

An Image Processing Approach to Surface Matching

Nathan Litke
Caltech

Marc Droske
University of Bonn

Martin Rumpf
University of Bonn

Peter Schröder
Caltech

Abstract

Establishing a correspondence between two surfaces is a basic ingredient in many geometry processing applications. Existing approaches, which attempt to match two meshes directly in 3D, can be cumbersome to implement and it is often hard to produce accurate results in a reasonable amount of time. In this paper, we present a new variational method for matching surfaces that addresses these issues. Instead of matching two surfaces directly in 3D, we apply well-established matching methods from image processing in the parameter domains of the surfaces. A matching energy is introduced that can depend on curvature, feature demarcations or surface textures, and a regularization energy controls length and area changes in the induced non-rigid deformation between the two surfaces. The metric on both surfaces is properly incorporated into the formulation of the energy. This approach reduces all computations to the 2D setting while accounting for the original geometries. Consequently a fast multiresolution numerical algorithm for regular image grids can be used to solve the global optimization problem. The final algorithm is robust, generically much simpler than direct matching methods, and very fast for highly resolved triangle meshes.

Categories and Subject Descriptors (according to ACM CCS): G.1.8 [Numerical Analysis]: Elliptic equations; Finite element methods. I.3.5 [Computer Graphics]: Curve, surface, solid and object representations; Geometric algorithms, languages, and systems; Physically based modeling.

Keywords: Surface matching, deformation energy, non-linear elasticity, digital geometry processing.

1. Introduction

A correspondence between two surfaces is a function that maps one surface onto the other. The need for a correspondence can be found in numerous geometry processing applications, for example:

- ◇ **Data fitting:** Fitting a canonical surface model to triangle mesh data from a 3D scanning system (possibly with dropouts);
- ◇ **Statistical analysis:** Bringing a corpus of geometric models into an appropriate common parameterization to apply statistical tools such as principal component analysis;
- ◇ **Comparison and quality assurance:** Comparing a scan of a physical object with a CAD description;
- ◇ **Attribute transfer:** Mapping displacements or textures (among many other possible examples) between surfaces.

Existing methods to establish such correspondences can be very expensive computationally or lack the high accuracy needed when very detailed matches are desired. Because of the many local minima in the energy landscape which expresses the relationship between the surfaces, extensive guidance from a user is often needed to produce reasonable results. In this paper, we propose a new method for establishing a correspondence between two surfaces which addresses these limitations. Our approach draws upon the extensive and mature body of work in global *image* matching (see the brief review below). To do so we map appropriate geometric attributes (metric, mean curvature, textures, *etc.*) into the parameter plane, induced by some smooth parameterization of the surfaces (cf. Fig-

ure 1). The standard image matching energies are then appropriately modified to correctly account for the geometry of the original surfaces. What remains is an energy minimization problem in 2D which can then be solved effectively even for highly detailed meshes through well established multiscale methods. Detailed control over the match can be achieved through additional feature energies as desired.

1.1. Previous Work

The body of relevant literature is quite broad and we will not attempt to give a complete account of it here. Most relevant for our setting is the 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 .

Image Matching In image processing, registration is often approached as a variational problem. One asks for a deformation ϕ 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

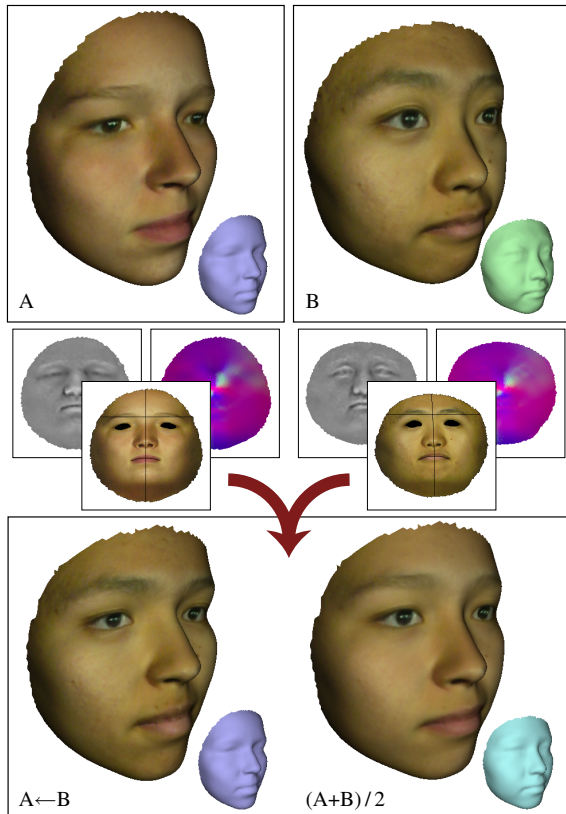


Figure 1: A brief overview of our surface matching process: The two surfaces to be matched are parameterized (top); we generate images of mean curvature, the metric and user-defined feature sets on a 256×256 grid (middle); an optimal matching deformation is then found through a global energy relaxation 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 minutes.

