$

which is the assertion in our theorem.

First, we observe that for any symmetric matrix A , the matrix $g^{-1}A$ is g -self adjoint. Indeed, $(g^{-1}A)^* = g^{-1}A^T g^{-1}g = g^{-1}A$. Thus, we deduce that $g^{-1}A$ and $g^{-1}B$ diagonalize with respect to the scalar product “ \cdot_g ” and we obtain

$$\begin{aligned} g^{-1}A &= \lambda_1 r_1 (g r_1)^T + \lambda_2 r_2 (g r_2)^T \\ g^{-1}B &= \mu_1 s_1 (g s_1)^T + \mu_2 s_2 (g s_2)^T \end{aligned}$$

with $g r_i \cdot r_j = \delta_{ij}$ and $g s_i \cdot s_j = \delta_{ij}$. Due to our assumption $\iota_{g^{-1}A} = \iota_{g^{-1}B}$ we achieve $\mu_i = \lambda_i$ after a possible reordering. Furthermore, there exists a g -orthogonal matrix Q with $Q s_i = r_i$. Thus we get

$$\begin{aligned}
A &= gg^{-1}A = \sum_{i=1,2} \lambda_i g Q s_i (g Q s_i)^T \\
&= \sum_{i=1,2} \lambda_i g Q s_i s_i^T Q^T g \\
&= \sum_{i=1,2} \lambda_i Q^{-T} g s_i s_i^T g Q^{-1} \\
&= Q^{-T} g \sum_{i=1,2} \lambda_i s_i (g s_i)^T Q^{-1} \\
&= Q^{-T} g g^{-1} B Q^{-1} \\
&= Q^{-T} B Q^{-1}
\end{aligned}$$

Here we have used the fact that $gQ^{-1} = Q^T g$ holds for g -orthogonal matrices Q and that Q^{-1} is g -orthogonal if Q is g -orthogonal, which implies $gQ = g(Q^{-1})^{-1} = Q^{-T}g$. Now, (A.1) applied to the g -orthogonal matrix Q^{-1} leads to

$$\tilde{W}(A) = \tilde{W}(Q^{-T} B Q^{-1}) = \tilde{W}(B),$$

which is what we have claimed above. Finally, we obtain for the principal invariants

$$\begin{aligned}
\text{tr}(g^{-1} D\phi^T D\phi) &= \text{tr}(D\phi g^{-1} D\phi^T), \\
\det(g^{-1} D\phi^T D\phi) &= (\det g)^{-1} (\det D\phi)^2.
\end{aligned}$$

Lemma A.2 *Suppose \mathcal{M} is a smooth surface, and x is an initial parameterization of \mathcal{M} over a set $\omega \subset \mathbb{R}^2$. A deformation ϕ which turns the parameterization $x[\phi]$ into an isometry is a local minimizer of the elastic energy $E[\cdot]$ with the density defined in (2.5) if*

$$\alpha_l > 0, \quad \alpha_a > 0, \quad \alpha_c \geq 0.$$

Proof Let us take derivatives of the energy integrand W with respect to the inverse metric $q = Ag^{-1}A^T$ in the direction of an arbitrary symmetric matrix $C \in \mathbb{R}^{2,2}$:

$$\begin{aligned} \partial_q W(\iota_q)(C) &= \alpha_l \operatorname{tr}(C) + \alpha_a \left(\det q - \left(\frac{\alpha_l}{\alpha_a} + 1 \right) (\det q)^{-1} \right) \operatorname{tr}(q^{-1}C) \\ &\quad + \alpha_c \left(2 \frac{\operatorname{tr}(q)}{\det(q)} \operatorname{tr}(C) - \frac{\operatorname{tr}(q)^2}{\det(q)} \operatorname{tr}(q^{-1}C) \right), \\ \partial_q^2 W(\iota_q)(C, C) &= \alpha_a \left(\det q + \left(\frac{\alpha_l}{\alpha_a} + 1 \right) (\det q)^{-1} \right) \operatorname{tr}(q^{-1}C)^2 \\ &\quad + \alpha_a \left(\left(\frac{\alpha_l}{\alpha_a} + 1 \right) (\det q)^{-1} - \det q \right) \operatorname{tr}(q^{-1}Cq^{-1}C) \\ &\quad + \alpha_c \left(2 \frac{\operatorname{tr}(C)^2}{\det(q)} - 2 \frac{\operatorname{tr}(q)\operatorname{tr}(C)\operatorname{tr}(q^{-1}C)}{\det(q)} + \frac{\operatorname{tr}(q)^2}{\det(q)} \operatorname{tr}(q^{-1}C)^2 + \frac{\operatorname{tr}(q)^2}{\det(q)} \operatorname{tr}(q^{-1}Cq^{-1}C) \right) \end{aligned}$$