the infinite dimensional space of deformations [Bro92, vdEPV93]. 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 [BB82] and more recent, significant extensions by Grenander and Miller [GM98]. 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 [Coh93] considered polyconvex elastic functionals and Droske and Rumpf [DR04] 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 [LDSS99]. Typically these methods are driven by user-supplied feature correspondences which are then used to drive a mutual parameterization. The main difficulty is the proper alignment of selected features during the parameterization process [KS04, PSS01, SAPH04] 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 [BV99] used cylindrical scans resulting in height “images” which were matched through a modified optical flow. Allen *et al.* [ACP03] 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.* [ZSCS04] 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 [GV04] 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

In this paper, we present a new method for matching surfaces with the following characteristics:

- ◊ We develop a variational approach based on minimizing bending and stretching in the matching deformation (cf. Figure 13).
- ◊ We provide user control over the match through feature lines which are mapped as sets onto corresponding feature lines rather than through point-wise constraints (cf. Figure 5).
- ◊ Our method decouples the discretization of the surfaces from the discretization of the matching deformation.
- ◊ We allow for a partial correspondence of the surfaces, *i.e.*, regions of the surfaces — in particular boundary regions — are not required to be in correspondence with regions on the other surface (cf. Figure 10).
- ◊ Existence and global injectivity of the matching deformations is established, such that the resulting deformations are smooth and bijective.

The reliability and robustness of our method is ensured by a multiresolution strategy. The algorithm proceeds from a coarse scale matching of the overall shapes to a fine scale identification of all surface details. Consequently most iterations of the algorithm are spent on very coarse matches which are solved efficiently.

Currently, a limitation of the method is that the surfaces to be matched must be homeomorphic to a disc such that we can build a single parameterization in the plane. Also, our energy formulation is not symmetric: matching surface \mathcal{M}_A to surface \mathcal{M}_B might give a different correspondence than matching \mathcal{M}_B to \mathcal{M}_A . Possible ways to overcome these limitations are discussed in Section 4.

Organization The paper is organized as follows. In the next section we introduce a variational formulation of the surface matching problem. We derive the basic components of the energy and detail how the geometry of the original meshes enters into the formulation in the image domain. Section 3 describes the matching algorithm in detail with the choice of the initial parameterizations, the user-guided selection of feature sets, the treatment of triangle meshes and the multiscale finite element method. Section 4 demonstrates the performance of our method with a variety of examples where high quality matches are critical. The paper concludes with a comparison to previous approaches and a discussion of possible improvements of the method.

2. Surface Matching

Our goal is to correlate two surface patches, \mathcal{M}_A and \mathcal{M}_B , through 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 . In doing so, we want to avoid the general difficulty of formulating these maps directly in \mathbb{R}^3 and the particularly tedious algorithmic issues in the application, where the two surface patches are given as distinct, arbitrary triangulations. Instead, we match parameter domains covered with geometric and user-defined feature characteristics. 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 for images 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 [GM98, DR04], 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 here, 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.

A Physical Interpretation of the Surface Deformation Consider the first surface to be a thin shell which we press into a mould of the second surface (cf. Figure 2). One can distinguish between stresses induced by stretching and compression, and stressed 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 Eq. (1) will measure the induced in-plane stresses, and the concrete energy

density in Eq. (2) 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 Eq. (3). Finally, the matching of feature sets in Eq. (4) will provide user-specified landmarks to guide the surface deformation. In what follows, we will develop the variational approach step-by-step.




Figure 2: A physical interpretation of $\phi_{\mathcal{M}}$ 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 $\phi_{\mathcal{M}}$.

2.1. Measuring Distortion via a Parameterization

To begin with, let us set up proper parameterizations. 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 suppose $x : \omega \rightarrow \mathcal{M}$ is a parameterization of \mathcal{M} on a parameter domain ω . 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 a bijection and a metric g is properly defined on ω ,

$$g = Dx^T Dx$$

where $Dx \in \mathbb{R}^{3,2}$ is the Jacobian of the parameterization x . The metric g acts on tangent vectors v, w on the parameter domain 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 the surface. Thus it follows that the metric describes how length, area and angles are distorted under the parameterization function.

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. 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.

2.2. Measuring Distortion via a Deformation

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 $\phi_{\mathcal{M}} : \mathcal{M}_A \rightarrow \mathcal{M}_B$ defined by

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

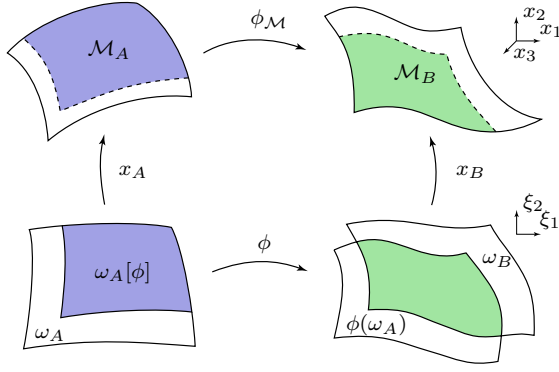


Figure 3: The matching function $\phi_{\mathcal{M}} := x_B \circ \phi \circ x_A^{-1}$ is a mapping between the corresponding shaded regions of the two surfaces. The partial correspondence is defined on $\omega_A[\phi] := \omega_A \cap \phi^{-1}(\omega_B)$.

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 3.

Now let us focus on the distortion from the surface \mathcal{M}_A onto the surface \mathcal{M}_B . In elasticity, the 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 the previous section. 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 where products are given 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 deformation,

$$E_{\text{reg}}[\phi] = \int_{\omega_A} W(\text{tr } \mathcal{G}[\phi], \det \mathcal{G}[\phi]) \sqrt{\det g_A} d\xi. \quad (1)$$

This simple class of energy functionals was rigorously derived in [CLR04] from a set of natural axioms for measuring the distortion of a single parameterization. In particular, the following energy density

$$W(a, d) = \alpha_l a + \alpha_a (d + (1 + \frac{\alpha_l}{\alpha_a}) d^{-1}) \quad (2)$$

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} . The weights $\alpha_l, \alpha_a > 0$ are typically chosen by the user according to the relative importance of length and area distortion.

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 [DLOR05], 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 by [GHDS03] 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

$$E_{\text{bend}}[\phi] = \int_{\omega_A} (h_B \circ \phi - h_A)^2 \sqrt{\det g_A} d\xi \quad (3)$$

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

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 [Cia88]. 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,

$$E_{\mathcal{F}}[\phi] = \mathcal{H}^1(\mathcal{F}_{\mathcal{M}_A} \setminus \phi_{\mathcal{M}}^{-1}(\mathcal{F}_{\mathcal{M}_B})) + \mathcal{H}^1(\mathcal{F}_{\mathcal{M}_B} \setminus \phi_{\mathcal{M}}(\mathcal{F}_{\mathcal{M}_A})) \quad (4)$$

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 Eq. (4) that involves the *distance on the surface to the feature sets*, and define

$$\begin{aligned} \tilde{E}_{\mathcal{F}}^{\epsilon}[\phi] = & \int_{\omega_A} (\eta^{\epsilon} \circ d_A(\xi)) (\theta^{\epsilon} \circ d_B(\phi(\xi))) \sqrt{\det g_A} d\xi + \\ & \int_{\omega_B} (\eta^{\epsilon} \circ d_B(\xi)) (\theta^{\epsilon} \circ d_A(\phi^{-1}(\xi))) \sqrt{\det g_B} d\xi \end{aligned} \quad (5)$$

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 Eq. (5) as

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

2.5. 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 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):

$$E_{bend}[\phi] = \int_{\omega_A[\phi]} (h_B \circ \phi - h_A)^2 \sqrt{\det g_A} \, d\xi, \quad (6)$$

$$E_{\mathcal{F}}[\phi] = \mathcal{H}^1(\omega_A[\phi] \cap \mathcal{F}_{\mathcal{M}_A} \setminus \phi_{\mathcal{M}}^{-1}(\mathcal{F}_{\mathcal{M}_B})) + \mathcal{H}^1(\mathcal{F}_{\mathcal{M}_B} \setminus \phi_{\mathcal{M}}(\omega_A[\phi] \cap \mathcal{F}_{\mathcal{M}_A})) \quad (7)$$

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

$$\chi_{\mathcal{A}}^\epsilon(\xi) = \max(1 - \epsilon^{-1} \text{dist}(\xi, \mathcal{A}), 0) \quad (8)$$

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} (3) and the already regularized energy $\tilde{E}_{\mathcal{F}}^\epsilon$ (5) and denote the resulting energies by E_{bend}^ϵ and $E_{\mathcal{F}}^\epsilon$, respectively.