One can verify that $q = \mathbb{I}$ is a local extremum of $W(\iota_q)$. Under our usual assumptions every diffeomorphism $\phi : \omega \rightarrow \phi(\omega)$ induces a positive definite inverse metric $q = g^{-1}[\phi]$. Hence, $q = \mathbb{I}$ is a local minimizer of $W(\iota_q)$ if $\partial_q^2 W(\iota_{\mathbb{I}})$ is positive definite. This leads to

$$\partial_q^2 W(\iota_{\mathbb{I}})(C, C) = (\alpha_l + 2\alpha_a) \operatorname{tr}(C)^2 + \alpha_l \operatorname{tr}(C^2) + 6\alpha_c \operatorname{tr}(C)^2 > 0$$

for all symmetric matrices $C \neq 0$. Here we have used the fact that $\operatorname{tr}(C^2) > 0$ for any symmetric matrix C with $C \neq 0$, and we have taken into account that $\partial_q W(\iota_{\mathbb{I}})(C) = 0$ to obtain the inequality. It follows that the inequality holds if (2.6) is satisfied.

A.2 Existence of Optimal Parameterizations

Theorem A.3 (Existence of an Optimal Parameterization) *Let \mathcal{M} be a smooth surface patch, which is parameterized by an initial parameterization x over a set $\omega \subset \mathbb{R}^2$. Furthermore, consider the energy in (2.5) over a set of admissible deformations ϕ , which are continuous and locally bijective. Then for parameters $\alpha_l, \alpha_a, \alpha_c$ which satisfy (2.6), there exists an admissible, minimizing deformation $\tilde{\phi}$.*

Proof Based on a classical result in the calculus of variations the proof is quite straightforward. Let us briefly sketch the main ingredients. We follow the well-known procedure of the direct methods in the calculus of variations [12, 35, 36], in particular the proof of the existence of minimizers for such polyconvex energy functions by Ball [4] (see also the monograph by Ciarlet [9], Theorem 7.7-1). Here, we only give a brief summary of the arguments. From the bi-Lipschitz assumption

on the initial parameterization we deduce that $g^{-1} \in L^\infty(\omega)$ and it is uniformly positive definite. Hence, $\sqrt{\text{tr}(Ag^{-1}A^T)}$ is actually a norm on linear mappings $A \in \mathbb{R}^{2,2}$.

The growth conditions imply that for every constant C there exists a ball $B_R(0)$ in \mathcal{A} with radius R such that $E[\phi] > C$ outside $B_R(0)$. Thus, there exists a minimizing sequence $(\phi^k)_{k \in \mathcal{N}}$ with $\|D\phi^k\|_{L^2(\omega)}$ and $\|\det D\phi^k\|_{L^2(\omega)}$ being uniformly bounded. Due to the constraint on the zero moment and a generalized Poincaré inequality, the $H^{1,2}(\omega)$ norm of ϕ^k is also uniformly bounded. Because of the reflexivity of $H^{1,2}(\omega) \times L^2(\omega)$ we can extract a weakly convergent subsequence from the minimizing sequence.

Applying well-known results on the weak continuity of the determinant we get that $(\phi^k, \det D\phi^k)$ converges weakly to $(\phi, \det D\phi)$. Furthermore, Ball proved that for the limit deformation $\det D\phi > 0$ holds almost everywhere in ω . Due to the convexity of W in both arguments we finally obtain a result on weak sequentially-lower semicontinuity and thus establish the existence of an admissible, minimizing deformation.