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 energies $E_{bend}^\epsilon, E_{reg}$ and $E_{\mathcal{F}}^\epsilon$, respectively, and define the global energy

$$E^\epsilon[\phi] = \beta_{bend} E_{bend}^\epsilon[\phi] + \beta_{reg} E_{reg}[\phi] + \beta_{\mathcal{F}} E_{\mathcal{F}}^\epsilon[\phi] \quad (9)$$

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 (1), (6) and (7):

$$E[\phi] = \beta_{bend} E_{bend}[\phi] + \beta_{reg} E_{reg}[\phi] + \beta_{\mathcal{F}} E_{\mathcal{F}}[\phi] \quad (10)$$

A matching deformation ϕ that minimizes Eq. (10) for a given set of parameters $\beta = (\beta_{bend}, \beta_{reg}, \beta_{\mathcal{F}})$ is called β -optimal (or simply described as “optimal” in what follows). In [DLOR05], we provide a *proof of existence, global injectivity, and regularity of β -optimal matching deformations*. Because of this proof, we can expect to obtain smooth deformations that are free of folds and singularities. The next section describes our surface matching procedure for finding optimal matches between triangulated surface patches.

3. The Matching Algorithm

So far we have 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 (cf. Figure 4). 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 [GGH02]).

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 Eq. (9).

3.1. Surface Patch Parameterization

We are interested in low distortion parameterizations to ensure 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. These considerations favor parameterization methods which support natural boundary conditions (see [FH05] for a comprehensive survey). We use the method of Clarenz *et al.* [CLR04] since it simultaneously optimizes the parameterization for low angle, area and length distortion (cf. Figure 4).

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 Eq. (5) and the hierarchical nature of the minimization algorithm this rough alignment is entirely sufficient.




Figure 4: A surface triangulation \mathcal{M} and its parameter domain ω . Distortion in the metric g is depicted by a density plot on the right, measured by Eq. (2) with $a = \text{tr}(g^{-1})$, $d = \det(g^{-1})$, $\frac{\alpha_L}{\alpha_a} = 100$.

3.2. Feature Set Construction

The user can control the match by identifying sets of similar features on the surface patches. The cost functional $E_{\mathcal{F}}^{\epsilon}[\cdot]$ from Eq. (5) 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 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 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 5b).
- ◇ Lines of symmetry drawn as features allow deformations tangential to the feature boundary, but prevent deformations that are transverse to it (Figure 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

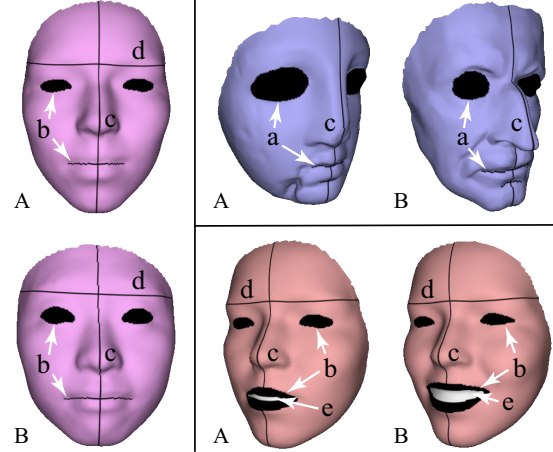


Figure 5: Examples of user-defined feature sets: (a) coarse registration of geometric features; (b) aligning texture features; (c) lines of symmetry; (d) preventing smooth, rigid regions from sliding; and (e) increasing the elasticity of highly deformable regions.

deformable face versus the forehead, Figure 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 5e). We achieve this by decreasing β_{reg} in the selected region.

The distance maps used in the definition of the feature energy are discretized by an upwind scheme for the corresponding Eikonal equations [OS88] (cf. Figure 6c). We used the particular upwind finite element algorithm of Bornemann [BR05] 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.

3.3. 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 Eq. (8). 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 (cf. Figure 6).

For h_A we use the magnitude of the mean curvature normal [DMSB99, MDSB02]. The sign is chosen according to the surface normal, which we take to be a weighted sum of triangle normals around a vertex. Other measurements for the mean curvature would work equally well (see [MDSB02] 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 [SSGH01]). The distance map for \mathcal{F}_A is described above and illustrated in Figure 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 (cf. Figure 6d).

3.4. Multiscale Finite Element Formulation

The total energy $E[\cdot]$ is highly non-linear. In particular the bending energy causes many local minima in the energy landscape over




Figure 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; (d) the distance map for the characteristic function χ_A^ϵ of the domain, with the boundary $\partial\omega_A$ in blue.

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 [AWS99], allowing for a robust and efficient global minimization on complicated energy landscapes (Figure 7).

```

// 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 7: Pseudo-code for the multiscale algorithm.

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 [NW99] (Figure 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 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 $k+1$. The prolongation of ϕ^k to ϕ^{k+1} is performed using bi-linear interpolation (Figure 7, **prolongate**).

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 7, **filter**). A Gaussian filter with respect to the surface metric of width σ_k is used to define the smoothing on the surfaces \mathcal{M}_A and \mathcal{M}_B . This is in fact an approximation for the mean curvature flow on the surface [DMSB99]. 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 space-varying filter kernels to the corresponding h_A and h_B images. Figure 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.




Figure 8: The mean curvature function h_A is extended to the full image domain ω and successively restricted to coarser grids from the multigrid pyramid.

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 $k = 1, \dots, m$. For the parameter domains discretized with a 256×256 grid in our examples, we have a maximal number of $m = 8$ scales. Figure 9 shows a representative energy decay during the global minimization algorithm, which exhibits a characteristic staircase-like behavior at the change in levels.

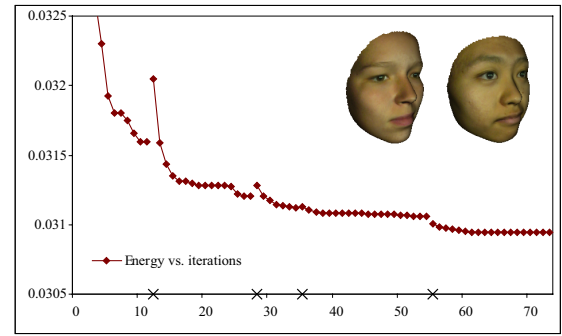


Figure 9: Plot of the energy decay for the sequence of deformations computed in the minimization algorithm on multiple scales. The staircase-like behavior results from the jumps onto the next finer scale (marked “x”), which may increase the energy temporarily.

Finally all these discretized quantities are used in the computation of the appropriate integrals in the finite element method. A discrete deformation $\phi^k : \omega \rightarrow \mathbb{R}^2$ on level k is a vector-valued function, whose components are piecewise affine, continuous functions. We represent ϕ^k with an array of vectors in \mathbb{R}^2 at the nodes of the image grid, and use bi-linear interpolation to evaluate ϕ^k at arbitrary points $\xi \in \omega$. For these deformations we have to evaluate the energy functional $E^\epsilon[\cdot]$. The integrals are evaluated over each pixel of the image grid using the filtered images at that level. Any pixel-wise integral $\int_C f d\xi$ is evaluated using a nodal quadrature rule. Local folds are detected when $\det(\phi^k) \leq 0$ in the energy computation, and these inadmissible deformations are explicitly disallowed in the minimization algorithm. Besides evaluating the discrete energy for a given deformation ϕ^k , we have to compute the gradient of this energy in the conjugate gradient method, which requires the differentiation of the discrete energy with respect to the discrete deformation. All the necessary expressions are provided in the Appendix.

Once we have computed the discrete deformation ϕ , the discrete surface matching deformation is given by

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

which is defined on $x_A(\omega_A[\phi])$. We can now apply $\phi_{\mathcal{M}}$ directly to the triangulated surface \mathcal{M}_A through its texture coordinates. For