Corollary A.4 (Existence of an Optimal, Area-Preserving Parameterization) *Under the assumptions of Theorem A.3, and for the modified set of admissible deformations ϕ with $\det g^{-1}[\phi] = 1$, there exists a minimizing deformation $\tilde{\phi}$.*

Proof We can reformulate the additional constraint requiring that $\det D\phi = \sqrt{\det g}$. But $\det D\phi$ is weakly continuous on the original set of admissible deformations. Hence the modified admissible set is weakly closed. Thus, we can proceed as in the proof of Theorem A.3. We only have to ensure that there is at least one area preserving, homeomorphic parameterization $x[\bar{\phi}]$ for some deformation $\bar{\phi}$. But the existence of such a parameterization has been established by Moser and Dacorogna [13, 37].

A.3 Variations of the Parameterization Energy

The concrete formulas for the derivatives of the energy density $\hat{W}(A) = W(\iota_{Ag^{-1}A^T})$ defined in (2.5) are provided here. The fully discrete $\nabla \bar{E}$ and $\nabla^2 \bar{E}$ are given by

$$\begin{aligned} (\nabla \bar{E})_{ik} &= \sum_{T \in \omega_h} |T| \partial_A \hat{W}(A(T))(B^{ik}(T)) \sqrt{\det G(T)}, \\ (\nabla^2 \bar{E})_{ikjl} &= \sum_{T \in \omega_h} |T| \partial_A^2 \hat{W}(A(T))(B^{ik}(T), B^{jl}(T)) \sqrt{\det G(T)}, \end{aligned}$$

where $A(T) = D\Phi|_T$ as above and $B^{rs}(T) = D\Phi^{rs}|_T$. Furthermore,

$$\begin{aligned}
\partial_A \hat{W}(A)(C) &= \alpha_l \frac{\partial a}{\partial A}(C) + \alpha_a (1 - \beta d^{-2}) \frac{\partial d}{\partial A}(C) \\
&\quad + \alpha_c \left(\frac{2a}{d} \frac{\partial a}{\partial A}(C) - \frac{a^2}{d^2} \frac{\partial d}{\partial A}(C) \right), \\
\partial_A^2 \hat{W}(A)(C, H) &= \alpha_l \frac{\partial^2 a}{\partial A^2}(C, H) + \beta (2d^{-3}) \frac{\partial d}{\partial A}(C) \frac{\partial d}{\partial A}(H) \\
&\quad + \alpha_a (1 - \beta d^{-2}) \frac{\partial^2 d}{\partial A^2}(C, H) \\
&\quad + \alpha_c \left(\frac{2a}{d} \frac{\partial^2 a}{\partial A^2}(C, H) + \frac{2}{d} \frac{\partial a}{\partial A}(C) \frac{\partial a}{\partial A}(H) \right. \\
&\quad - \frac{2a}{d^2} \frac{\partial a}{\partial A}(C) \frac{\partial d}{\partial A}(H) - \frac{a^2}{d^2} \frac{\partial^2 d}{\partial A^2}(C, H) \\
&\quad \left. - \frac{2a}{d^2} \frac{\partial d}{\partial A}(C) \frac{\partial a}{\partial A}(H) + \frac{2a^2}{d^3} \frac{\partial d}{\partial A}(C) \frac{\partial d}{\partial A}(H) \right),
\end{aligned}$$

and using the fact that $\partial_A(A^{-1})(C) = -A^{-1}CA^{-1}$,

$$\begin{aligned}
\frac{\partial a}{\partial A}(C) &= 2 \operatorname{tr}(Ag^{-1}C^T), \\
\frac{\partial^2 a}{\partial A^2}(C, H) &= 2 \operatorname{tr}(Hg^{-1}C^T), \\
\frac{\partial d}{\partial A}(C) &= 2 \det^2(A) \det^{-1}(g) \operatorname{tr}(A^{-1}C), \\
\frac{\partial^2 d}{\partial A^2}(C, H) &= 2 \det^2(A) \det^{-1}(g) (2 \operatorname{tr}(A^{-1}C) \operatorname{tr}(A^{-1}H) - \operatorname{tr}(A^{-1}HA^{-1}C)).
\end{aligned}$$

Appendix B

Surface Matching

B.1 Shape Operators

We define the shape operator $S \in \mathbb{R}^{2,2}$ of the surface \mathcal{M} with respect to the parameterization x as the representation of the normal variation in the metric, *i.e.*

$$(g S v) \cdot w = Dn v \cdot D x w,$$

where n is the normal field on \mathcal{M} . Thus $S_A = g_A^{-1} (D x_A)^T D n_A$. For $S_B[\phi]$ at the deformed position $\phi_{\mathcal{M}}(x)$ with respect to the metric g_A , we obtain

$$\begin{aligned} (g_A S_B[\phi] v) \cdot w &= D(n_B \circ \phi) v \cdot D(x_B \circ \phi) w \\ &= (D n_B \circ \phi) D \phi v \cdot (D x_B \circ \phi) D \phi w \end{aligned}$$

and deduce from this

$$\begin{aligned} S_B[\phi] &= g_A^{-1} D \phi^T (D x_B \circ \phi)^T (D n_B \circ \phi) D \phi \\ &= g_A^{-1} D \phi^T (g_B \circ \phi) (S_B \circ \phi) D \phi, \end{aligned}$$

where S_B is the shape operator of the surface \mathcal{M}_B taking into account the original metric g_B . Up to a factor of $\frac{1}{2}$, the average change of the normal variation on the tangent space $\mathcal{T}_x\mathcal{M}_A$ is given by

$$\begin{aligned}\mathrm{tr}(S_B[\phi] - S_A) &= \mathrm{tr}(g_A^{-1} D\phi^T (g_B \circ \phi) (S_B \circ \phi) D\phi - S_A) \\ &= \mathrm{tr}((S_B \circ \phi) D\phi g_A^{-1} D\phi^T (g_B \circ \phi) - S_A) \\ &= \mathrm{tr}((S_B \circ \phi) D\phi \mathcal{G}[\phi] D\phi^{-1} - S_A)\end{aligned}$$

Let us neglect for the moment the impact of the metric distortion $\mathcal{G}[\phi]$ and suppose that $\phi_{\mathcal{M}}$ is an isometry and thus $\mathcal{G}[\phi] = \mathbb{I}$. In doing so, we have completely decoupled bending from tangential distortion. Finally, this gives $\mathrm{tr}(S_B \circ \phi) - \mathrm{tr}(S_A)$ as a local measure for bending.

B.2 Variations of the Matching Energy

Here we collect the different contributions of the matching energy in (2.16) and compute variations in directions ϑ . We deal with an energy

$$E[\phi] = \int_{\omega} \chi^\epsilon I d\xi$$

where $\chi^\epsilon = \chi_{\omega_A}^\epsilon \chi_{\omega_B}^\epsilon(\phi)$ is the regularized characteristic function and the actual energy integrand I splits into different contributions $I_{bend} + I_{reg} + I_{\mathcal{F}}$ resulting from the energies E_{reg} (2.7), E_{bend} (2.9) and $\tilde{E}_{\mathcal{F}}^\epsilon$ (2.11):

$$\begin{aligned}I_{bend} &:= (h_B(\phi) - h_A)^2 \sqrt{\det g_A} \\ I_{reg} &:= W(\mathrm{tr}\mathcal{G}[\phi], \det \mathcal{G}[\phi]) \sqrt{\det g_A} \\ I_{\mathcal{F}} &:= \eta^\epsilon(d_A)\theta^\epsilon(d_B(\phi))\sqrt{\det g_A} + \\ &\quad \eta^\epsilon(d_B(\phi))\theta^\epsilon(d_A)\sqrt{\det g_B(\phi)} \det D\phi.\end{aligned}$$

The variation of the energy in a direction ϑ is given by

$$\begin{aligned}E'[\phi](\vartheta) &:= \left. \frac{d}{d\epsilon} E[\phi + \epsilon\vartheta] \right|_{\epsilon=0} \\ &= \int_{\omega} (\partial_\phi \chi^\epsilon \cdot \vartheta) I + \chi^\epsilon (\partial_{D\phi} I : D\vartheta + \partial_\phi I \cdot \vartheta) d\xi,\end{aligned}$$

where “ \cdot ” indicates the Euclidean scalar product and “ $:$ ” is a scalar product on matrices with $A : B = \text{tr}(A^T B)$. Hence, the discrete L^2 energy gradient is the vector of these integrals over directional derivatives of the different integrands in directions of the finite element basis on the space of deformations.