example, to morph between the two surfaces, we need to assign a 3D displacement vector to each vertex in \mathcal{M}_A (cf. Figure 13). Given the texture coordinates of a vertex in \mathcal{M}_A , we evaluate ϕ (using bi-linear interpolation) in $\omega_A[\phi]$ to get a 2D displacement vector which gives us a texture coordinate with respect to ω_B . This is then used to determine the corresponding point on \mathcal{M}_B . If we wish to transfer a texture from \mathcal{M}_B to \mathcal{M}_A (cf. Figure 1), we pass over all the pixels in $\omega_A[\phi]$ and evaluate ϕ to find the corresponding pixel color in ω_B . The result is the appropriately deformed texture from \mathcal{M}_B on \mathcal{M}_A .




Figure 10: 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.

4. Results and Conclusions

Figures 1, 10, 12 and 13 show examples of surface matching. Since matching lies at the heart of many applications, the aim here is to demonstrate the quality and robustness of our results rather than focus on a particular set of applications. The accompanying animations put these matches in motion by morphing between the pairs of surfaces. Certain features cannot have exact correspondences, e.g., a mole on one face in Figure 1, and the crack in the model in Figure 10. However, our results produce faithful correspondences between the geometric and texture features which the models have in common.

Figure 1 shows a match between two similar surfaces that exhibit a two-fold symmetry and possess texture features that should be matched exactly (i.e., the eyes and mouth). The accuracy of the match $\phi_{\mathcal{M}}$, especially around these features, is shown by transferring the texture of the second surface onto the first with $\phi_{\mathcal{M}}^{-1}$ and blending the surfaces with $x \leftarrow \frac{1}{2}x + \frac{1}{2}\phi_{\mathcal{M}}(x)$, $x \in \mathcal{M}_A$, respectively.

Our algorithm is also good at matching surfaces with gross differences and mismatched features, as in Figures 10 and 13. This often requires very large deformations and typically only partial matches are possible. We find that using coarse features as hints and relying on the bending energy for accuracy at higher resolutions is a good strategy in this case.

Figure 12 shows an example of facial animation. Here the goal is

\mathcal{M}_A	\mathcal{M}_B	V_A	V_B	Iter.	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 1: Performance figures for the matching examples (V_A, V_B denote # vertices for the triangle meshes $\mathcal{M}_A, \mathcal{M}_B$ respectively).

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.

Table 1 shows the matching pairs, the sizes of the triangle meshes and the image grid, and the computation times for the energy minimization. Performance figures were measured on a 3.6 GHz Intel Pentium IV Xeon processor. The time to generate the geometry morphs and texture transfers in the examples is negligible. We find that the majority of time is spent in the alignment of surface curvatures at finer scales. In an extreme case shown in Figure 13, the large dissimilarities between the surfaces actually contributes very little at finer scales, and thus the optimization is quite fast. For all our examples, the parameters $\beta_{bend} = 1, \beta_{reg} = 0.01, \beta_{\mathcal{F}} = 5$ gave the most aesthetically pleasing results.

4.1. Summary

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 described in Section 1.1 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 [LDSS99, PSS01] or involves a difficult algorithm to get a good partitioning [KS04, SAPH04]. The combinatoric complexity of the meshes 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 these methods cannot easily consider the alignment of surface curvatures and other attributes, for example from texture maps. Our approach 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 [ACP03, ZSCS04]. 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 13) and different poses (cf. Figure 12), 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 11). It is well known in image matching [GM98] 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.




Figure 11: 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. Future Work

One of our goals in future research is to extend this method to matching surfaces of arbitrary (but identical) topological type. Generally speaking, this can be done without much difficulty whenever a common parameter domain can be found for both surfaces. Our matching energies remain unchanged, except that the metrics are now taken with respect to the new parameter domain. For example, we can take this approach to treat closed surfaces such as spheres, for which parameterization algorithms exist [GGS03, PH03]. For arbitrary genus shapes, it would be possible to apply our method to an automatically generated atlas by adding inter-chart continuity conditions to the matching energy, similar to [KLS03], or through a global conformal parameterization of the surface [GY02, KSS05]. Furthermore, the lack of symmetry in our variational formulation can be overcome by simultaneously dealing with both deformations — from \mathcal{M}_A to \mathcal{M}_B and its inverse from \mathcal{M}_B to \mathcal{M}_A — with a constraint on these two independent transformations.

Currently surface morphing is implemented based on a linear blend between the identity and the actual matching deformation $\phi_{\mathcal{M}}$. With a proper metric on the space of deformations at hand, there may be more natural geodesic paths along which to perform the actual morphing (cf. the corresponding work in imaging [CRM96, CHHY00]).

Finally, the application of image processing methodology is not limited to surface matching problems. Other variational problems in modeling and manipulating geometry might benefit from this approach as well. Our method provides a template for a general embedding of such problems in the variational framework of image processing.

Acknowledgements This work was supported in part by NSF (DMS-0220905, DMS-0138458, ACI-0219979), DFG (SPP 1114), DOE (W-7405-ENG-48/B341492), the Packard Foundation, nVidia, Microsoft, the Center for Integrated Multiscale Modeling and Simulation, Alias and Pixar. Special thanks to Sharif Elcott, Ilja Friedel, Liliya Kharevych and Weiwei Yang, and to our reviewers for their helpful comments.

References

- [ACP03] ALLEN B., CURLESS B., POPOVIĆ Z.: The space of human body shapes: reconstruction and parameterization from range scans. *ACM Transactions on Graphics* 22, 3 (July 2003), 587–594.
- [AWS99] ALVAREZ L., WEICKERT J., SÁNCHEZ J.: A scale-space approach to nonlocal optical flow calculations. In *Second International Conference, Scale-Space '99* (sep 1999), Lecture Notes in Computer Science; 1682, pp. 235–246.
- [BB82] BAJCSY R., BROIT C.: Matching of deformed images. In *Proc. 6th Int. Joint Conf. Patt. Recogn.* (1982), pp. 351–353.
- [BR05] BORNEMANN F., RASCH C.: Finite-Element Discretization of Static Hamilton-Jacobi Equations Based on a Local Variational Principle. *Preprint*, 2005.
- [Bro92] BROWN L. G.: A survey of image registration techniques. *ACM Computing Surveys* 24, 4 (dec 1992), 325–376.
- [BV99] BLANZ V., 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.
- [CHHY00] COHEN I., HERLIN I., HUOT E., YAHIA H.: Matching structures by computing minimal paths on a manifold. *Journal of Visual Communication and Image Representation* 13 (2000), 302–312.
- [Cia88] CIARLET P. G.: *Three-Dimensional Elasticity*. Elsevier, New York, 1988.
- [CLR04] CLARENZ U., LITKE N., RUMPF M.: Axioms and variational problems in surface parameterization. *Computer Aided Geometric Design* 21, 8 (2004), 727–749.
- [Coh93] COHEN I.: Nonlinear variational method for optical flow computation. In *Proc. of the Eighth Scandinavian Conference on Image Analysis* (1993), pp. 523–530.
- [CRM96] CHRISTENSEN G. E., RABBITT R. D., MILLER M. I.: Deformable templates using large deformation kinematics. *IEEE Trans. Medical Imaging* 5, no. 10 (1996), 1435–1447.
- [DLOR05] DROSKE M., LITKE N., OLISCHLÄGER N., RUMPF M.: Variational methods for surface matching: Existence and qualitative properties. *In preparation*, 2005.
- [DMSB99] DESBRUN M., MEYER M., SCHRÖDER P., 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.
- [DR04] DROSKE M., RUMPF M.: A variational approach to non-rigid morphological image matching. *SIAM Appl. Math.* 64, 2 (2004), 668–687.
- [FH05] FLOATER M. S., HORMANN K.: Surface parameterization: a tutorial and survey. In *Advances in Multiresolution for Geometric Modelling*. Springer, 2005, pp. 157–186.
- [GGH02] GU X., GORTLER S. J., HOPPE H.: Geometry Images. *ACM Transactions on Graphics* 21, 3 (2002), 355–361.
- [GGS03] GOTSMAN C., GU X., SHEFFER A.: Fundamentals of Spherical Parameterization for 3D Meshes. *ACM Transactions on Graphics* 22, 3 (July 2003), 358–363.

- [GHDS03] GRINSPUN E., HIRANI A. N., DESBRUN M., SCHRÖDER P.: Discrete shells. In *2003 ACM SIGGRAPH / Eurographics Symposium on Computer Animation* (Aug. 2003), pp. 62–67.
- [GM98] GRENANDER U., MILLER M. I.: Computational anatomy: An emerging discipline. *Quarterly Appl. Math.* *LVI*, no. 4 (1998), 617–694.