Hence, the computation of these derivatives with respect to the arguments ϕ and $D\phi$ is a straightforward, albeit involved, application of the chain rule. In particular, for the variation of the regularized characteristic function we obtain

$$\partial_\phi \chi^\epsilon \cdot \vartheta = \chi_{\omega_A}^\epsilon \nabla \chi_{\omega_B}^\epsilon(\phi) \cdot \vartheta,$$

for the integrand of the bending energy (which does not depend on $D\phi$)

$$\partial_\phi I_{bend} \cdot \vartheta = 2(\nabla h_B(\phi) \cdot \vartheta) (h_B(\phi) - h_A) \sqrt{\det g_A},$$

and for the integrand of the regularization energy (2.8)

$$\begin{aligned} \partial_\phi I_{reg} \cdot \vartheta &= \left(\alpha_l \partial_\phi a \cdot \vartheta + \alpha_a \left(1 - \frac{\alpha_a - \alpha_l}{\alpha_a d^2}\right) \partial_\phi d \cdot \vartheta \right) \sqrt{\det g_A}, \\ \partial_{D\phi} I_{reg} : D\vartheta &= \left(\alpha_l \partial_{D\phi} a : D\vartheta + \alpha_a \left(1 - \frac{\alpha_a - \alpha_l}{\alpha_a d^2}\right) \partial_{D\phi} d : D\vartheta \right) \sqrt{\det g_A}, \end{aligned}$$

with

$$\begin{aligned} \partial_{D\phi} a : D\vartheta &= 2\text{tr}(g_A^{-1} D\phi^T g_B(\phi) D\vartheta), \\ \partial_\phi a \cdot \vartheta &= \text{tr}(g_A^{-1} D\phi^T (\nabla g_B(\phi) \cdot \vartheta) D\phi), \\ \partial_{D\phi} d : D\vartheta &= 2(\det g_A)^{-1} \det g_B(\phi) (\det D\phi)^2 \text{tr}(D\phi^{-1} D\vartheta), \\ \partial_\phi d \cdot \vartheta &= (\det g_A)^{-1} \det g_B(\phi) (\det D\phi)^2 \cdot \text{tr}(g_B^{-1}(\phi) (\nabla g_B(\phi) \cdot \vartheta)). \end{aligned}$$

Here we define $\nabla g_B(\phi) \cdot w \in \mathbb{R}^{2,2}$ by

$$(\nabla g_B(\phi) \cdot w)_{ij} = \sum_k \partial_k (g_B)_{ij}(\phi) w_k.$$

Finally, we obtain for the integrand of the feature energy

$$\begin{aligned}
\partial_\phi I_{\mathcal{F}} \cdot \vartheta &= \eta^\epsilon(d_A) (\theta^\epsilon)'(d_B(\phi)) (\nabla d_B(\phi) \cdot \vartheta) \sqrt{\det g_A} \\
&\quad + (\eta^\epsilon)'(d_B(\phi)) (\nabla d_B(\phi) \cdot \vartheta) \theta^\epsilon(d_A) \sqrt{\det g_B(\phi)} \det D\phi \\
&\quad + \eta^\epsilon(d_B(\phi)) \theta^\epsilon(d_A) \frac{\text{tr}(g_B^{-1}(\phi) (\nabla g_B(\phi) \cdot w))}{2\sqrt{\det g_B(\phi)}} \det D\phi, \\
\partial_{D\phi} I_{\mathcal{F}} : D\vartheta &= \eta^\epsilon(d_B(\phi)) \theta^\epsilon(d_A) \sqrt{\det g_B(\phi)} \det D\phi \text{tr}(D\phi^{-1} D\vartheta).
\end{aligned}$$

Bibliography

- [1] ALLEN, B., CURLESS, B., AND POPOVIĆ, Z. The space of human body shapes: reconstruction and parameterization from range scans. *ACM Transactions on Graphics* 22, 3 (July 2003), 587–594.
- [2] ALVAREZ, L., WEICKERT, J., AND SÁNCHEZ, J. A scale-space approach to nonlocal optical flow calculations. In *Second International Conference, Scale-Space '99* (1999), Lecture Notes in Computer Science; 1682, pp. 235–246.
- [3] BAJCSY, R., AND BROIT, C. Matching of deformed images. In *Proc. 6th Int. Joint Conf. Pattern Recogn.* (1982), pp. 351–353.
- [4] BALL, J. M. Convexity conditions and existence theorems in nonlinear elasticity. *Arch. Rat. Mech. Anal.* 63 (1977), 337–403.
- [5] BLANZ, V., AND VETTER, T. A morphable model for the synthesis of 3D faces. In *Proceedings of SIGGRAPH 99* (1999), Computer Graphics Proceedings, Annual Conference Series, pp. 187–194.
- [6] BORNEMANN, F., AND RASCH, C. Finite-element discretization of static Hamilton-Jacobi equations based on a local variational principle. *Preprint*, 2005.
- [7] BROWN, L. G. A survey of image registration techniques. *ACM Computing Surveys* 24, 4 (dec 1992), 325–376.
- [8] CHENG, X. On the nonlinear inexact Uzawa algorithm for saddle-point problems. *SIAM J. Numer. Anal.* 37, 6 (2000), 1930–1934.
- [9] CIARLET, P. G. *Three-Dimensional Elasticity*. Elsevier, New York, 1988.
- [10] CLARENZ, U., LITKE, N., AND RUMPF, M. Axioms and variational problems in surface parameterization. *Computer Aided Geometric Design* 21, 8 (2004), 727–749.

- [11] COHEN, I. Nonlinear variational method for optical flow computation. In *Proc. of the Eighth Scandinavian Conference on Image Analysis* (1993), pp. 523–530.
- [12] DACOROGNA, B. *Direct Methods in the Calculus of Variations*. Appl. Math. Sciences 78. Springer–Verlag, Berlin, 1989.
- [13] DACOROGNA, B., AND MOSER, J. On a partial differential equation involving the Jacobian determinant. *Ann. Inst. H. Poincaré Anal. Non Linéaire* 7, 1 (1990), 1–26.
- [14] DEGENER, P., MESETH, J., AND KLEIN, R. An adaptable surface parameterization method. In *Proc. 12th Int. Meshing Roundtable* (oct 2003), pp. 201–213.
- [15] DESBRUN, M., MEYER, M., AND ALLIEZ, P. Intrinsic parameterizations of surface meshes. In *Eurographics 2002* (2003), pp. 209–218.
- [16] DESBRUN, M., MEYER, M., SCHRÖDER, P., AND BARR, A. H. Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proceedings of SIGGRAPH 99* (Aug. 1999), Computer Graphics Proceedings, Annual Conference Series, pp. 317–324.
- [17] DROSKE, M., LITKE, N., OLISCHLÄGER, N., AND RUMPF, M. Variational methods for surface matching: Existence and qualitative properties. *In preparation*, 2005.
- [18] DROSKE, M., AND RUMPF, M. A variational approach to non-rigid morphological image matching. *SIAM Appl. Math.* 64, 2 (2004), 668–687.
- [19] FLOATER, M. S., AND HORMANN, K. Surface parameterization: a tutorial and survey. In *Advances in Multiresolution for Geometric Modelling*. Springer, 2005, pp. 157–186.
- [20] GRENANDER, U., AND MILLER, M. I. Computational anatomy: an emerging discipline. *Quarterly Appl. Math.* LVI, 4 (1998), 617–694.
- [21] GRINSPUN, E., HIRANI, A. N., DESBRUN, M., AND SCHRÖDER, P. Discrete shells. In *2003 ACM SIGGRAPH / Eurographics Symposium on Computer Animation* (Aug. 2003), pp. 62–67.
- [22] GU, X., GORTLER, S. J., AND HOPPE, H. Geometry images. *ACM Transactions on Graphics* 21, 3 (2002), 355–361.