- [GV04] GU X., VEMURI B. C.: Matching 3D shapes using 2D conformal representations. In *MICCAI (1)* (2004), pp. 771–780.
- [GY02] GU X., YAU S.-T.: Computing conformal structures of surfaces. *Communications in Informations and Systems* (2002), 121–146.
- [KLS03] KHODAKOVSKY A., LITKE N., SCHRÖDER P.: Globally smooth parameterizations with low distortion. *ACM Transactions on Graphics* *22*, 3 (July 2003), 350–357.
- [KS04] KRAEVOY V., SHEFFER A.: Cross-parameterization and compatible remeshing of 3D models. *ACM Transactions on Graphics* *23*, 3 (Aug. 2004), 861–869.
- [KSS05] KHAREVYCH L., SPRINGBORN B., SCHRÖDER P.: Discrete conformal mappings via circle patterns. *In preparation*, 2005.
- [LDSS99] LEE A., DOBKIN D., SWELDENS W., SCHRÖDER P.: Multiresolution mesh morphing. In *Proceedings of SIGGRAPH 99* (Aug. 1999), Computer Graphics Proceedings, Annual Conference Series, pp. 343–350.
- [MDSB02] MEYER M., DESBRUN M., SCHRÖDER P., BARR A. H.: Discrete differential-geometry operators for triangulated 2-manifolds. *Vis. Math.* (2002).
- [NW99] NOCEDAL J., WRIGHT S. J.: *Numerical Optimization*. Springer Series in Operations Research. Springer-Verlag, New York, 1999.
- [OS88] OSHER S. J., SETHIAN J. A.: Fronts propagating with curvature dependent speed: Algorithms based on Hamilton–Jacobi formulations. *J. of Comp. Physics* *79* (1988), 12–49.
- [PH03] PRAUN E., HOPPE H.: Spherical Parametrization and Remeshing. *ACM Transactions on Graphics* *22*, 3 (July 2003), 340–349.
- [PSS01] PRAUN E., SWELDENS W., SCHRÖDER P.: Consistent mesh parameterizations. In *Proceedings of ACM SIGGRAPH 2001* (Aug. 2001), Computer Graphics Proceedings, Annual Conference Series, pp. 179–184.
- [SAPH04] SCHREINER J., ASIRVATHAM A., PRAUN E., HOPPE H.: Inter-surface mapping. *ACM Transactions on Graphics* *23*, 3 (Aug. 2004), 870–877.
- [SSGH01] SANDER P. V., SNYDER J., GORTLER S. J., HOPPE H.: Texture mapping progressive meshes. *Proceedings of SIGGRAPH 2001* (2001), 409–416.
- [vdEPV93] VAN DEN ELSEN P. A., POL E.-J. J., VIERGEVER M. A.: Medical image matching: A review with classification. *IEEE Engineering in Medicine and Biology* *12* (1993), 26–39.
- [ZSCS04] ZHANG L., SNAVELY N., CURLLESS B., SEITZ S. M.: Spacetime faces: high resolution capture for modeling and animation. *ACM Transactions on Graphics* *23*, 3 (Aug. 2004), 548–558.

Appendix: Variations of the Matching Energy

An implementation of the formulas presented here is available from the authors upon request. Let us collect the different contributions of the matching energy in Eq. (9) and compute variations in directions ϑ . We deal with an energy

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

where $\chi^{\epsilon} = \chi_{\omega_A}^{\epsilon} \chi_{\omega_B}^{\epsilon}(\phi)$ is the regularized characteristic function. The energy integrand I splits into different contributions $I_{bend} + I_{\mathcal{F}}$ resulting from the energies E_{bend}^{ϵ} (3) and $\tilde{E}_{\mathcal{F}}^{\epsilon}$ (5) and I_{reg} corresponds to the energy E_{reg} (1):

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

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

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

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.

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, we obtain for the variation of the regularized characteristic function

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

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

$$\begin{aligned} \partial_{\phi} I_{reg} \cdot \vartheta &= \left(\alpha_1 \partial_{\phi} a \cdot \vartheta + \alpha_a \left(1 - \frac{\alpha_a - \alpha_l}{\alpha_a d^2} \right) \partial_{\phi} d \cdot \vartheta \right) \sqrt{\det g_A}, \\ \partial_{D\phi} I_{reg} : D\vartheta &= \left(\alpha_1 \partial_{D\phi} a : D\vartheta + \right. \\ &\quad \left. \alpha_a \left(1 - \frac{\alpha_a - \alpha_l}{\alpha_a d^2} \right) \partial_{D\phi} d : D\vartheta \right) \sqrt{\det g_A}, \end{aligned}$$

with

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

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

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

$$\partial_{D\phi} I_{\mathcal{F}} : D\vartheta = \eta^{\epsilon}(d_B(\phi)) \theta^{\epsilon}(d_A) \sqrt{\det g_B(\phi)} \det D\phi \text{tr}(D\phi^{-1} D\vartheta).$$




Figure 12: 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 skin 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 13: 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).