- [23] GU, X., AND VEMURI, B. C. Matching 3D shapes using 2D conformal representations. In *MICCAI (1)* (2004), pp. 771–780.
- [24] GU, X., AND YAU, S.-T. Computing conformal structures of surfaces. *Communications in Informations and Systems* (2002), 121–146.
- [25] HORMANN, K., AND GREINER, G. MIPS: An efficient global parametrization method. In *Curve and Surface Design: Saint-Malo 1999* (2000), pp. 153–162.
- [26] KHAREVYCH, L., SPRINGBORN, B., AND SCHRÖDER, P. Discrete conformal mappings via circle patterns. *In preparation*, 2005.
- [27] KRAEVOY, V., AND SHEFFER, A. Cross-parameterization and compatible remeshing of 3D models. *ACM Transactions on Graphics* 23, 3 (Aug. 2004), 861–869.
- [28] LEE, A., DOBKIN, D., SWELDENS, W., AND SCHRÖDER, P. Multiresolution mesh morphing. In *Proceedings of SIGGRAPH 99* (Aug. 1999), Computer Graphics Proceedings, Annual Conference Series, pp. 343–350.
- [29] LÉVY, B., PETITJEAN, S., RAY, N., AND MAILLOT, J. Least squares conformal maps for automatic texture atlas generation. *ACM Transactions on Graphics* 21, 3 (2002), 362–371.
- [30] LITKE, N., CLARENZ, U., RUMPF, M., AND SCHRÖDER, P. A variational approach to optimal parameterization. *In preparation*, 2005.
- [31] LITKE, N., DROSKE, M., RUMPF, M., AND SCHRÖDER, P. An image processing approach to surface matching. *To appear in the Proceedings of the Symposium on Geometry Processing 2005*.
- [32] MARSDEN, J. E., AND HUGHES, T. J. R. *Mathematical Foundations of Elasticity*. Prentice–Hall, Englewood Cliffs, 1983.
- [33] MERCAT, C. Discrete Riemann surfaces and the Ising model. *Communications in Mathematical Physics* 218, 1 (2001), 177–216.
- [34] MEYER, M., DESBRUN, M., SCHRÖDER, P., AND BARR, A. H. Discrete differential-geometry operators for triangulated 2-manifolds. *Vis. Math.* (2002).

- [35] MORREY, C. Quasi-convexity and lower semicontinuity of multiple integrals. *Pac. J. Math.* 2 (1952), 25–53.
- [36] MORREY, C. *Multiple Integrals in the Calculus of Variations*. Springer-Verlag, New York, 1966.
- [37] MOSER, J. On the volume elements on a manifold. *Trans. Amer. Math. Soc.* 120 (1965), 286–294.
- [38] NOCEDAL, J., AND WRIGHT, S. J. *Numerical Optimization*. Springer Series in Operations Research. Springer-Verlag, New York, 1999.
- [39] OGDEN, R. W. *Non-Linear Elastic Deformations*. John Wiley, 1984.
- [40] OSHER, S. J., AND SETHIAN, J. A. Fronts propagating with curvature dependent speed: Algorithms based on Hamilton–Jacobi formulations. *J. of Comp. Physics* 79 (1988), 12–49.
- [41] PARKE, F. I. Computer generated animation of faces. In *ACM'72: Proceedings of the ACM annual conference* (1972), pp. 451–457.
- [42] PARKE, F. I. *A parametric model for human faces*. PhD thesis, University of Utah, Salt Lake City, Utah, 1974.
- [43] PRAUN, E., SWELDENS, W., AND SCHRÖDER, P. Consistent mesh parameterizations. In *Proceedings of ACM SIGGRAPH 2001* (Aug. 2001), Computer Graphics Proceedings, Annual Conference Series, pp. 179–184.
- [44] SANDER, P. V., SNYDER, J., GORTLER, S. J., AND HOPPE, H. Texture mapping progressive meshes. *Proceedings of SIGGRAPH 2001* (2001), 409–416.
- [45] SCHREINER, J., ASIRVATHAM, A., PRAUN, E., AND HOPPE, H. Inter-surface mapping. *ACM Transactions on Graphics* 23, 3 (Aug. 2004), 870–877.
- [46] SHEFFER, A., AND DE STURLER, E. Parameterization of faceted surfaces for meshing using angle based flattening. *Engineering with Computers* 17, 3 (2001), 326–337.
- [47] SORKINE, O., COHEN-OR, D., GOLDENTHAL, R., AND LISCHINSKI, D. Bounded-distortion piecewise mesh parameterization. In *IEEE Visualization* (2002), pp. 355–362.

- [48] VAN DEN ELSEN, P. A., POL, E.-J. J., AND VIERGEVER, M. A. Medical image matching: A review with classification. *IEEE Engineering in Medicine and Biology* 12 (1993), 26–39.
- [49] WOLBERG, G. *Digital Image Warping*. IEEE Computer Society Press, Los Alamitos, CA, 1990.
- [50] ZHANG, L., SNAVELY, N., CURLESS, B., AND SEITZ, S. M. Spacetime faces: high resolution capture for modeling and animation. *ACM Transactions on Graphics* 23, 3 (Aug. 2004), 548–558.